

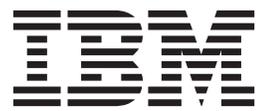
IBM Tivoli Storage Manager
Version 6.3.3

Problem Determination Guide



IBM Tivoli Storage Manager
Version 6.3.3

Problem Determination Guide



Note:

Before using this information and the product it supports, read the information in "Notices" on page 271.

This edition applies to Version 6.3.3 of IBM Tivoli Storage Manager (product number 5608-E01, 5608-E02, 5608-E03, and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters. This edition replaces GC23-9789-03.

© **Copyright IBM Corporation 1993, 2012.**

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

Contents

About this publication vii

Who should read this guide.	vii
Publications	vii
Tivoli Storage Manager publications	viii
Tivoli Storage FlashCopy Manager publications.	x
Support information.	x
Getting technical training	xi
Searching knowledge bases	xi
Contacting IBM Software Support	xiii

Chapter 1. Help facilities 1

Backup-archive client help.	1
Accessing help for the Windows service configuration utility (dsmcutil)	2
Server or storage agent help	2
Accessing server or storage agent help for commands	2
Accessing help for messages	3
Command-line interface help for the client	3
Reporting a problem with a help topic.	3

Chapter 2. Resolving Tivoli Storage Manager client problems 5

Examining error messages	5
Examining the server activity log messages	5
Identifying when and where the problem can occur	5
Reproducing the problem	6
Collecting documentation to resolve problems with the client application	6
Determining why the dsmdc or dsmdm or dsmdj program does not start	7
Resolving problems with client option sets	9
Resolving problems with client option sets	9
Resolving password authentication problems	10
Resolving LDAP-authenticated password problems	11
Verifying the configuration for password authentication	12
Resolving problems with the LDAP directory server	12
Auditing the LDAP directory server to clean up the server	14
Error messages for LDAP authenticated passwords.	15
Resolving client scheduling problems.	16
Determining the status of a scheduled event	17
Checking for errors in the server activity log	17
Starting and stopping the client service	18
Resolving errors when including or excluding client files during backup processing	19
Identifying files that are included or excluded by the server client option set	19
Excluding files automatically from backup processing.	20
Excluding files with the EXCLUDE.DIR statement	22

Determining whether compression, encryption, and subfile backup statements include or exclude	23
Using delimiters to include or exclude files.	23
Resolving errors due to the incorrectly coded include or exclude list	24
Resolving Snapshot Difference problems.	24
Resolving snapshot directory problems for NetApp or N-Series file system volumes.	26
Resolving login problems when using the encrypted file system on AIX operating systems	27
Resolving image backup errors	27
Resolving Linux image backup errors.	27
Resolving image backup failures when using Linux snapshot	29
Resolving errors during AIX JFS2 snapshot-based backup-archive and image backup.	30
Support solutions for the Tivoli Storage Manager API	31
Gathering API information before calling IBM support.	31
Gathering API files before calling IBM support	32
Determining if data is sent to the Tivoli Storage Manager storage agent rather than the server	34
Running applications that use the API as a non-root user	35
Journal Based Backup problem determination	36
Determining if a backup will be journal-based.	37
Running the journal daemon in the foreground	38
Using the Journal Database Viewing utility	38
Resolving problems with open file support and the logical volume snapshot agent	39
Examining the Windows system event log	39
Obtaining trace data	40
Configuring the system for a full memory dump	40
Forcing a memory dump for a system stoppage when you suspect a logical volume snapshot agent problem	41
Best practices for open file support	41
Using Windows Volume Shadow Copy Services	41
Defining VSS transient errors	41
Defining Windows VSS test flags	42
Windows 2003 VSS fixes	43
Volume Shadow Copy Services tuning	43
Gathering VSS diagnostic information for Microsoft assistance	43
Troubleshooting errors using a VSS trace	44
Running VSS API calls with the vsreq.exe sample program	44
Comparing Tivoli Storage Manager and Ntbackup.exe interaction with VSS	45
Show commands for the backup-archive client	45

Chapter 3. Resolving server problems 49

Recreating the problem	49
Deployment engine error during initialization	49

Checking the server activity log file	50
Checking system error log files for device errors	50
Reverting server options or settings	51
Restarting the scheduling service	51
Resolving server space issues	51
Allocating additional server memory	52
Changing the copy frequency	52
Resolving RELABEL operation errors	53
Resolving installation and upgrade problems	53
Installation log files	53
Resolving a stopped uninstallation process	58
Server installation failure: deployment engine initialization	58
Manually recreating server instances due to connectivity problems	58
Resolving a server upgrade failure	59
Resolving GSKit installation problems	60
Resolving server stoppages	61
Resolving a stoppage or loop	62
Resolving wait state problems with external user repository servers	63
Finding the server error file (dsmserv.err)	64
Finding the system image (core file)	64
Retrieving library files for core analysis	65
Retrieving system log files	66
Retrieving the activity log	66
Detecting errors after a server service starts and stops	66
sqllib/db2dump directory causes shutdown	67
Resolving database errors	68
Resolving database manager starting problems	69
Tracing the User ID/Password plug-in	70
Resolving DB2 automatic database backup errors	70
Limiting DB2 memory allocation	71
Retrieving DB2 version information	71
Locating DB2 diagnostic log files	71
Locating DB2 log files after an upgrade	73
Resolving a missing or incorrect database ID file problem	73
Resolving problems with the BACKUP DB and the RESTORE DB commands	74
Characteristics of the \$\$_TSMDBMGR_\$\$ user ID	79
Resolving database reorganization problems	79
Analyzing the process symptoms to resolve problems	79
Reviewing process messages to determine the state of server operations	80
Analyzing the ANR1221E error message	85
Analyzing the ANR2317W error message	86
Analyzing error messages ANR1330E and ANR1331E	87
Resolving error messages CTGTRV009E and CTGTRV011E	90
Files are not expired after reducing versions	91
Process symptoms indicate migration errors	91
Resolving storage pool issues	92
"ANR0522W Transaction failed..." message received	92
Storage pool experiences high volume usage after increasing MAXSCRATCH value	93

Storage pool is set to use collocation, but volumes contain data that is not collocated	93
Resolving storage problems for active data pools	94

Chapter 4. Resolving communication problems 95

Resolving errors created when connecting to the server	95
Resolving failed connections by clients or administrators	95
Resolving Secure Sockets Layer errors	96
Recovering the key database file password	98

Chapter 5. Resolving Administration Center problems 99

Re-establishing a connection between the Administration Center and a Tivoli Storage Manager server	99
Resolving Tivoli Integrated Portal user authority problems	101
Adding a user ID with access to the Administration Center	102
Resolving a user ID access problem with the Administration Center	103
Resolving a stopped Tivoli Integrated Portal server	103
Running the collector tool to obtain problem-analysis information	103
Diagnosing log-entry problems by using the log analyzer tool (showlog)	105
Resolving excessive memory consumption problems with the Tivoli Integrated Portal server	106
Configuring the IP address to align with the Administration Center	106
Resolving server access problems	107
Resolving Administration Center health monitor problems	108
Health monitor conditions that can cause an unknown server status	109
Resolving health monitor conditions that can cause a warning or critical storage status	109
Health monitor conditions that can cause a warning or critical database status	110
Determining when to resynchronize the ADMIN_CENTER administrator ID password	110
Administration Center Support utility	111
Responding to Administration Center task failure messages	111
Responding to Administration Center messages about unexpected results	112
Checking the server activity log to resolve Administration Center problems	112
Resolving errors caused by starting or stopping a wizard or portlet	113
Resolving problems caused by internal errors	114
Determining the source of a message	114
Tivoli Storage Manager messages	115
Locating Tivoli Integrated Portal error messages	115
Resolving Tivoli Storage Manager server command-definition file problems	116

Resolving backup-archive client deployment problems	117
Configuring the server for automatic backup-archive client deployments	117
Resolving automatic deployment scheduling problems	119
Restarting the client operating system during a deployment	119
Backup-archive client deployment status is unknown.	120

Chapter 6. Resolving storage agent problems 121

Checking the server activity log for storage agent information	121
Resolving an error caused by reading or writing to a device	121
Resolving problems caused by changing storage agent options	122
Resolving problems caused by changing server options or settings.	122
Storage agent LAN-free setup	122
Resolving the issue of data being sent directly to the server	122
Resolving a disqualified LAN-free-enabled storage pool.	123
Ensuring that data is transferred using a LAN-free environment	124

Chapter 7. Using trace to resolve problems 125

Trace classes for the Administration Center	125
Enabling Administration Center trace	127
Enabling a trace for the server or storage agent	129
Enabling a stack trace for specific messages for the server or storage agent	130
Trace classes for a server or storage agent	131
Show commands for the server or storage agent	145
Enabling a trace for the Tivoli Storage Manager device driver	158
Tracing from the server console	158
Tracing data from a command shell for AIX, Oracle Solaris, and Windows	159
Tracing to detect a code page conversion failure	160
Tracing data for the client	160
Client and Journal Daemon traceflags	161
Client trace classes	162
Enabling a backup-archive client trace	166
Determining if data is encrypted or compressed during backup-archive through trace	175
Tracing data for the API.	176
Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on an AIX or Linux system	177
Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on a Windows operating system.	179

Chapter 8. Resolving Tivoli Monitoring for Tivoli Storage Manager problems . 181

Resolving installation issues with Tivoli Monitoring for Tivoli Storage Manager	181
---	-----

Reviewing the logs.zip file to resolve installation failures	181
History configuration failed	183
Tivoli Enterprise Portal configuration fails on SUSE Linux Enterprise 11 during installation.	185
Resolving errors with the Tivoli Enterprise Portal component of IBM Tivoli Monitoring	186
No data, or old data, is displayed in the Tivoli Enterprise Portal workspaces	186
Empty workspaces when agent queries run too long	188
Cannot log on to the Tivoli Enterprise Portal	190
Missing monitoring agent from the Tivoli Enterprise Portal	191
Start and end times display 00:00:00 in Node Activity workspace	192
Removing agent instances does not stop them	192
Resolving historical data reporting problems in the Warehouse Proxy workspace	193
Unable to start Tivoli Enterprise Portal using a non-English administrator name	193
Resolving errors with the Tivoli Common Reporting reports	194
Resolving no data in Tivoli Common Reporting reports	194
No current data displayed in the Tivoli Common Reporting reports.	199
Missing server name in the Tivoli Common Reporting report prompt page.	200
No data is available message displayed in a Cognos report	202
Empty data set message in Tivoli Common Reporting	204
Error message RQP-DEF-0177 is displayed in a Cognos report	205
Unable to find data source 'TDW' message is displayed in a Cognos report	206
Missing BIRT or Cognos reports	209
No group membership exists to access Tivoli Common Reporting reports.	210
Start and end times display 12:00:00 AM in a Cognos report	211
Reports do not display in Japanese, Korean, or Traditional Chinese languages	211
Resolving a java.net.ConnectException: CTGRI0011E error.	212
Resolving web browser issues with the Administration Center or Tivoli Enterprise Portal	213
Java Web start applet does not start when Tivoli Enterprise Portal server is running on a Windows system	214
File system is full or almost full	215
Resetting passwords for Tivoli Monitoring for Tivoli Storage Manager	216
Verifying the WAREHOUS data for each report	219

Chapter 9. Resolving data storage problems 223

Resolving unreadable data problems	223
Checking the server activity log to resolve data storage issues	223

Checking HELP for messages issued for a data storage problem	223
Recreating the data storage problem.	223
Resolving data storage errors related to reading or writing to a device	224
Changing the storage hierarchy to resolve data storage problems	224
Changing the server policies to resolve data storage problems	224
Resolving a data storage backup or copy problem that occurs only with a specific node	225
Resolving a data storage problem that occurs only for a specific volume	225
Hints and tips for storage	225
Device driver hints and tips	225
Hard disk drives and disk subsystems hints and tips.	231
Tape drives and libraries hints and tips.	234
SAN hints and tips	235
NDMP filer-to-Tivoli Storage Manager server operation hints and tips	250
Resolving SCSI device problems	251
Resolving sequential media volume (tape) errors through messages ANR0542W or ANR8778W	251

Appendix A. Getting call stack information from a core file	253
--	------------

Appendix B. Installing and running the tsmdiag utility	255
---	------------

Appendix C. IBM Global Security Kit return codes.	259
--	------------

Appendix D. Accessibility features for the Tivoli Storage Manager product family.	269
--	------------

Notices	271
Trademarks	273

Glossary	275
---------------------------	------------

Index	277
------------------------	------------

About this publication

This publication helps you determine the source of problems with the servers and clients in your IBM® Tivoli® Storage Manager environment.

Before using this publication, ensure that you are familiar with the following areas:

- Your Tivoli Storage Manager server and client operating systems
- The communication protocols installed on your client and server computers

Troubleshooting information for the following component products can be found in their product-specific documentation:

- Data Protection for Lotus® Domino® UNIX, Linux, and OS/400®
- Data Protection for Lotus Domino for Microsoft Windows
- Data Protection for SAP, Installation and User's Guide for Oracle
- Data Protection for SAP, Installation and User's Guide for DB2®
- Data Protection for Microsoft Exchange
- Data Protection for Microsoft SQL Server

Who should read this guide

This guide is written for anyone administering or managing IBM Tivoli Storage Manager. Similarly, information provided by this guide can be useful to business partners and anyone with the responsibility to support Tivoli Storage Manager.

You should be familiar with Tivoli Storage Manager and the operating systems used for the Tivoli Storage Manager environment.

This document references error logs, trace facilities, and other diagnostic information for Tivoli Storage Manager. These trace facilities and diagnostic tools are not a programming interface for the product. Tivoli Storage Manager product development and support use these tools for diagnosing and debugging problems. For this guide, these are provided only to aid in diagnosing and debugging any problems. Trace facilities are subject to change without notice and might vary, depending upon the version and release of the product or the platform on which the product is being run. Information that is referenced within this guide might not be supported or applicable to all versions or releases of the product. Changes are periodically made to the information herein. IBM might make improvements and changes in the products and the programs described in this publication at any time without notice.

Publications

Publications for the IBM Tivoli Storage Manager family of products are available online. The Tivoli Storage Manager product family includes IBM Tivoli Storage FlashCopy® Manager, IBM Tivoli Storage Manager for Space Management, IBM Tivoli Storage Manager for Databases, and several other storage management products from IBM Tivoli.

To search all publications, go to the Tivoli Storage Manager information center at <http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3>.

You can download PDF versions of publications from the Tivoli Storage Manager information center or from the IBM Publications Center at <http://www.ibm.com/shop/publications/order/>.

Go to Tivoli Documentation Central to find information centers that contain official product documentation for current and previous versions of Tivoli products, including the Tivoli Storage Manager product family. You can find Tivoli Documentation Central at <https://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Home>.

You can also order some related publications from the IBM Publications Center website at <http://www.ibm.com/shop/publications/order/>. The website provides information about ordering publications from countries other than the United States. In the United States, you can order publications by calling 1-800-879-2755.

Tivoli Storage Manager publications

The following tables list the publications that make up the Tivoli Storage Manager library.

Table 1. Tivoli Storage Manager server publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for AIX Installation Guide</i>	GC23-9781
<i>IBM Tivoli Storage Manager for AIX Administrator's Guide</i>	SC23-9769
<i>IBM Tivoli Storage Manager for AIX Administrator's Reference</i>	SC23-9775
<i>IBM Tivoli Storage Manager for HP-UX Installation Guide</i>	GC23-9782
<i>IBM Tivoli Storage Manager for HP-UX Administrator's Guide</i>	SC23-9770
<i>IBM Tivoli Storage Manager for HP-UX Administrator's Reference</i>	SC23-9776
<i>IBM Tivoli Storage Manager for Linux Installation Guide</i>	GC23-9783
<i>IBM Tivoli Storage Manager for Linux Administrator's Guide</i>	SC23-9771
<i>IBM Tivoli Storage Manager for Linux Administrator's Reference</i>	SC23-9777
<i>IBM Tivoli Storage Manager for Oracle Solaris Installation Guide</i>	GC23-9784
<i>IBM Tivoli Storage Manager for Oracle Solaris Administrator's Guide</i>	SC23-9772
<i>IBM Tivoli Storage Manager for Oracle Solaris Administrator's Reference</i>	SC23-9778
<i>IBM Tivoli Storage Manager for Windows Installation Guide</i>	GC23-9785
<i>IBM Tivoli Storage Manager for Windows Administrator's Guide</i>	SC23-9773
<i>IBM Tivoli Storage Manager for Windows Administrator's Reference</i>	SC23-9779
<i>IBM Tivoli Storage Manager for z/OS® Media Installation and User's Guide</i>	SC27-4018
<i>IBM Tivoli Storage Manager Upgrade and Migration Guide for V5 Servers</i>	GC27-4017
<i>IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack®</i>	SC27-2828

Table 2. Tivoli Storage Manager storage agent publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for SAN for AIX Storage Agent User's Guide</i>	SC23-9797

Table 2. Tivoli Storage Manager storage agent publications (continued)

Publication title	Order number
<i>IBM Tivoli Storage Manager for SAN for HP-UX Storage Agent User's Guide</i>	SC23-9798
<i>IBM Tivoli Storage Manager for SAN for Linux Storage Agent User's Guide</i>	SC23-9799
<i>IBM Tivoli Storage Manager for SAN for Oracle Solaris Storage Agent User's Guide</i>	SC23-9800
<i>IBM Tivoli Storage Manager for SAN for Windows Storage Agent User's Guide</i>	SC23-9553

Table 3. Tivoli Storage Manager client publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for UNIX and Linux: Backup-Archive Clients Installation and User's Guide</i>	SC23-9791
<i>IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User's Guide</i>	SC23-9792
<i>IBM Tivoli Storage Manager Using the Application Programming Interface</i>	SC23-9793
<i>IBM Tivoli Storage Manager for Space Management for UNIX and Linux: User's Guide</i>	SC23-9794
<i>IBM Tivoli Storage Manager HSM for Windows Administration Guide</i>	SC23-9795

Table 4. Tivoli Storage Manager data protection publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Installation and User's Guide</i>	GC27-4010
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for UNIX and Linux Installation and User's Guide</i>	SC27-4019
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for Windows Installation and User's Guide</i>	SC27-4020
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Installation and User's Guide</i>	GC27-4009
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino UNIX and Linux Installation and User's Guide</i>	SC27-4021
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino for Windows Installation and User's Guide</i>	SC27-4022
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for DB2</i>	SC33-6341
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for Oracle</i>	SC33-6340
<i>IBM Tivoli Storage Manager for Virtual Environments Installation and User's Guide</i>	SC27-2898
<i>IBM Tivoli Storage Manager for Microsoft SharePoint Guide</i>	N/A

Table 5. IBM Tivoli Storage Manager troubleshooting and tuning publications

Publication title	Order number
<i>IBM Tivoli Storage Manager Problem Determination Guide</i>	GC23-9789

Table 5. IBM Tivoli Storage Manager troubleshooting and tuning publications (continued)

Publication title	Order number
<i>IBM Tivoli Storage Manager Performance Tuning Guide</i>	GC23-9788
<i>IBM Tivoli Storage Manager Client Messages and Application Programming Interface Return Codes</i>	SC27-2878
<i>IBM Tivoli Storage Manager Server Messages and Error Codes</i>	SC27-2877
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Messages</i>	GC27-4011
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Messages</i>	GC27-4012
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle Messages</i>	SC27-4014
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino Messages</i>	SC27-4015
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Messages</i>	SC27-4016

Note: You can find information about IBM System Storage® Archive Manager at the Tivoli Storage Manager v6.3.0 information center.

Tivoli Storage FlashCopy Manager publications

The following table lists the publications that make up the Tivoli Storage FlashCopy Manager library.

Table 6. Tivoli Storage FlashCopy Manager publications

Publication title	Order number
<i>IBM Tivoli Storage FlashCopy Manager for UNIX and Linux Installation and User's Guide</i>	SC27-4005
<i>IBM Tivoli Storage FlashCopy Manager for Windows Installation and User's Guide</i>	SC27-4006
<i>IBM Tivoli Storage FlashCopy Manager for VMware Installation and User's Guide</i>	SC27-4007
<i>IBM Tivoli Storage FlashCopy Manager Messages</i>	GC27-4008

Support information

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

Start at the IBM Support Portal: <http://www.ibm.com/support/entry/portal/>. You can select the products that you are interested in and search for a wide variety of relevant information.

Getting technical training

Information about Tivoli technical training courses is available online.

Go to the following websites to sign up for training, ask questions, and interact with others who use IBM storage products.

Tivoli software training and certification

Choose from instructor led, online classroom training, self-paced Web classes, Tivoli certification preparation, and other training options at <http://www.ibm.com/software/tivoli/education/>

Tivoli Support Technical Exchange

Technical experts share their knowledge and answer your questions in webcasts at http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html.

Storage Management community

Interact with others who use IBM storage management products at <http://www.ibm.com/developerworks/servicemanagement/sm/index.html>

Global Tivoli User Community

Share information and learn from other Tivoli users throughout the world at <http://www.tivoli-ug.org/>.

IBM Education Assistant

View short "how to" recordings designed to help you use IBM software products more effectively at <http://publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp>

Searching knowledge bases

If you have a problem with your Tivoli Storage Manager family product, there are several knowledge bases that you can search.

Begin by searching the Tivoli Storage Manager Information Center at <http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3>. From this website, you can search the current Tivoli Storage Manager documentation.

Searching the Internet

If you cannot find an answer to your question in the IBM Tivoli Storage Manager information center, search the Internet for the information that might help you resolve the problem.

To search multiple Internet resources, go to the IBM support website at <http://www.ibm.com/support/entry/portal/>.

You can search for information without signing in. Sign in using your IBM ID and password if you want to customize the site based on your product usage and information needs. If you do not already have an IBM ID and password, click **Sign in** at the top of the page and follow the instructions to register.

From the support website, you can search various resources including:

- IBM technotes.
- IBM downloads.
- IBM Redbooks® publications.
- IBM Authorized Program Analysis Reports (APARs). Select the product and click **Downloads** to search the APAR list.

If you still cannot find a solution to the problem, you can search forums and newsgroups on the Internet for the latest information that might help you find problem resolution.

An independent user discussion list, ADSM-L, is hosted by Marist College. You can subscribe by sending an email to listserv@vm.marist.edu. The body of the message must contain the following text: `SUBSCRIBE ADSM-L your_first_name your_family_name`.

To share your experiences and learn from others in the Tivoli Storage Manager and Tivoli Storage FlashCopy Manager user communities, go to Service Management Connect (<http://www.ibm.com/developerworks/servicemanagement/sm/index.html>). From there you can find links to product wikis and user communities.

Using IBM Support Assistant

IBM Support Assistant is a complimentary software product that can help you with problem determination. It is available for some Tivoli Storage Manager and Tivoli Storage FlashCopy Manager products.

To learn about which products are supported, go to the IBM Support Assistant download web page at <http://www.ibm.com/software/support/isa/download.html>.

IBM Support Assistant helps you gather support information when you must open a problem management record (PMR), which you can then use to track the problem. The product-specific plug-in modules provide you with the following resources:

- Support links
- Education links
- Ability to submit problem management reports

You can find more information at the IBM Support Assistant website:

<http://www.ibm.com/software/support/isa/>

You can also install the stand-alone IBM Support Assistant application on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use. Find add-ons for specific products at <http://www.ibm.com/support/docview.wss?uid=swg27012689>.

Finding product fixes

A product fix to resolve your problem might be available from the IBM software support website.

You can determine what fixes are available by checking the IBM software support website at <http://www.ibm.com/support/entry/portal/>.

- If you previously customized the site based on your product usage:
 1. Click the link for your product, or a component for which you want to find a fix.
 2. Click **Downloads**, and then click **Fixes by version**.
- If you have not customized the site based on your product usage, click **Downloads** and search for your product.

Receiving notification of product fixes

You can receive notifications about fixes, flashes, upgrades, and other news about IBM products.

To sign up to receive notifications about IBM products, follow these steps:

1. From the support page at <http://www.ibm.com/support/entry/portal/>, click **Sign in to create, manage, or view your subscriptions** in the **Notifications** pane.
2. Sign in using your IBM ID and password. If you do not have an ID and password, click **register now** and complete the registration process.
3. Click **Manage all my subscriptions** in the **Notifications** pane.
4. Click the **Subscribe** tab and then click **Tivoli**.
5. Select the products for which you want to receive notifications and click **Continue**.
6. Specify your notification preferences and click **Submit**.

Contacting IBM Software Support

You can contact IBM Software Support if you have an active IBM subscription and support contract and if you are authorized to submit problems to IBM.

To obtain help from IBM Software Support, complete the following steps:

1. Ensure that you have completed the following prerequisites:
 - a. Set up a subscription and support contract.
 - b. Determine the business impact of your problem.
 - c. Describe your problem and gather background information.
2. Follow the instructions in “Submitting the problem to IBM Software Support” on page xiv.

Setting up a subscription and support contract

Set up a subscription and support contract. The type of contract that you need depends on the type of product you have.

For IBM distributed software products (including, but not limited to, IBM Tivoli, Lotus, and Rational® products, as well as IBM DB2 and IBM WebSphere® products that run on Microsoft Windows or on operating systems such as AIX or Linux), enroll in IBM Passport Advantage® in one of the following ways:

- **Online:** Go to the Passport Advantage website at <http://www.ibm.com/software/lotus/passportadvantage/>, click **How to enroll**, and follow the instructions.
- **By telephone:** You can call 1-800-IBMSERV (1-800-426-7378) in the United States. For the telephone number to call in your country, go to the IBM Software Support Handbook web page at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you must understand and assess the business impact of the problem you are reporting.

Severity 1	Critical business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.
Severity 2	Significant business impact: The program is usable but is severely limited.
Severity 3	Some business impact: The program is usable with less significant features (not critical to operations) unavailable.
Severity 4	Minimal business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem has been implemented.

Describing the problem and gathering background information

When explaining a problem to IBM, it is helpful to be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? For example, hardware, operating system, networking software, and so on.
- Are you using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submitting the problem to IBM Software Support

You can submit the problem to IBM Software Support online or by telephone.

Online

Go to the IBM Software Support website at [http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_\(general\)](http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_(general)). Sign in to access IBM Service Requests and enter your information into the problem submission tool.

By telephone

For the telephone number to call in your country, go to the IBM Software Support Handbook at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

Chapter 1. Help facilities

IBM Tivoli Storage Manager has several outlets for resolving problems that you might have with the server or the backup-archive client.

Backup-archive client help

Use the help command to display information about commands, options, and messages. If you use the help command on the initial command line, no server contact is made and no password is needed.

Syntax

```
▶▶—dsmc help—▶▶
```

<i>command-name</i> [<i>subcommand-name</i>]
<i>option-name</i>
<i>TOC-section-number</i>
[<i>ANS</i>] <i>message-number</i>

Entering the **HELP** command with no arguments causes help to display the complete table of contents. Either with the initial command or when HELP displays a prompt, you can enter the following parameters.

Parameters

command-name [*subcommand-name*]

Specifies a command name and, optionally, a subcommand name or their abbreviation. For example: **backup image**, or **b i**. In this case the combination should be unique. Non-unique abbreviations result in the display of the first section of the entire help file matching the abbreviation. This parameter is optional.

option-name

Specifies the name of an option. For example: **domain** or **do**. This parameter is optional.

TOC-section-number

Specifies a table of contents section number. For example: 1.5.3. This parameter is optional.

[*ANS*]*message-number*

Specifies a message number with or without its prefix. For example: **ans1036** or **1036**. This parameter is optional. The severity code is never necessary. Entering **ans1036E** results in a not-found response.

Important: If you enter arguments that do not fit these descriptions, you might get unexpected results (or no results) displayed. If you enter more than two arguments, your help request is rejected. Where a command name and an option name are the same, for example: **incremental** (command) and **incremental** (option), you can only get help on the option by entering its table-of-contents section number.

The requested help text is displayed in one or more sections, depending on the number of display lines that are available in your command window. When

enough lines are displayed to fill the display space, or when the end of the requested help text is displayed, a prompt is presented with instructions on what can be entered at that prompt. To continue displaying text for your current selection, press **Enter** or press the “d” key to scroll down. To scroll up in the current selection, press the “u” key and press **Enter**. Use the “q” key to quit the help facility. Other choices might be presented, so read all instructions.

Proper display of the help text requires a usable display width of 72 characters. A display width less than 72 characters causes sentences that are 72 characters wide to wrap to the next line. This can cause the displayed help text to begin somewhere within the section rather than at the beginning. The lines that are not displayed can be viewed by using the scrolling function of the terminal to move up.

Accessing help for the Windows service configuration utility (dsmcutil)

Windows

You must issue the **DSMCUTIL** command to obtain help information.

When you issue **DSMCUTIL HELP**, you are prompted to enter **C** for a basic command summary or **F** for full help information.

The basic command summary is a listing of available **dsmcutil** sub-commands and options. This information is most useful if you are already familiar with using **dsmcutil**.

The full help information displays more comprehensive information within the Windows help utility.

Server or storage agent help

The server and storage agent both include a help facility. The help facility provides descriptions and syntax for server commands and a full description of server messages.

Accessing server or storage agent help for commands

Issue the **HELP** command to access help for the server or storage agent.

To display command-line help for server commands that have unique names, you can type `help commandName`, where *commandName* is the name of the server command for which you want information. For example, to display help for the **REGISTER NODE** command, type `help register node`. Command syntax and parameter descriptions are displayed in the output.

You can also type `help` followed by the topic number for the command. Topic numbers are listed in the table of contents for command-line help, for example:

```
3.0 Administrative commands
  3.46 REGISTER
    3.46.1 REGISTER ADMIN (Register an administrator)
    3.46.2 REGISTER LICENSE (Register a new license)
    3.46.3 REGISTER NODE (Register a node)
```

To display help about the **REGISTER NODE** command, type:

```
help 3.46.3
```

Use topic numbers to display command-line help for subcommands. **DEFINE DEVCLASS** is an example of a command that has subcommands. For example, you can specify the **DEFINE DEVCLASS** command for 3570 device classes and for 3590 device classes:

```
3.0 Administrative commands
...
3.13.10 DEFINE DEVCLASS (Define a device class)
    3.13.10.1 DEFINE DEVCLASS (Define a 3570 device class)
    3.13.10.2 DEFINE DEVCLASS (Define a 3590 device class)
...
```

To display help for the **DEFINE DEVCLASS** command for 3570 device classes, type:
help 3.13.10.1

Accessing help for messages

Issue the help command to access help for messages.

Issue the following command for help on a server message: `HELP message number` where *message number* is the message for which you want information. If you specify the message number without including the message prefix, for example `HELP 0445`, it assumes the message prefix ANR and reports the help information for ANR0445W. If the message number is specified with the prefix, for example `HELP ANR0445`, it reports the help information for that message. Issue `HELP ANR0445` to view the following example output for that message:

```
ANR0445W Protocol error on session session number for node client node name
(client platform) - maximum group transaction size exceeded.
Explanation: The server detects a protocol error on the specified session
because the client has attempted to group more than the maximum database
update operations in a single database transaction.
System Action: The server ends the client session.
User Response: Correct the programming error in the client program if it has
been written by your installation using WDSF verbs. Otherwise, contact your
service representative.
```

Command-line interface help for the client

The command-line client interface includes a help facility that provides descriptions and syntax for client commands and options and a full description of client messages.

Help information for the graphical user interface (GUI) and web GUI clients is available through the **Help** menu item.

Reporting a problem with a help topic

When you want to report a problem with the help system, you must first collect specific information.

1. Record what you clicked to get the help. For example, if you clicked the question mark for a portal, record the name of the portal.
2. View the source of the help pop-up window. On most browsers, a right mouse-click shows you a menu with a **View Source** option. Select **View Source** to view the HTML source code for that window. Write down the title of that window, which is the URL or the name of the file that the help system is trying to show.

Chapter 2. Resolving Tivoli Storage Manager client problems

Resolving problems with the client application can involve connecting to the server, changing policy settings, reproducing the error, or several other possible options.

Examining error messages

You can examine the error messages that are generated during program operation to help resolve problems that might occur.

The QUIET processing option for the IBM Tivoli Storage Manager client suppresses all messages. Restart the client with the QUIET option turned off to allow all messages to be issued, which will provide a more complete understanding of the problem.

Check for any ANSnnnnx messages issued to the console, `dsmsched.log`, or `dsmerror.log`. Additional information for ANSnnnnx messages is available in either the *Client Messages and Application Programming Interface Return Codes* or from the client HELP facility.

Examining the server activity log messages

Use the **QUERY ACTLOG** command to view the server activity log file and the messages issued for this IBM Tivoli Storage Manager client session.

The messages from the server activity log might provide additional information about the symptoms for the problem or might provide information about the actual cause of the problem that the client encountered.

Identifying when and where the problem can occur

Problems with IBM Tivoli Storage Manager client processing often occur only when you are performing specific operations, at certain times, or only on certain client computers.

To further isolate when and where a problem occurs, determine the following answers:

- Does this problem occur for a single client, some clients, or all clients for a given server?
- Does this problem occur for all clients running on a specific operating system?
- Does this problem occur for specific files, for files that are in a specific directory, for files on a specific drive, or for all files?
- Does this problem occur for clients on a specific network, subnet, or all parts of the network?
- Does this problem occur only for the command-line client, the GUI client, or the web client?
- Does Tivoli Storage Manager always fail when processing the same file or directory, or is it different from run to run?

Reproducing the problem

When you reproduce a problem as part of problem determination, try to minimize the impact that the process has on the IBM Tivoli Storage Manager.

You can help Tivoli Storage Manager support by minimizing the complexity of the environment in which you want to recreate the problem. The following options can be used to minimize the complexity of the environment:

- Use a minimal options file consisting of only TCPSEVERADDRESS, TCPPOINT, and NODENAME.
- If the problem occurs for a file during incremental backup, try to reproduce the problem with a selective backup of just that file.
- If the problem occurs during a scheduled event, try to reproduce the problem by manually running the command.

Collecting documentation to resolve problems with the client application

The support staff at IBM is better able to resolve a problem if you can supply them with relevant documentation. The IBM Tivoli Storage Manager backup-archive client creates information in a number of different sources.

The IBM Support Assistant for the Tivoli Storage Manager backup-archive client provides product-specific web search tools as well as a data collector. The data collector can gather various documentation when reporting a problem to IBM technical support.

Tip: Tivoli Storage Manager has a built-in help facility within the client command line. Issue the **dsmc help** command to access the command-line client's help facility. The help facility is a menu-driven interface with information that includes the command reference, option reference, and extended information about client messages.

Tivoli Storage Manager client problems and configuration information might be found in one or more of the following documents:

- Error log. The client error log file is `dsmerror.log`.
- Scheduler log. The error log for the client scheduler is `dsmsched.log`.
- Web client log. The error log for the web client is `dsmwebcl.log`.
- Options files. The backup-archive client might use a combination of files for its configuration. These files are `dsm.opt`, `dsm.sys` for AIX®, HP-UX, Linux, or Oracle Solaris systems, and the include-exclude file.
- Trace data. If tracing was active, the file containing the trace data could be provided to support.
- Application dump. If the Tivoli Storage Manager backup-archive client stops running, many platforms will generate an application dump. The operating system provides the application dump.
- Memory dump. If the Tivoli Storage Manager backup-archive client stops, a memory dump can be generated that can then be used to help with diagnosis. The type of system determines how the memory dump occurs, and the operating system provides the memory dump.

The **DSMC QUERY SYSTEMINFO** command is available and collects most of this information in the `dsminfo.txt` file. The following items can help you to determine Tivoli Storage Manager problems:

- A list of all the software installed on the client system. The client might experience problems due to interactions with other software on the computer or because of the maintenance levels of software that the client uses.
- Client option sets defined on the server that apply to this client node. Issue the **QUERY CLOPTSET** command to search for the client option sets.
- Server options. There are a number of server options that are used to manage the interaction between the backup-archive client and server. An example of one such server option is `TXNGROUPMAX`.
- Information about this node as it is defined to the server. To collect this information, issue the **QUERY NODE *nodeName* F=D** command using an administrative client connected to the server.
- Schedule definitions for the schedules that apply to this node. The schedule definitions can be queried from the server when you issue the **QUERY SCHEDULE** command.
- The policy information configured for this node on the Tivoli Storage Manager server. The policy information can be queried from the server when you issue the **QUERY DOMAIN, QUERY POLICYSET, QUERY MANAGEMENTCLASS, or QUERY COPYGROUP** commands.

Determining why the `dsmc` or `dsmadm` or `dsmj` program does not start

The IBM Tivoli Storage Manager backup-archive client uses the `dsmc`, `dsmadm`, or `dsmj` program in its startup procedure. When one of these programs does not start, the backup-archive client does not start.

The `dsmc`, `dsmadm`, and `dsmj` programs have the following definitions:

dsmc The backup-archive command-line client.

dsmadm

The administrative command-line client.

AIX **HP-UX** **Linux** **Solaris** **dsmj**

The backup-archive client graphical user interface (GUI). The Oracle Java™ runtime version is checked when you first start the `dsmj` GUI. In some cases, this check is done improperly and the `dsmj` startup might fail with a “bad number” message.

AIX **HP-UX** **Linux** **Solaris** Processing stops and the following message is displayed if the `dsmc` or `dsmadm` or `dsmj` program does not start:

```
ANS1398E Initialization functions cannot open one of the
Tivoli Storage Manager logs or a related file: /dsmerror.log. errno = 13,
The file access permissions do not allow the specified action.
```

Remember: **AIX** **HP-UX** **Linux** **Solaris** The `dsmerror.log` file is used only as an example file in the message.

Client applications do not run without being able to write to a log file, and the system denies “write” access to the log file named in the message. If the log file does not exist, it is created with default permissions. The following rules apply:

- The name and the directory specified by the `ERRORLOGNAME` option are used.

- If the option is absent, the name `dsmerror.log` in the directory specified in the **DSM_LOG** environment variable, if present, is used. Otherwise, the name `dsmerror.log` in the current working directory is used.

The following issues are applicable if the default permissions are used:

- A log file created by the root user might not be written to by any other user
- The root user must set the proper permissions or access control lists (ACLs) to allow free use of the client application by all users who need to use it

If the log file is successfully created, an error-free session will leave a zero-length (empty) log file.

The Tivoli Storage Manager client does not try to create log files in the root directory. Message ANS1398E is displayed when the method in the first rule, above, directs the log file to be created in the root directory.

If a log file exists and can be located, Tivoli Storage Manager uses the method from the first rule. It can also be in the root directory, if you so choose. Furthermore, whatever permissions you give that log file will be preserved by Tivoli Storage Manager code.

Create your log file in advance of first use, ensuring that all eligible users have write access to it. Define the **ERRORLOGNAME** option or the **DSM_DIR** environment variable to designate your predefined log file.

Attention: A system log file error indicates that you cannot write to the `dsmerror.log` file. Certain background Tivoli Storage Manager applications might not start due to errors writing to `dsmerror.log`. When these errors occur, a number of errors are recorded in the Windows system event log file ('system log' on other platforms).

For example:

```
C:\Program Files\Tivoli\Tsm\baclient>net start "TSM Sched"
The server scheduling service is starting.
The server scheduling service could not be started.
A service specific error occurred: 12.
```

Additional setup steps are required for non-root users in order for them to be able to run Tivoli Storage Manager applications or Tivoli Storage Manager for Data Protection applications. You will receive the ANS1398E error if you try to run Tivoli Storage Manager applications using an error log file which has already been generated by root, that is left with default permissions. For data protection clients, you might only receive a Tivoli Storage Manager API error. Here is one method for setting up `dsmerror.log` for use by non-root users:

1. Set **ERRORLOGNAME** in `dsm.sys`. For example, `errorLogName /var/msgs/tsm/dsmerror.log`
2. Generate **dsmerror.log**. `dsmc q sess`
3. Modify the permissions on `dsmerror.log` to allow writing by all users. `chmod 666 /var/msgs/tsm/dsmerror.log`

Resolving problems with client option sets

With client option sets, administrators can specify additional options that might not be included in the option file of the backup-archive client. The backup-archive client uses these options during a backup, archive, restore, or retrieve process.

An administrator for IBM Tivoli Storage Manager can create a set of client options to be used by a client node on Tivoli Storage Manager. The client options are defined on the Tivoli Storage Manager server. The client options that are specified in the client option set are used in conjunction with the client options file.

The order in which the options are processed can be controlled. Multiple options can be defined and then assigned a sequence number, with these options then processed from low to high sequence. The following example displays the **INCLEXCL** options:

Option	Sequence number	Override	Option Value
INCLEXCL	0	No	exclude 'sys:\backup*'
INCLEXCL	1	No	include 'sys:\system*'
INCLEXCL	2	No	include 'sys:\tmp*'

This sequence results in the exclusion of all files in the `sys:\backup*` path, while the files in the `sys:\system*` and `sys:\tmp*` paths are backed up.

Resolving problems with client option sets

Use client option sets to resolve various problems, from having critical environments where restoring is a high priority, to using a database that does not stop.

Tip: Trace settings for the client option sets are specified in the IBM Tivoli Storage Manager option file for all backup-archive clients.

The following scenarios show you how you can take advantage of the client option set.

Scenario 1:

Having an environment where restoring is a high priority.

Use the `COLLOCATEBYFILESPEC` option so that all filespec data is stored on as few tapes as possible, which enhances restore processing by using less tape mounts. You do not want the client to be able to override this option. Issue the following Tivoli Storage Manager server command:

```
Define cloptset crit_rest description="Critical Restore Option Sets"  
Define clientopt crit_rest collocatebyfilespec yes force=yes  
Update node dale cloptset=crit_rest
```

Scenario 2:

Using workstations that are on a slow network with limited space for data on the Tivoli Storage Manager server.

Use the compression option to limit the amount of data that is sent and stored. Issue the following Tivoli Storage Manager server command:

```
Define cloptset space_rest description="Space Restriction Option Sets"
Define clientopt space_rest compressalways no force=yes
Define clientopt space_rest compression yes force=yes
Update node mark cloptset=space_rest
```

Scenario 3:

Using a database that does not stop.

A problem exists with the database because the files are open and the server cannot back them up. Exclude all files and subdirectories from Tivoli Storage Manager backups and add the files and subdirectories to the existing "space_rest" client option set. Issue the **EXCLUDE DIR** command and specify the directory path that is to be excluded. Issue the following Tivoli Storage Manager server command:

```
Define clientopt space_rest inclxcl "exclude.dir c:\notes\data"
```

Scenario 4:

Finishing backups using a fast network and wanting to make the best possible use of Tivoli Storage Manager client resources.

Set the **RESOURCEUTILIZATION** option to the maximum amount. Issue the following Tivoli Storage Manager server command:

```
Define cloptset unix_srv description="UNIX Server Option Sets"
Define clientopt unix_srv resourceutilization 10 force=yes
```

Resolving password authentication problems

A client authentication error can occur if the password expires. The error can also occur, however, if either the server or the client is renamed, or if the IBM Tivoli Storage Manager administrative user identification password expires.

If you receive the **ANS1025E** error message during an interactive session, the probable cause is an incorrect password. Change the password using either of the following two methods:

- The administrator can reset the password by issuing the **UPDATE NODE** command.
- Issue the **DSMC QUERY SESSION** command and, when prompted, enter the new password.

If the **ANS1025E** error message occurs during a noninteractive session such as central scheduling, ensure that the client option is **PASSWORDACCESS GENERATE**. The **PASSWORDACCESS GENERATE** option causes the client to store the password locally. The password is encrypted and stored either in the registry for Windows clients or, for Macintosh clients, in a file named **TSM.PWD**. Do not edit the registry or the **TSM.PWD** file. Instead, see the following actions:

- Ensure that **PASSWORDACCESS GENERATE** is set in the option file.
- Issue a **DSMC QUERY SESSION** command. This command will force-set the locally-stored password.
- If the previous actions do not resolve the problem, update the node's password by issuing the **UPDATE NODE** administrative command.
- Reissue the **DSMC QUERY SESSION** command, providing the new password.

To see the password expiration setting for a particular node, issue the **QUERY NODE F=D** administrative command. Look for the *Password Expiration Period* field.

Important: Expiration does not apply to node or administrator passwords that authenticate with an LDAP directory servers. If you want to change the settings for these passwords, you must do it on the LDAP directory server.

Remember: If the password expiration period field is blank, the default password expiration period (90 days) applies.

1. To change the password expiration period for a particular node, issue the administrative **UPDATE NODE** command with the option **PASSEXP=*n***, where *n* is the number of days. A value of 0 disables the password expiration.

If a Windows client node is unable to connect to the server after having been renamed, verify that the node name was changed in both the client options file and Windows registry. When the client scheduler runs as a foreground process using the **DSMC SCHED** command, Tivoli Storage Manager uses the nodename in the client options file to contact the server. However, when the scheduler runs as a Windows service, Tivoli Storage Manager uses the node name in the Windows registry.

2. Issue the **DSMCUTIL UPDATE SCHEDULE** command to achieve the following results:
 - With the *node* parameter, address how to change the nodename used with the Tivoli Storage Manager scheduler service on Windows
 - With the *validate:yes* parameter, contact the Tivoli Storage Manager server to authenticate (and store the updated password)

For more information, see the “Change processing options used by the scheduler service” topic in the *Windows Backup-Archive Clients Installation and User’s Guide*.

Resolving LDAP-authenticated password problems

With a password that authenticates with an LDAP directory server, you can use uppercase and lowercase letters, special characters, and numbers. A password that authenticates with the Tivoli Storage Manager is not case-sensitive.

Most problems arising from password authentication can be attributed to the connection between the Tivoli Storage Manager server and the LDAP directory server.

Before you can use the LDAP-authenticated password, you must configure the LDAP directory server to be in synch with the Tivoli Storage Manager server. Make sure that the access control list on the LDAP directory server grants full authority to one user over the base distinguished name (Base DN). That user cannot be a part of that Base DN. For example, if your Base DN is `cn=tsmdata`, the user who administers the passwords and IDs must not have their distinguished name in that Base DN.

Related reference:

“Resolving Secure Sockets Layer errors” on page 96

Verifying the configuration for password authentication

You must ensure that you properly configured the Tivoli Storage Manager server and the LDAP directory server.

Make sure that you complete the following steps before you set up the authentication to an LDAP directory server for any administrator or node:

1. Open the `dsmserv.opt` option file and add the `LDAPURL` option, which contains the server and the base distinguished name (Base DN).

```
LDAPURL ldap://server.dallas.outbound/cn=project_x
```

Remember: Each URL must begin with `ldap://`

2. Issue the **SET LDAPUSER** command to designate a user who can change or reset passwords. The **LDAPUSER** is the administrator for the LDAP directory server space that is used for Tivoli Storage Manager passwords.
3. Designate a password for the **LDAPUSER** by issuing the **SET LDAPPASSWORD** command.

The characters that you can use for a password are listed here:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z  
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z  
0 1 2 3 4 5 6 7 8 9  
~ ! @ # $ % ^ & * _ - + = ` | ( ) { } [ ] : ; < > , . ? /
```

Important: If you are issuing a **REGISTER NODE** or **REGISTER ADMIN** command, place special characters within quotation marks.

```
register admin JohnDoe "pa$$w0rd<new>" authentication=ldap
```

Issue one of the following commands to determine which nodes or administrator IDs are configured to use passwords that authenticate with an LDAP directory server:

```
query node * authentication=ldap  
query admin * authentication=ldap
```

Example: Changing the authentication method

You can change the authentication method for an administrator ID or node by issuing an **UPDATE** command. For example:

```
update node node_x n0de_x|n0de_y authentication=ldap  
update admin ralph SeCre#Tw0rd authentication=ldap
```

Resolving problems with the LDAP directory server

If you are having problems with password authentication, verify that you completed all of the configuration steps correctly. Have you defined the base distinguished name (Base DN) on the LDAP directory server? Did you set the **LDAPURL** option?

After you install the IBM Tivoli Storage Manager server V6.3.3 and later, you must configure the LDAP directory server to communicate with the Tivoli Storage Manager server.

If you have connection problems, complete the following steps with an LDAP utility such as `ldapsearch` or `ldp.exe`:

1. Test forward- and reverse-DNS lookup of the LDAP server system on the Tivoli Storage Manager server system.

2. Test the network connection between the Tivoli Storage Manager server operating system and the LDAP directory server operating system.
3. Connect to the LDAP directory server with the host name and port that you specified in the **LDAPURL** option.
4. Establish a Transport Layer Security (TLS) connection by issuing the **StartTLS** option.
5. Use simple bind authentication to authenticate with the parameters that you defined for the **LDAPUSER** and the **LDAPPASSWORD**.
6. Search the LDAP directory server for the BaseDN that you specified in the **LDAPURL** option.

An LDAP server administrator might use the **ldapsearch** utility, as follows, to troubleshoot LDAP directory authentication problems:

Using OpenLDAP (specify the certificate file using the **TLS_CACERT option in the **ldap.conf** file)**

Without SSL/TLS

```
ldapsearch -H <hostname>
-D <LDAPUSER> -W -s base -b
<BaseDN from LDAPURL> -v -x objectclass="*"

```

With SSL/TLS

```
ldapsearch -H <hostname>
-D <LDAPUSER> -W -s base -b
<BaseDN from LDAPURL> -v -x -ZZ objectclass="*"

```

Using the IBM Tivoli Directory Server LDAP client (installed with AIX or downloaded from ibm.com)

Without SSL/TLS

```
ldapsearch -h <hostname>
-D <LDAPUSER> -w ? -s base -b
<BaseDN from LDAPURL> -v objectclass="*"

```

With SSL/TLS

```
ldapsearch -h <hostname>
-D <LDAPUSER> -w ? -s base -b
<BaseDN from LDAPURL> -v -Y -x -K "cert.kdb" objectclass="*"

```

For the previous commands, the following parameters apply:

- *hostname* = the URL from the **LDAPURL** option, for example
`ldap://ldap.ibm.com:389/`
- *LDAPUSER* = the parameters from the **SET LDAPUSER** command, for example
`cn=tsmsserver,cn=users,dc=ibm,dc=com`
- *BaseDN from LDAPURL* = the Base DN from the **LDAPURL** option, for example
`"OU=tsm,DC=ibm,DC=com"`

See the *Tivoli Storage Manager Administrator's Guide* for more information about how to configure the LDAP directory server.

Resolving problems with locked nodes and administrators

The passwords that authenticate with the LDAP directory server can be locked if you go over the limit on wrong passwords or by administrator actions.

If you cannot unlock a locked password, try the following steps:

1. Return the password to the Tivoli Storage Manager server by issuing the following example command:

```
update node node_x new_pw authentication=local
```

2. Clean out the LDAP directory server by issuing the following example command:

```
audit ldapdirectory fix=yes wait=no
```

This command removes nodes or administrator IDs that are stored on the LDAP directory server that do not authenticate passwords with an LDAP directory server.

3. Log-out from the node.
4. Issue the following command:

```
update node node_x newest_pw authentication=ldap
```
5. Log-on to the node with the new password.

Auditing the LDAP directory server to clean up the server

Keeping the LDAP directory server in synch with the Tivoli Storage Manager server makes it easier to know what you are working with. An LDAP directory server might have hundreds of entries that are no longer used. The LDAP directory server might also lack certain admin or node entries that are supposed to be on the LDAP directory server.

An audit can report to you on which administrator ID or node entries are recorded as authenticating their passwords with the LDAP directory server. You can audit the LDAP directory server to clear out unused passwords, administrators, and nodes. The Tivoli Storage Manager controlled namespace on the LDAP directory server can get out-of-sync with what the Tivoli Storage Manager server is storing.

If the administrator for the LDAP directory server manually changed entries on the external directory, those entries would not be in-sync. The Tivoli Storage Manager server can also get out-of-sync with the LDAP server when you use the default **SYNCLDAPDELETE=NO** command during a **REMOVE**, **RENAME**, or **UPDATE** command. The **AUDIT LDAPDIRECTORY** command deletes any entries from the LDAP directory server that do not correlate to the Tivoli Storage Manager database. The command also issues warnings to help you fix items.

Warnings are issued if passwords that authenticate with the LDAP directory server are stored in the Tivoli Storage Manager database but not in the LDAP namespace. From the warnings, you can use the **UPDATE NODE** or **UPDATE ADMIN** command to correct the issue.

Example: Auditing the LDAP directory server

If the Tivoli Storage Manager namespace on the LDAP directory server is not in-sync with the Tivoli Storage Manager database, issue the following command:

```
AUDIT LDAPDIRECTORY FIX=YES
```

The command produces a list of all nodes and administrators that are removed from the LDAP directory server. Also produced is a list of all nodes and

administrators that are missing from the LDAP directory server. If you want to see what is out-of-sync, use the default **FIX=NO** setting to report on the discrepancies between the servers.

Note: Do not use the **FIX=YES** setting if multiple Tivoli Storage Manager servers share LDAP directory namespace.

Error messages for LDAP authenticated passwords

When you authenticate passwords with an LDAP directory server, common errors can occur over the connection between the Tivoli Storage Manager server and the LDAP directory server.

The following table contains errors that you might find if your configuration is not correct:

Table 7. Errors that might occur when you authenticate passwords with an LDAP directory server

Error messages	Resolution
ANR3114E LDAP error 118 (The SSL library cannot be loaded) ANR3116E LDAP SSL/TLS error 118 (Unknown SSL error) ANR3103E Failure occurred while initializing LDAP directory services	The library path might not be set properly. Make sure that you are using the correct version of the Global Security ToolKit (GSKit). See the <i>Installation Guide</i> topic, "Creating the server instance."
ANR3114E LDAP error 116 (Failed to connect to the SSL server) ANR3116E LDAP SSL/TLS error 406 (I/O error) ANR3103E Failure occurred while initializing LDAP directory services ANR2732E Unable to communicate with the LDAP directory server	The level of GSKit might be incorrect on the Tivoli Directory Server. Upgrade GSKit to the correct level. See the technote (http://www.ibm.com/support/docview.wss?uid=swg21469388). For Active Directory, disable automatic root certificates updates with Windows Update if an internet connection is not available. See the following websites: Windows 2003 automatic certificate updates (http://technet.microsoft.com/en-us/library/cc738920%28v=WS.10%29.aspx) Windows 2008 automatic certificate updates (http://technet.microsoft.com/en-us/library/ee126107%28v=WS.10%29.aspx)
ANR3114E LDAP error 52 (DSA is unavailable) ANR3103E Failure occurred while initializing LDAP directory services ANR2732E Unable to communicate with the LDAP directory server	The Active Directory server does not have a certificate available for Transport Layer Security (TLS)/SSL. Create a signed certificate that can be used by Microsoft Active Directory.

Table 7. Errors that might occur when you authenticate passwords with an LDAP directory server (continued)

Error messages	Resolution
<p>ANR3114E LDAP error 116 (Failed to connect to SSL server)</p> <p>ANR3116E LDAP SSL/TLS error 414 (Bad certificate)</p> <p>ANR3103E Failure occurred while initializing LDAP directory services</p> <p>ANR2732E Unable to communicate with the LDAP directory server</p>	<p>The LDAP directory server certificate is not trusted. Add the root certificate authority (CA) certificate to the Tivoli Storage Manager server key database file (cert.kdb) and verify that the certificates are not expired.</p>
<p>ANR3094E The distinguished name (DN) that is specified in the LDAPURL option does not exist on the LDAP directory server</p> <p>ANR3103E Failure occurred while initializing LDAP directory services</p>	<p>If the DN exists, the LDAPUSER might not have full access control rights to the Base DN that is specified in the LDAPURL option.</p>
<p>ANR3114E LDAP error 50 (Insufficient access)</p> <p>ANR1885E LDAP directory service initialization: Permission was denied when the LDAP directory entry was accessed as LDAPUSER</p> <p>ANR3103E Failure occurred while initializing LDAP directory services</p> <p>ANR1885E SET LDAPPASSWORD: Permission was denied when the LDAPUSER entry was accessed</p>	<p>The LDAPUSER does not have full access control rights to the base DN that is specified in the LDAPURL option.</p>
<p>ANR3114E LDAP error 116 (Failed to connect to SSL server)</p> <p>ANR3116E LDAP SSL/TLS error 420 (Socket closed)</p>	<p>For Tivoli Directory Server, the SSL_TIMEOUT_MILLISEC is not set high enough. See the technote (http://www-01.ibm.com/support/docview.wss?rs=767&uid=swg21233758).</p>
<p>ANR3114E LDAP error 4 (Size limit exceeded)</p>	<p>Increase the LDAP server search size limit to accommodate the total number of LDAP-authenticated nodes and administrators.</p>

Resolving client scheduling problems

The administrator for the IBM Tivoli Storage Manager can schedule tasks to run automatically.

If you are experiencing problems with your client scheduler, the following diagnostic steps are available to help you determine the cause of the problem:

- Additions and changes to the Tivoli Storage Manager client options are not recognized by the client scheduler until the next scheduled start. Deletions made to the client option set do not take effect until you restart the scheduler.

- Additions, deletions, and changes made to the client-acceptor-managed schedules are recognized at the next scheduled start.
- Use the **SHOW PENDING** diagnostic tool to show schedules, nodes, and the next scheduled run time.
- From the client options file, view the `dsm.sys` stanza for the node and the `MANAGEDSERVICES`, `PRESCHEDCMD`, and `POSTSCHEDCMD` option values for information after a node misses a scheduled event.

Determining the status of a scheduled event

The server maintains a record of all scheduled events. The records are useful for managing IBM Tivoli Storage Manager schedules on numerous client computers.

Perform the following steps to view the event records on a server:

1. Issue the **QUERY EVENT** command.
2. Issue the following query to view all of the event results for the previous day:


```
query event * * begindate=today-1 begintime=00:00:00
enddate=today-1 endtime=23:59:59
```
3. Issue the following query to limit the query results to exception cases:


```
query event * * begindate=today-1 begintime=00:00:00
enddate=today-1 endtime=23:59:59 exceptiononly=yes
```

The query results include a status field that gives a summary of the result for a specific event. By using the `format=detailed` option you can also see the result of an event that is the overall return code passed back by the Tivoli Storage Manager client. See the **QUERY EVENT** command for scheduled and completed events.

Checking for errors in the server activity log

If a scheduled event is missed but other consecutive scheduled events for that node show a result of Completed, check for errors in the server activity log and the client schedule log.

When checking the server activity log, narrow the query results down to the time frame surrounding the scheduled event. Begin the event log query at a time shortly before the start window of the scheduled event in question. For example, investigate the following suspect event:

```
Scheduled Start Actual Start Schedule Name Node Name Status
-----
08/21/2003 08:27:33 HOURLY NODEA Missed
```

Afterwards you could issue one of the following queries:

```
query actlog begind=08/21/2003 begint=08:25:00
query actlog begind=08/21/2003 begint=08:25:00 originator=client node=nodea
```

The IBM Tivoli Storage Manager client keeps a detailed log of all scheduled activities. Check the client's local schedule log if queries of the server's activity log cannot explain a failed scheduled event.

You must have access to the client computer in order to inspect the schedule log file. The schedule log file is typically stored in the same directory as the `dsmerror.log` file in the `dsmsched.log` file. The location of the log file can be specified using client options, so you might need to refer to the options file to see if the `SCHEDLOGNAME` option was used to relocate the log file. On Windows, the schedule log can also be relocated by an option setting which is part of the schedule service definition. You can issue the **DSMCUTIL QUERY** command to check if

this option was set. When you locate the schedule log, search through the file to find the time period corresponding with the start date and time of the scheduled event in question. The following list displays common search parameters:

- If you are investigating a missed event, check the details of the previous event, including the time at which the previous event completed.
- If you are investigating a failed event, look for error messages that explain the failure (such as the server session limit being exceeded).
- When an explanation is still not clear, the last place to check is the client's error log file (usually named `dsmerror.log`).

Starting and stopping the client service

Starting and stopping the client service can sometimes help to resolve client scheduling problems.

Tip: When you manage a large number of clients running scheduler processes, you also might want to be able to start and stop the client service from a remote computer. The client for Windows provides a utility to assist with remote management of the scheduler service. For other platforms, standard operating system utilities are required.

To remotely manage the Windows client scheduler service using **DSMCUTIL** with the `/computer:` option, you must have administrative rights in the domain of the target computer. To determine whether the scheduler service is running on a remote computer, check the “Current Status” field from a query similar to the following query:

```
dsmcutil query /name:"TSM Client Scheduler" /computer:ntserv1.ibm.com
```

Issue the following queries to restart a scheduler service that is missing schedules:

```
dsmcutil stop /name:"TSM Client Scheduler" /computer:ntserv1.ibm.com
dsmcutil start /name:"TSM Client Scheduler" /computer:ntserv1.ibm.com
```

Consequently, if you use the client acceptor daemon (CAD) to manage the scheduler, you might have to restart the CAD service or stop the scheduler service and restart the CAD service with the following queries:

```
dsmcutil query /name:"TSM Client Scheduler" /computer:ntserv1.ibm.com
dsmcutil query /name:"TSM Client Acceptor" /computer:ntserv1.ibm.com
dsmcutil stop /name:"TSM Client Scheduler" /computer:ntserv1.ibm.com
dsmcutil stop /name:"TSM Client Acceptor" /computer:ntserv1.ibm.com
dsmcutil start /name:"TSM Client Acceptor" /computer:ntserv1.ibm.com
```

AIX **HP-UX** **Linux** **Solaris** For AIX, HP-UX, Linux, or Oracle Solaris, you can write a shell script to search for and stop running IBM Tivoli Storage Manager schedulers or client acceptor processes, and then restart the processes. Software products such as Symark's Power Broker allow Tivoli Storage Manager administrators limited access to servers for the purpose of managing the scheduler processes, and copying off the schedule log file.

AIX **HP-UX** **Linux** **Solaris** The following example shell script shows you how to recycle the Tivoli Storage Manager scheduler process:

```
#!/bin/ksh
# Use the following script to kill the currently running instance
# of the TSM scheduler, and restart the scheduler in nohup mode.
#
# This script will not work properly if more than one scheduler
# process is running.
```

```

# If necessary, the following variables can be customized to allow an
# alternate options file to be used.
# export DSM_DIR=
# export DSM_CONFIG=
# export PATH=$PATH:$DSM_DIR
# Extract the PID for the running TSM Scheduler
PID=$(ps -ef | grep "dsmc sched" | grep -v "grep" | awk {'print $2'});
print "Original TSM scheduler process using PID=$PID"
# Kill the scheduler
kill -9 $PID
# Restart the scheduler with nohup, redirecting all output to NULL
# Output will still be logged in the dsmsched.log
nohup dsmc sched 2>&1 > /dev/null &
# Extract the PID for the running TSM Scheduler
PID=$(ps -ef | grep "dsmc sched" | grep -v "grep" | awk {'print $2'});
print "New TSM scheduler process using PID=$PID"

```

Resolving errors when including or excluding client files during backup processing

The include-exclude processing option impacts which files are sent to the server for a backup or archive operation. Several reasons are possible if you implicitly or explicitly indicate that a file is included or excluded during backup processing and it is not processed correctly.

Identifying files that are included or excluded by the server client option set

The IBM Tivoli Storage Manager administrator can include or exclude files on behalf of the client. Include or exclude statements that come from the server will override include and exclude statements that are entered in the local client option file.

Contact the Tivoli Storage Manager server administrator to correct the problem.

You can issue the backup-archive client **DSMC QUERY INCLEXCL** command to identify the files that are included or excluded by the server client options set. The output from this command shows “Operating System” as the source file for files that were automatically excluded from backup processing. In our example, the users indicate that they want all files that end with a “.o” extension to be included in the local options file, but the server sends the client an option to exclude all files that end with a “.o” extension. The server-provided option prevails.

```

tsm> q inclexcl
*** FILE INCLUDE/EXCLUDE ***
Mode Function Pattern (match from top down) Source File
-----
Excl All /.../*.o Server
Incl All /.../*.o dsm.sys

```

Options that are passed to the Tivoli Storage Manager client from the Tivoli Storage Manager server are provided in groups, meaning that if the INCLUDE and EXCLUDE options are supported on the Tivoli Storage Manager server, that all INCLUDE options would be sent in a group and all EXCLUDE options would be sent in a group. You could not intermix these options to get wanted results of including some files from excluded directories. Using the INCLEXCL option allows you to intermix and order the INCLUDE and EXCLUDE options.

Excluding files automatically from backup processing

The backup application does not back up particular files because they are not necessary for backup, or IBM Tivoli Storage Manager uses the files for internal processing.

If particular files must be included in the backup processing, Tivoli Storage Manager can include them if you put *INCLUDE* statements in the client options set on the server.

Important: Because some files were explicitly identified as files not being backed up, including them in the server client options set is not recommended.

Issue the backup-archive client **DSMC QUERY INCLEXCL** command to identify the files that were not backed up. The output from the **DSMC QUERY INCLEXCL** command shows "Operating System" as the source file for files that were automatically excluded from backup processing.

```
tsm> q inclexcl
*** FILE INCLUDE/EXCLUDE ***
Mode Function Pattern (match from top down) Source File
-----
Excl All C:\WINDOWS\Registration\*.clb Operating System
Excl All C:\WINDOWS\netlogon.chg Operating System
```

See Table 8 on page 21 for the files that are automatically excluded.

Table 8. Files automatically excluded during backup processing

Platform	Files Excluded
<p>Windows</p>	<ul style="list-style-type: none"> • Files enumerated in the HKLM\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup registry key • The client staging directory C:\ADSM.SYS • Removable storage manager (RSM) database files (these files are processed in the system object or system state backup) • Internet information server (IIS) metafiles (these files are processed in the system object or system state backup) • Registry files (these files are processed in the system object or system state backup) • Client trace files • System files <p>Windows system files are silently excluded from the system drive backup processing and cannot be included.</p> <p>To process these Windows system files, you must issue a DSMC BACKUP SYSTEMOBJECT command (Windows 2000 and Windows XP) or a DSMC BACKUP SYSTEMSTATE command (Windows 2003 and Windows Vista).</p> <p>The Windows system files are excluded from the system drive backup processing because they are usually sent during the system object or system state backups. System files are boot files, catalog files, performance counters, and files protected by the Windows system file protection (sfp). These files are not processed during backup of the system drive. However, the files are excluded from the system drive processing internally instead of relying on explicit exclude statements due to the sheer number of exclude statements that would be needed to represent all of these files. Backup performance can be adversely affected.</p> <p>You can issue the backup-archive client DSMC QUERY SYSTEMINFO command to identify the Windows system files. The output of this command is written to the dsminfo.txt file.</p> <pre>(partial contents of the dsminfo.txt file) ===== SFP c:\windows\system32\ahui.exe (protected) c:\windows\system32\apphelp.dll (protected) c:\windows\appatch\apphelp.sdb (protected) c:\windows\system32\asycfilt.dll (protected)</pre>
<p>AIX HP-UX</p> <p>Linux Solaris</p>	<p>Client trace file</p>
<p>Macintosh</p>	<ul style="list-style-type: none"> • Volatile, temporary, and device files used by the operating system • Client trace files

Excluding files with the EXCLUDE.DIR statement

EXCLUDE.DIR statements exclude all directories and files under the parent directory.

If you want to include all files based on a file pattern, regardless of their location within a directory structure, do not use EXCLUDE.DIR statements.

AIX **HP-UX** **Linux** **Solaris** For example, consider this set of include-exclude statements:

```
exclude.dir /usr
include /.../*.o
```

The INCLUDE statement in this example indicates that all files with a “.o” extension should be included, but the preceding EXCLUDE.DIR statement will exclude all files in the /usr directory, even if they have a “.o” extension. This fact would be true, regardless of the order of the two statements.

If you want to back up all the files ending with “.o” use the following syntax:

```
exclude /usr/.../*
include /.../*.o
```

When using wildcards in include-exclude, use “*” if you want all the files rather than “*.*.*,” which means to include-exclude all files containing at least one dot (.) character, while * means to include-exclude all files. If you use “*.*”, files containing no dot characters (such as C:\MYDIR\MYFILE), the file is not filtered.

If you want to perform a selective backup of a single file from the command-line client, it is not affected by the EXCLUDE.DIR option.

If you issue a selective backup from the command-line client of a single file, the file is processed, even if there is an EXCLUDE.DIR statement which excludes one of the parent directories.

AIX **HP-UX** **Linux** **Solaris** For example, consider the following include-exclude statement that is used in subsequent command-line actions:

```
exclude.dir /home/spike
```

The following selective backup always results in the file being processed:

```
dsmc selective /home/spike/my.file
```

If you issue a selective backup using a wildcard, no files are processed because the directory is excluded:

```
dsmc selective "/home/spike/my.*"
```

Important: A subsequent incremental backup of the /home file system renders inactive the “/home/spike/my.file” file.

EXCLUDE.DIR statements should not be terminated with a directory delimiter.

The following examples show incorrect EXCLUDE.DIR statements, due to a terminating directory delimiter:

```
exclude.dir /usr/ (AIX, HP-UX, Linux, Oracle Solaris)
exclude.dir c:\directory\ (Windows)
exclude.dir Panther>User: (Macintosh)
```

The following examples show the correct coding of EXCLUDE.DIR:

```
exclude.dir /usr (AIX, HP-UX, Linux, Oracle Solaris)
exclude.dir c:\directory (Windows)
exclude.dir Panther:User (Macintosh)
```

Determining whether compression, encryption, and subfile backup statements include or exclude

Include and exclude statements for compression (INCLUDE.COMPRESS), encryption (INCLUDE.ENCRYPT), and subfile backup (INCLUDE.SUBFILE) do not imply that the file will be included for backup processing.

You can use the INCLUDE and EXCLUDE statements in combination with the COMPRESS, ENCRYPT, and SUBFILE statements to produce your wanted results.

AIX **HP-UX** **Linux** **Solaris** Consider the following example:

```
exclude /usr/file.o
include.compress /usr/*.o
```

This statement indicates that the /usr/file.o file is excluded from backup processing. The INCLUDE.COMPRESS statement indicates that “if a file is a candidate for backup processing and matches the pattern /usr/*.o; then compress the file.” The INCLUDE.COMPRESS statement should not be interpreted as “backup all files that match the pattern /usr/*.o and compress them.” If you want to back up the /usr/file.o file in this example, you must remove the exclude statement.

Using delimiters to include or exclude files

When the volume or directory delimiters are not correct, it might cause *INCLUDE* and *EXCLUDE* statements to malfunction.

A platform-specific INCLUDE or EXCLUDE statement contains syntax for “everything” and “all files under a specific directory.”

If you want to use an *INCLUDE* statement for “all files under a specific directory,” ensure that the slashes and volume delimiters are correct. If you want to exclude all of the files under a directory called “home,” or simply all files, see the following examples:

Using the backwards slash “\” and the volume delimiter “:” (Windows)

```
*include everything in the c:\home directory
include c:\home\...\*
*include everything
include *:\...\*
```

Using the forward slash “/” (AIX, HP-UX, Linux, Oracle Solaris, and Macintosh OS X)

```
*include everything in the /home directory
include /home/...\*
*include everything
include /...\*
```

Resolving errors due to the incorrectly coded include or exclude list

Due to the complexity or number of INCLUDE or EXCLUDE statements, you might experience the unintentional inclusion or exclusion of a file.

Configure the client with the **INCLEXCL** trace flag to help determine why a file was included or excluded.

For example, when you believe that the `c:\home\file.txt` file should be included in the backup processing. The trace shows that there is an EXCLUDE statement that excludes this file:

```
polbind.cpp (1026): File 'C:\home\file.txt' explicitly excluded by pattern  
'Excl All c:\home\*.txt'
```

Using the backup-archive client **DSMC QUERY INCLEXCL** command shows that this statement is in the IBM Tivoli Storage Manager server client options set:

```
tsm> q inclexcl  
*** FILE INCLUDE/EXCLUDE ***  
Mode Function Pattern (match from top down) Source File  
-----  
Excl All c:\home\*.txt Server
```

Resolving Snapshot Difference problems

You can perform faster incremental backups of N-Series and NetApp Filer volumes if you use the NetApp Snapshot Difference application programming interface (API).

Prerequisites

To use the Snapshot Difference feature, you must first set up a NetApp user ID and password on the IBM Tivoli Storage Manager client. The user ID and password are necessary for Tivoli Storage Manager to connect to the Filer. Set up a user ID/password with root authority for AIX, or one with administrative authority for Windows. Set the authority level to be the same as the authority level used when you map or mount the filer volume. Ensure that you use the fully-qualified host name or the dotted IP address format for the Filer name. Issue the backup-archive client **SET PASSWORD** command to save this user ID/password information.

Remember: The `dsmc SET PASSWORD` command is extended to save “Filer” type passwords.

The Snapshot Difference feature compares two snapshots (base and differential) and returns a list of files that were modified, deleted, or added between the two. Tivoli Storage Manager backs up this list of files instead of scanning the file system for changes.

The Snapshot Difference feature supports the following features, which are only applicable at the volume level:

- NetApp/N-Series filers running Data ONTAP release 7.3 or later
-  Common internet files system-attached (CIFS) volumes
- Both traditional and FlexVol filer volumes
- Java and web client GUI

- **AIX** Network file system (NFS) attached volumes

The Snapshot Difference feature does not support the following features:

- SAN-attached NetApp/N-Series volumes
- Native Windows GUI support (no MFC GUI)
- QTrees or subdirectories
- Vfiler

Windows

Verifying the Filer volume type

Tivoli Storage Manager expects the Common Internet Files System-attached (CIFS) security type to be Windows NT file system (NTFS). Use the NetApp FilerView and make sure that the CIFS security type is set to “ntfs.”

Snapshot Difference restrictions

The lack of Unicode support from NetApp prevents Tivoli Storage Manager from processing any files that use characters that are not in the 7-bit ASCII range. Tivoli Storage Manager can only back up names containing ASCII characters. Two Snapshot Difference behaviors were noted when testing with Unicode characters:

1. Snapshot Difference incremental command ends with return code 13001. This return code happens with the 'specials' and 'surrogate' ranges of Unicode for Snapshot Difference Filer volumes created with the UTF8 flag. This Snapshot Difference error happens more frequently without the UTF8 flag. Tivoli Storage Manager ends with error message ANS5283E “The operation was unsuccessful.” No files are backed up.
2. Snapshot Difference application programming interface (API) does not fail, but returns characters that are not part of the real name. Tivoli Storage Manager inspects the string to see if any character is outside of the 7-bit ASCII range. If so, Tivoli Storage Manager skips the file and logs the error to the dsmerror.log file.

The following are situations under which files/directories might not get backed up and no errors are reported:

- You exclude a file by adding an exclude rule in the include/exclude file. Tivoli Storage Manager performs a backup of the current snapshot with that exclude rule in effect. You did not change the file, but you do remove the rule that excluded the file. A snapshot-assisted incremental backup command with the `snaptdiff` option does not detect this incl/excl change because it only detects file changes between two snapshots. The files themselves have to be changed in order for the Snapshot Difference API to detect the change and for Tivoli Storage Manager to back up the file.
- You have added an include statement to the option file. Only if the file is detected to have changed by the Snapshot Difference API can that include option take effect. The files might not get backed up because Tivoli Storage Manager is not inspecting each file on the volume during backup.
- You explicitly delete a file from Tivoli Storage Manager inventory by issuing the **DSMC DELETE BACKUP** command. The Snapshot Difference API does not detect that a file was manually deleted from Tivoli Storage Manager by you. Therefore, the file remains unprotected in Tivoli Storage Manager storage. The file is unprotected until it is changed on the volume and the change is detected by the Snapshot Difference API. After the change is detected, the Snapshot Difference API signals Tivoli Storage Manager to, once again, back it up.

- Policy changes such as changing the policy from Mode=modified to mode=absolute are not detected. The entire file space is deleted from Tivoli Storage Manager inventory. The undetected policies cause Tivoli Storage Manager to create a snapshot to use as the source (base) and a full incremental backup is performed.

Running a full incremental backup without the `snappdiff` option solves these limitations. Tivoli Storage Manager does not control what constitutes a changed object. The changing of objects is now controlled by the Snapshot Difference API. Therefore, running a full incremental backup without the `SNAPDIFF` option ensures that all file changes are detected.

Trace flags that you can use for Snapshot Difference:

- `enter`
- `exit`
- `general`
- `snapshot`
- `hci`
- `hci_detail`
- `diskmap`
- `diskmap_detail`
- `hdw`
- `hdw_detail`
- `bacache`
- `snappdiffdb`

Set up a user ID and password for root on the filer `myFiler.ibm.com`.

```
dsmc set password -type=filer myFiler.ibm.com root
```

```
Please enter password for user id "root@myFiler.ibm.com": *****
Re-enter the password for verification:*****
ANS0302I Successfully done.
```

Set up a user ID and password for root on the filer `myFiler.ibm.com`.

```
dsmc set password -type=filer myFiler.ibm.com root secret
```

Resolving snapshot directory problems for NetApp or N-Series file system volumes

When a network file system (NFS) mounted or a Common Internet File System (CIFS) mapped volume is backed up, so are all snapshots within the snapshot directory. This backup includes unwanted snapshots that can occupy valuable space. The NFS-mounted or CIFS-mapped volumes can be either NetApp or N-Series.

To avoid backing up unwanted snapshots, use the Network Data Management Protocol (NDMP) backup method. You can also back up your data with the client `SNAPSHOTROOT` option or run an incremental backup with the **INCREMENTAL** command and the `SNAPDIFF` option. Alternatively, exclude the snapshot directory from any backup.

Resolving login problems when using the encrypted file system on AIX operating systems

AIX

During login processing, the encrypted file system (EFS) keystore opens automatically when the keystore password matches the user login password.

When the login password for AIX is different from the EFS keystore password, you must open the keystore manually before you start the IBM Tivoli Storage Manager client. Open the keystore by issuing the following command:

```
efskeymgr -o <cmd>
```

Start the Tivoli Storage Manager client in one of the following ways:

- Start the command-line client by issuing the `efskeymgr -o ./dsmc` command.
- Start the Java GUI client by issuing the `efskeymgr -o ./dsmj` command.

If you are using the Tivoli Storage Manager client web graphical user interface (GUI), you must synchronize the passwords. To synchronize the user password with the EFS keystore password, issue the following command:

```
efskeymgr -n
```

Resolving image backup errors

Image backup errors can occur with Linux images, Linux Snapshot images, or during AIX JFS2 Snapshot-based backup-archive and image backup.

Resolving Linux image backup errors

Linux

You can resolve Linux image backup errors by performing specific steps, depending on the type of error that occurs.

The following error was generated during image backup:

```

paris:#dsmc b image /dev/system/lv01
Backup Image Function Invoked.
ANS1228E Sending of object '/dev/system/lv01' failed
ANS1584E Error loading system library 'libdevmapper.so'
required for image operations for LVM2 volumes.
ANS1813E Image Backup processing of '/dev/system/lv01'
finished with failures.
Total number of objects inspected: 1
Total number of objects backed up: 0
Total number of objects updated: 0
Total number of objects rebound: 0
Total number of objects deleted: 0
Total number of objects expired: 0
Total number of objects failed: 1
Total number of bytes transferred: 0 B
Data transfer time: 0.00 sec
Network data transfer rate: 0.00 KB/sec
Aggregate data transfer rate: 0.00 KB/sec
Objects compressed by: 0%
Elapsed processing time: 00:00:29
paris# cat dsmerror.log
11/15/2006 13:07:53 ANS1228E Sending of object
'/dev/system/lv01' failed
11/15/2006 13:07:56 ANS1584E Error loading system
library 'libdevmapper.so' required for
image operations for LVM2 volumes.
11/15/2006 13:07:56 ANS1813E Image Backup processing
of '/dev/system/lv01' finished
with failures.

```

For this error, ensure that the system has the correct version of the library device mapper installed. Perform the following steps to determine the installed version:

1. Issue the # **DMSETUP VERSION** command. The output is similar to the following output:

```

Library version: 1.00.09-ioct1 (2004-03-31)
Driver version: 4.4.0

```

or

Issue the following command to determine the version using the rpm:

```
# rpm -q -a |grep device-mapper
```

The output is similar to the following output:

```
device-mapper-1.00.09-17.5
```

The library version must be Version 1.01 or later. If you have an earlier version, upgrade the device mapper rpm file at this website.

2. Verify the installation after the upgrade.

```

# rpm -Uvh device-mapper-1.01.01-1.6.i586.rpm
Preparing... ##### [100%]
1:device-mapper ##### [100%]
# rpm -q -a |grep device-mapper
device-mapper-1.01.01-1.6

```

You can also check the /lib directory to see that the correct versions are installed. A system with the correct levels would have the following information:

```

# ls -l /lib/libdev*
lrwxrwxrwx 1 root root 20 Jul 5 11:42 /lib/libdevmapper.so
->libdevmapper.so.1.01
-rwxr-xr-x 1 root root 24490 May 23 2005 /lib/libdevmapper.so.1.00
-rwxr-xr-x 1 root root 28216 May 23 2005 /lib/libdevmapper.so.1.01

```

Resolving image backup failures when using Linux snapshot

Linux

To resolve a failed Linux snapshot image backup, validate that the system is set up to create a snapshot.

Try to create a snapshot from a shell command prompt by issuing the following command:

```
/sbin/lvcreate -L 16384K -n <snapname eg. tsmsnap>-s  
<volume devname eg /dev/system/lv01>
```

If you receive the “Snapshot: Required device-mapper target(s) not detected in your kernel,” error, the **:dm_snapshot** kernel module is not loaded. This command could also fail for other reasons, which might result in similar IBM Tivoli Storage Manager behavior.

The following example shows the output generated when an image backup fails with error message ANS1258E, “The image snapshot operation failed.”

```
dsmerror.log :  
05/31/2006 15:14:36 ANS1259E The image snapshot operation failed.  
Diagnostic text: tsmStartSnapshot.  
05/31/2006 15:14:38 ANS1259E The image snapshot operation failed.  
Diagnostic text: tsmTerminateSnapshot.  
05/31/2006 15:14:38 ANS1228E Sending of object '/fs1' failed  
05/31/2006 15:14:38 ANS1258E The image snapshot operation failed.
```

Perform the following steps to load the modules:

1. Verify that the module is not loaded. See the following example command:

```
# lsmod |grep dm_  
dm_mod 112104 6
```

2. Load the module. See the following example command:

```
# modprobe dm_snapshot
```

3. Verify that the previous step is successful. See the following example command:

```
# lsmod |grep dm_  
dm_snapshot 44024 0  
dm_mod 112104 6 dm_snapshot  
#
```

4. Create a snapshot from the shell prompt. See the following example command:

```
# /sbin/lvcreate -L 16384K -n tsmsnap -s /dev/system/lv01  
Logical volume "tsmsnap" created
```

5. Remove the snapshot that was created in the previous step. See the following example command:

```
# lvremove /dev/system/tsmsnap  
Do you really want to remove active logical volume "tsmsnap"? [y/n]: y  
Logical volume "tsmsnap" successfully removed  
#
```

If you followed all of the steps, you might now be able to run snapshot image backups.

Restriction: If the **lvcreate** command fails with error “Insufficient free extents (0) in volume group...”, there is not enough space in the volume group for a snapshot volume.

Resolving errors during AIX JFS2 snapshot-based backup-archive and image backup

AIX

During IBM Tivoli Storage Manager termination, the Tivoli Storage Manager client deletes the AIX enhanced journaled file system (JFS2) snapshot that is created during the backup process. However, there are situations where AIX might fail the snapshot delete request made by Tivoli Storage Manager.

The following situations illustrate where a snapshot delete request might fail:

- The Control-c keystroke is issued during a Tivoli Storage Manager snapshot backup process. The JFS2 snapshot unmount request might fail with a “Device Busy” error, due to the Tivoli Storage Manager process being in the middle of accessing the snapshot.
- Two Tivoli Storage Manager snapshot backup requests are invoked concurrently for the same file system. For example, if the `dsmc backup image /fs1 backup` request is submitted from one console, and at the same time a `dsmc backup image /fs1 backup` request is issued from another console. If the process from the first console creates the first snapshot for /fs1 and the second process from the second console creates the second snapshot for /fs1, and if the second process finishes first and tries to delete the snapshot, AIX fails the delete request.
- Two Tivoli Storage Manager snapshot backup requests are invoked concurrently for two virtual mount points whose source file system is the same. For example, issuing `dsmc incr /fs1/level1/dir1` from one console and `dsmc incr /fs1/level2/level3/dir3` from a second console, concurrently.

AIX expects snapshot delete requests to be issued in a certain order, with the oldest snapshot deletion requested first and the next oldest snapshot deletion requested next, and so on. If Tivoli Storage Manager cannot honor the sequence due to concurrent processes creating snapshots for the same file system, AIX fails the delete requests. In the previous examples, Tivoli Storage Manager logs a warning message asking the user to delete the snapshots manually.

Issue the following commands, in order, to perform a manual snapshot deletion:

1. `snapshot -q -c ' ' <SRCFS>`
2. `df -k`
3. `umount -f /tsm*`
4. `rmdir /tsm*`
5. `snapshot -d /dev/tsm*` If the snapshot delete process fails with “Device Busy” or some other error message, issue the `umount -f <srcfs>` command to unmount the source file system. Retry snapshot delete.
6. `ls -l /dev/tsm*` If any /DEV/TSM* logical volumes remain, issue the `rmlv -f tsm*` command.
7. If you have an unmounted source file system, issue the `mount <srcfs>` command to mount it.

If any snapshots are not deleted during a previous Tivoli Storage Manager process, Tivoli Storage Manager tries to delete the snapshots during its next invocation because as older snapshots remain, AIX fails deletion requests for newer snapshots for a given file system. The following cases indicate where Tivoli Storage Manager does not try to delete older snapshots:

- If the snapshot was not created by Tivoli Storage Manager. Tivoli Storage Manager names its snapshots with a “tsm” prefix to distinguish them from other snapshots created for the same file system. If the snapshot was not created by Tivoli Storage Manager, an error message is generated that asks the user to delete the older snapshot and retry the operation.
- If the snapshot is created by Tivoli Storage Manager but is still mounted, the snapshot is being used by some other Tivoli Storage Manager process.
- If the snapshot is created by Tivoli Storage Manager, is not mounted, but is newly created, the snapshot might have just been created by some other Tivoli Storage Manager process.

In all such cases, you might have to perform a manual deletion. If any unused older snapshots are existing, subsequent Tivoli Storage Manager backups fail to delete snapshots.

Important: There are AIX defect fixes related to JFS2 snapshots in AIX 5.3.0.70 or later and AIX 6.1 or later. If the fixes are not applied, an AIX system shutdown can occur or Tivoli Storage Manager might stop during snapshot deletion and snapshot query processes. It might also cause data corruption during used-block-image backup. Therefore, Tivoli Storage Manager will not perform the following tasks:

- Snapshot monitoring
- Snapshot deletion
- Used-block-image backup (unless AIX is at the AIX 5.3.0.70 or later level)

In order to exploit these features, ensure that your operating system level is at AIX 5.3.0.70 or later and AIX 6.1 or later.

Support solutions for the Tivoli Storage Manager API

Resources are available to learn about or to diagnose the IBM Tivoli Storage Manager application programming interface (API).

API instrumentation is only activated if the `testflag INSTRUMENT: API` is set in the configuration file and the `dsmSetUp` and `dsmCleanUp` calls are used in the application.

See the *Using the Application Programming Interface* or <http://www.ibm.com/software/support/isa/> for more information.

Gathering API information before calling IBM support

You can significantly help to determine an application programming interface (API) problem by collecting information about your environment.

Gather as much of the following information as possible before contacting IBM Support:

- On what operating system is the problem being experienced?
- What is the exact level of the operating system, including all service packs and hot fixes that were applied?
- What is the exact level of the IBM Tivoli Storage Manager API?
- What is the exact level of the Tivoli Storage Manager server?
- What is the Tivoli Storage Manager server platform and operating system level?
- What is the exact level of the Tivoli Storage Manager storage agent (if LAN-free environment)?

- What is the Tivoli Storage Manager storage agent platform and operating system level (if LAN-free environment)?
- What applications are running on the system?
- What steps are required to recreate the problem? If you cannot recreate the problem, what steps caused the problem?

Gathering API files before calling IBM support

Log files and other important data are created by the IBM Tivoli Storage Manager application programming interface (API).

Gather as many of the following files before contacting IBM Support:

- Tivoli Storage Manager API error log file. The default API error log file is `dsierror.log`.
- Any trace files created for the API (the recommended trace flags are `api`, `api_detail`, or `verbdetail`).
- Output from any failed command or operation which might be either the console output redirected to a file or an actual screen image of the failure.
- The output from the Tivoli Storage Manager server **QUERY SYSTEM** command.
- Tivoli Storage Manager server activity log file. The Tivoli Storage Manager administrator can view this log file for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- If the API client is configured for LAN-free data movement, collect the options file for the Tivoli Storage Manager storage agent. The default name for this options file is `dsmsta.opt`.
- A short program or sections of the application source code invoking the Tivoli Storage Manager API function calls that are suspected to cause the problem.
- Tivoli Storage Manager API options file.

The following two options files are used on UNIX and OS/400 operating systems:

dsm.opt

The client options file

dsm.sys

The system options file

For Windows, find the `dsm.opt` default options file or the file referenced by the **DSMI_CONFIG** environment variable. For UNIX, the default options file is `dsm.sys` and is found in the directory referenced by the **DSMI_DIR** environment variable.

On other operating systems, the client options file (`dsm.opt`) contains all of the options. The following definitions are environment variables that describe the location of the option files and other API components:

DSMI_CONFIG

The fully-qualified name for the client options file.

DSMI_DIR

The `DSMI_DIR` variable points to the API installation directory and is also used to find the `dsm.sys` file (on UNIX). Wherever the `DSMI_DIR` is set, ensure that a `dsm.sys` file exists in the same directory.

DSMI_LOG

The `DSMI_LOG` variable points to the path for the `dsierror.log` file.

Tip: If this variable points to a directory for which the user does not have write permission, **dsmSetup** and **dsmInitEx** fail with return code DSM_RC_ACCESS_DENIED (106).

If the ERRORLOGNAME option is set in the options file (dsm.sys/dsm.opt), its value is used as the error log name instead of the default value dsierror.log.

Verifying that the API uses the correct option file

When you gather application programming interface (API) files, you must verify that the API uses the correct option file or server stanza in the dsm.sys file.

Perform the following steps to verify that the API uses the correct option file or server stanza:

1. Insert an erroneous option or value in the client option file or server stanza in dsm.sys. For example, if it is uncertain whether the API uses the srvr1.cmpron server, insert 'ERRONEOUS_OPTION 12345' into the srvr1.cmpron server stanza of the dsm.sys file. See the following example:

```
...
SERVERNAME srvr1.cmproff
COMPRESSION NO
TCPSERVERADDRESS computer.company.com

SERVERNAME srvr1.cmpron
COMPRESSION YES
ERRONEOUS_OPTION 12345
TCPSERVERADDRESS computer.company.com

SERVERNAME srvr1.pwdf1
PASSWORDACCESS GENERATE
PASSWORDDIR .
TCPSERVERADDRESS computer.company.com
...
```

2. Verify that the API detects the error. You can use the sample API program, **dapismp**, for this purpose.

```
# dapismp
...
Enter selection ==>0
Node name:node1
Owner name:
Password:
API Config file:
Session options:
User Name:
User pswd:
Are the above responses correct (y/n/q)?
Doing signon for node node1, owner, with password
*** Init failed: ANS0220E (RC400) An invalid option was found during option parsing.
```

If no error is reported, the wrong options file was updated.

3. Check the environment variable values that were previously mentioned or repeat steps 1 and 2 with a different options file/server stanza.
4. Remove the option inserted in step 1.

Determining if data is sent to the Tivoli Storage Manager storage agent rather than the server

You must know if your data is being sent to the IBM Tivoli Storage Manager storage agent, rather than a server. If the data is sent to the storage agent, you cannot recover it.

Perform the following steps to verify that data is being sent to the Tivoli Storage Manager storage agent, rather than to the server:

1. Add the following trace options to the client options file prior to backing up or archiving objects:

- TRACEFILE *<trace file name>*
- TRACEFLAGS api api_detail verbdetail

2. Examine the trace file after the operation and locate a statement that looks similar to the following statement:

```
dsmSendObj ENTRY:... objNameP: '<the file name>'
```

This statement is followed by the following trace statement:

```
tsmEndSendObjEx: Total bytes sent * *, encryptType is *** encryptAlg is  
*** compress is *, totalCompress is * * totalLFBytesSent * *
```

The trace statement indicates whether the object totalLFBytesSent was sent to the Tivoli Storage Manager storage agent. If totalLFBytesSent is 0 0, the data was sent directly to the Tivoli Storage Manager server.

Alternatively, your application itself can determine whether the data was sent through a LAN-free path by using the dsmEndSendObjEx function call and the dsmEndSendObjExOut_t data structure.

```
/*-----+  
| Type definition for dsmEndSendObjExOut_t  
+-----*/  
typedef      struct      dsmEndSendObjExOut_t  
{  
  dsUInt16_t   stVersion; /* structure version */  
  dsStruct64_t totalBytesSent; /* total bytes read from app */  
  dsmBool_t    objCompressed; /* was object compressed */  
  dsStruct64_t totalCompressSize; /* total size after compress */  
  dsStruct64_t totalLFBytesSent; /* total bytes sent LAN Free */  
  dsUInt8_t    encryptionType; /* type of encryption used */  
} dsmEndSendObjExOut_t;  
totalLFBytesSent - The total LAN-free bytes that were sent.
```

For example:

```
...  
    rc = dsmEndSendObjEx(&endSendObjExIn, &endSendObjExOut);  
    if (rc)  
    {  
        printf("*** dsmEndSendObjEx failed: ");  
        rcApiOut(dsmHandle, rc);  
    }  
    else  
    {  
        dI64toCh(&endSendObjExOut.totalLFBytesSent,t,10);  
        format_number(t,t2);  
        printf("LAN-free bytes sent: %s\n", t2);  
    }  
}
```

See *API Function Calls in Using the Application Programming Interface* for more details.

Running applications that use the API as a non-root user

AIX

HP-UX

Linux

Solaris

You must perform specific steps if you are a non-root user trying to run an application that uses the application programming interface (API).

Perform the following steps to allow a non-root user access to the API:

1. Set the **DSMI_CONFIG** environment variable. Verify that the non-root user has read permission for the client options file specified by **DSMI_CONFIG**. Otherwise, `dsmInit/dsmInitEx` fails with return code `DSM_RC_NO_OPT_FILE` (406). For example, the following options file is not readable by a non-root user, therefore the file permissions must be updated:

```
$ ls -l $DSMI_CONFIG
-rwx----- 1 root sys 86 Oct 7 13:07 /testfsapi/callmt_nr/dsm.opt
$ su root
Password:
# chmod a+r /testfsapi/callmt_nr/dsm.opt
# exit
$ ls -l $DSMI_CONFIG
-rwxr--r-- 1 root sys 86 Oct 7 13:07 /testfsapi/callmt_nr/dsm.opt
```

2. Set the **DSMI_DIR** environment variable to the API installation directory. Verify that the non-root user has read permission for the system options file specified by **DSMI_DIR/dsm.sys**.

```
$ export DSMI_DIR=/opt/tivoli/tsm/client/api/bin64
$ ls -l $DSMI_DIR/dsm.sys
-rw-r--r-- 1 root sys
4712 Oct 19 18:07 /opt/tivoli/tsm/client/api/bin64/dsm.sys
```

3. Set the **DSMI_LOG** environment variable. Verify that the non-root user has write permission for this directory. For example, the following **DSMI_LOG** directory is owned by a non-root user:

```
$ ls -ld $DSMI_LOG
drwxr-xr-x 2 apitest users 96 Oct 19 17:56 /testfsapi/callmt_nr/logs
```

If **PASSWORDACCESS GENERATE** is set in system options file `dsm.sys`, perform steps 4 and 5, otherwise go to step 6.

4. Check the ownership and permissions of the Trusted Communication Agent (TCA). This information is in the directory indicated by the **DSMI_DIR** environment variable. For example, the following TCA has the correct ownership and permissions:

```
$ ls -l $DSMI_DIR/dsmtca
-rwsr-xr-x 1 root bin 5021160 Oct 14 09:48
/opt/tivoli/tsm/client/api/bin64/dsmtca
```

Wrong permissions or ownership result in `DSM_RC_AUTH_FAILURE` (137) returned from `dsmInit`. Additionally, it is imperative that you use the same version of the API library and `dsmtca`. Mixed versions result in errors.

```
Error : calling program and dsmtca are not compatible
calling program build date : Mon Oct 18 21:15:59 2004 Mon Oct 18 21:15:59 2004
TCA build date : Wed Oct 13 16:48:03 2004 Wed Oct 13 16:48:03 2004
*** Init failed: ANS0282E (RC168) Password file is not available.
```

5. The root user must generate the TSM.PWD password file using either the IBM Tivoli Storage Manager backup-archive client or the `dapism` sample API application. Location of the password file is determined by the **PASSWORDDIR**

option in the `dsm.sys` system options file. In the following example, the sample API application generates the `TSM.PWD` password file for a node whose password is *oddesy*:

```
# dapismp
*****
* Welcome to the sample application for the Tivoli Storage Manager API. *
* API Library Version = 5.4.0.0 *
*****
Choose one of the following actions to test:
0. Signon
1. Backup
2. Restore
3. Archive
4. Retrieve
5. Queries
6. Change Password
7. Utilities : Deletes, Updates, Logevent, SetAccess, RetentionEvent
8. Set preferences, envSetUp
9. Exit to system
10. Restore/Retrieve Without Offset Prompt
11. Extended Signon
Enter selection ==>0
Node name:
Owner name:
Password:oddesy
API Config file:
Session options:
User Name:
User pswd:
Are the above responses correct (y/n/q)?
Doing signon for node, owner, with password oddesy
Handle on return = 1
Choose one of the following actions to test:
0. Signon
1. Backup
2. Restore
3. Archive
4. Retrieve
5. Queries
6. Change Password
7. Utilities : Deletes, Updates, Logevent, SetAccess, RetentionEvent
8. Set preferences, envSetUp
9. Exit to system
10. Restore/Retrieve Without Offset Prompt
11. Extended Signon
Enter selection ==>9
# ls -l TSM.PWD
-rw----- 1 root sys 121 Oct 19 18:28 TSM.PWD
Function call dsmInit returns DSM_RC_NO_PASS_FILE (168), if the password
file is not present in the directory specified by the PASSWORDDIR option.
```

6. If tracing is enabled, verify that the non-root user has write permission for the file indicated by issuing the `TRACEFILE` option.

Journal Based Backup problem determination

Journal Based Backup (JBB) is appropriate for backing up file systems with small or moderate amounts of change activity between backup cycles.

Determining if a backup will be journal-based

Before implementing a backup, you need to determine if it is going to be journal-based.

Perform the following steps to ensure that the backup is journal-based:

1. Configure the journal daemon to journal the file system that is being backed up. The journal daemon journalizes a file system after you list the file system in the `tsmjbbd.ini` configuration file. See the following configuration information:

```
[JournaledFileSystemSettings]
;
; List of journalized file systems
JournaledFileSystems=c:
```

2. Perform a full incremental backup on the corresponding file system while the file system is actively being journalized. This full incremental backup must set the “Last Backup Completed” date on the IBM Tivoli Storage Manager server file space in order for the journal to be set to valid. You can view the “Last Backup Completed” date by issuing the **QUERY FILESPACE** server command. After the journal is set to the valid state, subsequent backups by the same node to the same Tivoli Storage Manager server will be journal-based. If a backup uses a different node or a different server, the backup will be non-journal-based but the journal will remain valid for the original node and server, and backups to the original node and server will be journal-based. The following message is an example of what is written to the Windows Application Event Log when a journal is initially set to valid:

```
Journal set to valid for fs 'H:' and will be used for backup by
node GSHLAGER3 to server GSHLAGER2_SERVER1.
```

3. Ensure that the Tivoli Storage Manager node and server that the backup is using matches the node and server for which the journal is valid.
4. Use the Journal Database Viewing utility to determine the current state of a journal. If a valid journal is restarted, backups will be non-journal based until the journal is re-validated. The following message is written to the Windows Application Eventlog when a journal is restarted:

```
Journal database 'c:\tsmjjournal\tsmH__.jdb' for fs 'H:' has been
deleted and reset to the invalid state.
```

Restarting a valid journal

You can increase performance by restarting a valid journal.

The reasons for restarting a valid journal:

- Error conditions in the journal daemon
 - Buffer overflow errors caused by excessive change activity on the journal file system being monitored for changes
 - Journal database access errors (disk full errors, etc.)
- Request by a backup client
- Clients will issue a journal restart request when it is determined that a journal file system lacks integrity for one of the following reasons:
 - Server filesystem no longer exists
 - Server filesystem was deleted after the last backup
 - The node policy set was updated after the last backup
 - The Last Backup Completed or Last Backup Started dates are not valid (not set)

Running the journal daemon in the foreground

You can improve the diagnostic capabilities and your ability to test by running the journal daemon in the foreground, rather than as a Windows service.

Start the journal daemon from a Windows command prompt as follows:
tsmjbbd.exe i

Using the Journal Database Viewing utility

The Journal Database Viewing utility provides valuable information to help in problem determination.

The Journal Database Viewing utility provides the following information:

- The current state of the journal
- The file system tracked by the journal
- The journal activation time stamp
- The journal validation time stamp
- The maximum supported journal size
- The node and server for which the journal is valid
- The number of entries currently in the journal

Note: Use the viewing utility that matches your backup-archive client. For backup-archive clients V6.1.1 or later, you cannot view the contents of open journals with the viewing utility. An open journal is one that is currently opened by another process, such as the journal daemon. You can, however, view the contents of an open journal control record. The viewing utility comes with Tivoli Storage Manager backup-archive client V6.3.1 and later. For more information about the viewing utility, see the following technical articles:

- Dbviewb fails to retrieve information about the journal database (<http://www.ibm.com/support/docview.wss?uid=swg21296288>)
- Run the dbviewb.exe utility in batch mode (<http://www.ibm.com/support/docview.wss?uid=swg21248994>)

This utility also allows searching, inserting, or deleting specific entries in a journal database.

The syntax of this utility is:

```
dbviewb <fully qualified journal database basefile name>
dbviewb <fully qualified journal database basefile name> <i>
D:\tsm540c\debug\bin\winnt_unicode>dbviewb c:\tsmjjournal\tsmh__.jdb
IBM Tivoli Storage Manager
Journal Database Viewing Utility
Version 5, Release 4, Level 0.0
Last Update: Nov 28 2006
Querying Journal DB ...
Journal Database Information:
Database File c:\tsmjjournal\tsmh__.jdb
Database File Disk Size 81 KB (83754 Bytes)
Journal File System H:
Journal Activation Date Tue Nov 28 11:49:05 2006
Journal Validation Date Wed Nov 29 16:41:11 2006
Maximum Journal Size 8191 PB (9223372036854775807 Bytes)
Journal Type Change Journal
Journal State Valid
```

```

Valid for Server GSHLAGER2_SERVER1
Valid for Node GSHLAGER3
Number of DB Entries 22
D:\tsm540c\debug\bin\winnt_unicode>
D:\tsm540c\debug\bin\winnt_unicode>dbviewb c:\tsmjournal\tsmh__.jdb i
IBM Tivoli Storage Manager
Journal Database Viewing Utility
Version 5, Release 4, Level 0.0
Last Update: Nov 28 2006
Querying Journal DB ...
Journal Database Information:
Database File c:\tsmjournal\tsmh__.jdb
Database File Disk Size 81 KB (83754 Bytes)
Journal File System H:
Journal Activation Date Tue Nov 28 11:49:05 2006
Journal Validation Date Wed Nov 29 16:41:11 2006
Maximum Journal Size 8191 PB (9223372036854775807 Bytes)
Journal Type Change Journal
Journal State Valid
Valid for Server GSHLAGER2_SERVER1
Valid for Node GSHLAGER3
Number of DB Entries 22
Enter request on a single line, in the following format:
Req-Type [Entry-key]
Req-type might be one of the following:
Del Delete a row from the database. The fully-qualified case sensitive
file name is required.
Find Find the entry whose key is the argument.
List Print all the entries to stdout. No arguments are required.
Quit
Please enter your request: find H:\dbview.example\Dir3Depth1\F2.txt
Located Journal Database Record:
-----
Object Name : H:\dbview.example\Dir3Depth1\F2.txt
Action : Modify
Object Type : File
Inserted : Fri Dec 01 10:15:28 2006
Object Time : Fri Dec 01 14:15:28 2006
Hit Count : -2110169276
-----
Please enter your request: quit

```

Resolving problems with open file support and the logical volume snapshot agent

Several methods are available to determine the cause of problems with the open file support (OFS) and logical volume snapshot agent (LVSA).

Examining the Windows system event log

Windows

In IBM Tivoli Storage Manager, critical information is written for problem determination to the Windows system event log.

Examining the event log is the first step in isolating potential problems with the logical volume snapshot agent (LVSA) in the context of online image or open file backup.

Obtaining trace data

Windows

Many problems can be resolved through obtaining trace data from the `tsmlvsa.sys` driver to supplement what you typically get from a client service trace.

You can obtain the logical volume snapshot agent (LVSA) debug data by using the DebugView tool.

Perform the following steps to install the DebugView tool and to log data to a file:

1. Download the latest version of the DebugView tool from the Microsoft website.
2. Install the DebugView tool. The installation is a simple extraction of the files from the `DebugView.zip` file.
3. Place the following text line in the `dsm.opt` file:

```
TRACEFLAG SERVICE  
TRACEFILE trace.txt
```

Ensure that you direct the trace file to a location with a substantial amount (several GB) of free space.

4. Run the `Dbgview.exe` file before running the Tivoli Storage Manager client.
5. Configure `Dbgview.exe` to log the data to a file through the **File > Log to File** option.
6. Perform the operation that previously failed and retrieve the log file.

Configuring the system for a full memory dump

When a system bug check occurs, you must obtain a full memory dump to assist in the diagnosis of possible logical volume snapshot agent (LVSA) problems.

The failing system must be configured to take a full memory dump. The following steps show you the proper configuration:

1. Open the control panel.
2. Open the system icon.
3. Select the advanced tab.
4. Select the startup and recovery button.
5. Ensure that the following check boxes are selected in the “System Failure” section:
 - Write an event to the system log
 - Send an administrative alert
6. Ensure that the “Automatically reboot” check box is not selected.
7. In the “write debugging information” section, select “Complete Memory Dump.”

Important: Make a note of where the file will be written (`%SystemRoot%\MEMORY.DMP`). Ensure that the “Overwrite any existing file” check box is selected. Windows will ask you to restart so that the new setting can take effect.

8. Restart. If you did not select the “Overwrite any existing file” check box, after the restart you must rename the previous dump file (if any) to a new name.

When a bug check occurs, note the contents of the bug check screen. Collect the `memory.dmp` file upon reboot for examination.

Forcing a memory dump for a system stoppage when you suspect a logical volume snapshot agent problem

Windows

If you have ensured proper configuration and a memory dump is not taken, a memory dump might need to be forced when the system stops.

There are two methods you might employ:

1. If you have the opportunity to restart and recreate the stoppage, see Microsoft Knowledge Base article 244139 “Windows feature allows a Memory.dmp file to be generated with the keyboard.” This method requires a registry change and reboot to enable an on-demand memory dump when the right CTRL key is held and SCROLL LOCK key is pressed twice. This method might also be required if the BANG! tool is unable to cause a bug check and memory dump.
2. Perform the following steps if the system is stopped and you cannot afford a reboot and recreate:
 - a. Download and install BANG! from the Windows Driver Developers website. Follow instructions provided in the BANG! package/website.
 - b. Run BANG! and click the “Crash Now” button. The system typically gets a blue screen and generates a full memory dump.

Note: IBM does not support the BANG! utility. Any questions or problems regarding the BANG! utility should be reported to the Windows Driver Developers website.

Best practices for open file support

The IBM Tivoli Storage Manager client open file support (OFS) technotes outline current limitations and known problems. You can also find steps that might help you to diagnose problems in the setup and use of OFS.

Related information:

 <http://www.ibm.com/support/docview.wss?uid=swg21248971>

Using Windows Volume Shadow Copy Services

Windows

The IBM Tivoli Storage Manager Windows client uses the Volume Shadow Copy Services (VSS) of Windows 2003 and Windows Vista to perform system state and system services backup. VSS can also be used as a snapshot provider for open file support (OFS) and online image operations.

Defining VSS transient errors

Windows

The IBM Tivoli Storage Manager client considers several Volume Shadow Copy Services (VSS) errors to be transient. Transient errors are network errors or drives that are temporarily misbehaving that might require backup recovery.

When one of these errors occurs, the client will, by default, retry the VSS backup process three times at 30-second intervals. The number of retries and retry intervals can be configured using two test flags. The Tivoli Storage Manager client considers the following VSS errors to be transient:

VSS_E_MAXIMUM_NUMBER_OF_VOLUMES_REACHED
VSS_E_SNAPSHOT_SET_IN_PROGRES
VSS_E_MAXIMUM_NUMBER_OF_SNAPSHOTS_REACHED
VSS_E_PROVIDER_VETO VSS_E_UNEXPECTED
VSS_E_FLUSH_WRITES_TIMEOUT
VSS_E_HOLD_WRITES_TIMEOUT
VSS_E_WRITERERROR_TIMEOUT
VSS_E_WRITERERROR_RETRYABLE
VSS_E_WRITERERROR_OUTOFRESOURCES
VSS_E_WRITER_NOT_RESPONDING
VSS_E_VOLUME_IN_USE
VSS_E_PROVIDER_IN_USE
VSS_E_UNEXPECTED_PROVIDER_ERROR
VSS_E_UNEXPECTED_WRITER_ERROR

Defining Windows VSS test flags

Windows

The IBM Tivoli Storage Manager client uses two different test flags to configure the number of Volume Shadow Copy Services (VSS) retries and how long between retries.

The following test flags are used to set the retry and retry interval of Tivoli Storage Manager:

SETVSSMAXRETRY

Specifies the number of times the VSS backup process is retried if a transient error occurs. The default value is to retry three times.

SETVSSDELAY

Specifies the number of seconds to wait between retries of the VSS backup process, should a transient error occur. The default value is 60 seconds.

Option file example:

```
retry 10 times at 300 second intervals  
TESTFLAG SETVSSMAXRETRY:10  
TESTFLAG SETVSSDELAY:300
```

Windows 2003 VSS fixes

Windows

Several fixes for Microsoft Volume Shadow Copy Services (VSS) are available if you are experiencing difficulty with VSS.

IBM technote 1242128 lists the known fixes, but might not be current. Contact Microsoft for the most current VSS fixes.

Related information:

 [technote 1242128](#)

Volume Shadow Copy Services tuning

Windows

Several fixes for Microsoft Volume Shadow Copy Services (VSS) tuning are available if you are experiencing difficulty with VSS tuning.

Controlling the VSS diff area size

After you apply these fixes, one of the following events occurs:

- “The shadow copy of volume C: took too long to install”
- “The shadow copy of volume C: was stopped because the diff area file could not grow in time.”

Reduce the I/O load on this system to avoid these problems. If the events still occur, use the following registry key to control the size of the diff area used by VSS:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\VolSnap\  
MinDiffAreaFileSize : REG_DWORD: <size in MB> (the default size is 300, but can  
be increased to 3000).
```

Event log maximum size

Microsoft indicates that if the event logs are sufficiently large, the copy operation can take longer than the timeout for systems with high I/O load or high memory load. The log size is best when less than 64 MB.

Gathering VSS diagnostic information for Microsoft assistance

Windows

The IBM diagnostic information for Volume Shadow Copy Services (VSS) failures might not have what you need. You can find diagnostic information for VSS failures from the Microsoft support site.

If the VSS failure is outside of the scope of IBM Tivoli Storage Manager, gather the following information for Microsoft support:

- Windows application event log
- Windows system event log
- VSS trace

Examine the application and system event log files, focusing on the error events created by the VolSnap and VSS sources at the time of failure. You can extract the germane events from the log to isolate the problem to have a more productive interaction with Microsoft support.

Troubleshooting errors using a VSS trace

Windows

You can troubleshoot your Volume Shadow Copy Services (VSS) errors by conducting a VSS trace.

Perform the following steps to complete a VSS trace:

1. Create a tracing.reg file and change the TraceFile entry to point to a volume that is not going to have a shadow copy created. Use the contents at the bottom of this file to create the file. Note the double-backslash delimiter usage; you must enter "\\\" as the delimiter for each backslash in the path that you want to specify.
2. Double-click the file from within Windows Explorer to install tracing.reg.
3. Reproduce the problem.
4. Turn off tracing by deleting the "HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\VSS\Debug\Tracing" key.

The following contents are displayed in the tracefile.reg registry file:

```
Windows Registry Editor Version 5.00
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\VSS\Debug\Tracing]
"TraceFile"="c:\\trace.txt"
"TraceLevel"=dword:ffffff
"TraceEnterExit"=dword:00000001
"TraceToFile"=dword:00000001
"TraceToDebugger"=dword:00000000
"TraceFileLineInfo"=dword:00000001
"TraceForceFlush"=dword:00000000
```

Running VSS API calls with the vsreq.exe sample program

Windows

The Volume Shadow Copy Services (VSS) Software Developers Kit (SDK) contains the vsreq (VSS requester) sample program. The VSS requester program performs a sequence of VSS API calls like the calls that are performed by the IBM Tivoli Storage Manager backup-archive client.

You can compile and run vsreq.exe on the failing system to determine if vsreq and Tivoli Storage Manager encounter the same problem. If vsreq can reproduce the same problem as Tivoli Storage Manager, then the output of vsreq can be supplied to MS support to help in diagnosing the VSS problem.

In some cases, Microsoft provides an input/output (I/O) subsystem analysis tool ("yapt") to gather I/O performance data for analysis. vshadow is a tool that is also available as an alternative to vsreq.exe.

Comparing Tivoli Storage Manager and Ntbackup.exe interaction with VSS

Windows

Using the Ntbackup.exe executable file does not fully utilize Volume Shadow Copy Services (VSS) and cannot always be considered as a benchmark for IBM Tivoli Storage Manager interaction with VSS.

The known difference between Ntbackup.exe and Tivoli Storage Manager in the context of VSS is that Ntbackup.exe does not use VSS to back up the Active Directory (NTDS). Although Ntbackup.exe uses VSS to take a snapshot, it still uses the legacy NTDS backup API to read data from the disk. Tivoli Storage Manager uses the VSS interface to read NTDS data from the disk. If there is a problem with the VSS writer responsible for NTDS, it does not reveal itself with Ntbackup.exe.

Issue the **VSSADMIN LIST** command to query the VSS writer state to ensure that VSS is in a stable or ready state.

Show commands for the backup-archive client

SHOW commands are unsupported diagnostic commands that are used to display information about in-memory control structures and other runtime attributes. The **SHOW** commands are used by development and service only as diagnostic tools. Several **SHOW** commands exist for the backup-archive client.

Depending upon the information that a **SHOW** command displays, there might be instances where the information is changing or cases where it might cause the application (client, server, or storage agent) to stop running. The **SHOW** commands should only be used when development or service suggests it. The **SHOW** commands in table Table 9 are not all of the available **SHOW** commands.

Table 9. *SHOW* commands for the backup-archive client

SHOW Command	Description	Information
CACHE	Displays information about the subfile cache.	For Microsoft Windows clients that were configured to use subfile backup, this is useful for displaying information about the configured subfile cache.
CLUSTER	Displays information about the disk mappings in a Microsoft Cluster.	Useful for displaying information about the disk mapping (configuration) in a Microsoft Cluster environment.
DOMAIN	Displays information about the configured domains to use for incremental backup processing.	Useful for displaying and summarizing the DOMAIN, DOMAIN.IMAGE, and DOMAIN.NAS client options.
OPTIONS	Displays the client options.	Useful to determine the settings of client options.

Table 9. *SHOW* commands for the backup-archive client (continued)

SHOW Command	Description	Information
OPTTABLE	Displays information about options that are administered by the server versus those that are managed by the client option file.	The client might receive its option settings from either the client option file or from the server. To receive the option from the server, a client option set must be defined with the DEFINECLOPTSET command. This command helps you to determine whether the client is using an option configured from the option file or an option configured from a client option set defined on the server.
PLUGINS	Displays information about installed plug-ins for this client.	The client uses plug-ins to provide additional capabilities, such as image backup. This SHOW command displays the plug-ins that are installed for this client as well as attributes of the various plug-ins, such as their version, type, and location.
SESSION	Displays the capabilities this client is able to have for this connection to the server.	The client and server report and negotiate the capabilities that each has when a session is started between a client and a server. This SHOW command reports the capabilities available by this server and client.
SYSTEMOBJECT	For Windows 2000 and XP clients, displays the SYSTEM OBJECT data that is available on this client.	It is helpful in determining which SYSTEM OBJECT files are installed on this Windows client and those that could be backed up.
SYSTEMSERVICES	For Windows 2003 clients, displays the SYSTEM SERVICES data that is available on this client.	The SYSTEMSERVICES SHOW command is helpful in determining which SYSTEM SERVICES files are installed on this Windows client and those that could be backed up. Note: SYSTEMSERVICES is valid for Tivoli Storage Manager Version 5.4.
SYSTEMSTATE	For Windows 2003 and Windows Vista clients, displays the SYSTEM STATE data that is available on this client.	The SYSTEMSTATE SHOW command is helpful in determining which SYSTEM STATE files are installed on this Windows system and those that could be backed up.
TRACEFLAGS	Displays information about trace classes and aggregate trace classes for this client.	The TRACEFLAGS SHOW command is helpful in determining which trace classes and aggregate trace classes could be used for this client.

Table 9. *SHOW* commands for the backup-archive client (continued)

SHOW Command	Description	Information
VERSION	Displays the version and build date for this client.	The VERSION SHOW command is helpful in determining which client is running and when it was built.

Chapter 3. Resolving server problems

When working with IBM Tivoli Storage Manager, you might experience problems specific to the Tivoli Storage Manager server. The Tivoli Storage Manager server diagnostic tips that you can perform vary from simple actions such as restarting your server, to more involved procedures.

The following list contains some actions that you can perform to help diagnose server problems:

- Check the server activity log
- Recreate the problem
- Check error logs related to reading or writing to a device
- Change the server options
- Stop and start scheduling services
- Query the database or storage pool
- Trace the UNICODE trace class

Recreating the problem

Recreate the problem to isolate its cause to a specific sequence of events, if the problem can be easily or consistently recreated.

Many problems occur as a result of a combination of events. For example, expiration running along with nightly scheduled backups for 20 clients. In some cases, by changing the timing or order of implementation of events, you might prevent the problem from reoccurring. One way to change the timing is to run expiration at a time when the nightly scheduled backups for 20 clients is not running.

Deployment engine error during initialization

AIX HP-UX Linux Solaris

You might experience a deployment engine initialization error with the V6 server where no error messages are produced. If you have an invalid `/etc/cit/cit.ini` file, you are returned to the operating system prompt without advanced warning or an explanation.

An invalid `cit.ini` file is one in which the entries point to a nonexistent Common Inventory Technology file that is listed in the `/opt/tivoli/cit` directory. The following are indicators of this particular type of error that might occur during a deployment engine initialization:

- The `de_trace.log` file contains the following error:

```
ActionProcessor run() tst021dv.spr.bz IP
com.ibm.ac.tp.basehost.ActionErrorException:
action_processor_failure, action=externalCommand,
actionId=hwcollector
Caused by:
com.ibm.ac.si.ap.action.ExternalCommandActionException:
ACU0SI0050E External command action failed with return code
18.
```

```
Invocation string: /usr/ibm/common/acsi/bin/wscanhw, -o,  
/tmp/collector5198848052593475511.tmp, -c,  
/usr/ibm/common/acsi/dat/config/scanconfig.xml
```

- When you run the `/usr/ibm/common/acsi/bin/wscanvdp` executable file, it returns the following error:

```
# /usr/ibm/common/acsi/bin/wscanvdp  
CCG0051E The log manager could not read the properties file  
/CitTrace.properties.
```

- The `log.txt` file contains the following error:

```
FINER : THROW (from com.ibm.ac.coi.ext.ia.  
COIWrapperPluginImpl.processInventorySteps)  
Caused by: com.ibm.ac.coi.api.exception.COIRuntimeException:  
Exception received trying to execute the target install  
in the ANT Buildfile for the step: TSM_Inventory  
at: com.ibm.ac.coi.impl.utils.COIUtilsImpl.runPackageStep  
(COIUtilsImpl.java:126)  
at: com.ibm.ac.coi.ext.ia.COIWrapperPluginImpl.  
processInventorySteps(COIWrapperPluginImpl.java:1960)  
at: com.ibm.ac.coi.ext.ia.plugin.COIProcessInventorySteps.  
install(COIProcessInventorySteps.java:120)  
at: com.zerog.ia.installer.actions.CustomAction.installSelf  
(DashoA10*...:-1)  
at: com.zerog.ia.installer.AAMgrBase.a(DashoA10*...:-1)  
at: com.zerog.ia.installer.ConsoleBasedAAMgr.a  
(DashoA10*...:-1)
```

You can resolve this problem by completing the following steps:

1. Start the installer.
2. Rename the `/etc/cit/cit.ini` file. For example, `/etc/cit/cit.orig`.
3. When the installation is complete, return the file to its original name.

Checking the server activity log file

Check the server activity log file and look at the reports occurring 30 minutes before and 30 minutes after the time of the error.

To review the messages in the server activity log, issue the **QUERY ACTLOG** command. Often, other messages can offer additional information about the cause of the problem and how to resolve it.

Checking system error log files for device errors

If the problem is an error created by reading or writing data from a device, many systems and devices record information in a system error log.

If a device or volume that is being used by IBM Tivoli Storage Manager is reporting some sort of error to the system error log, it is likely a device issue. The error messages recorded in the system error log might provide enough information to resolve the problem.

The following are some examples of system error logs:

- `errpt` for AIX
- Event Log for Windows

Reverting server options or settings

If there were configuration changes to the server, try reverting the settings back to their original values and retry the failing operation.

If the operation succeeds, try to make one change at a time and retry the operation until the attribute change that caused the failure is identified.

Changes to options in the server options file, or configuration changes to the server using **SET** or **UPDATE** commands might cause failures for operations that had previously succeeded. Changes on the server to device classes, storage pools, and policies might also cause failures to operations that had previously succeeded.

Restarting the scheduling service

Scheduled client operations are influenced by the schedule definitions on the server as well as the scheduling service (dsmsched) that runs on the client computer itself.

Restart the scheduling service on the client if a schedule changes on the server.

Important: If the scheduling service is managed by the client acceptor, stop and restart only the client acceptor.

Resolving server space issues

The IBM Tivoli Storage Manager server's primary function is to store data. If it runs out of space in the database or storage pools, operations might fail.

To determine if the database is out of space, issue the **QUERY DB** command. If the percent utilized (used space) is at or near 100%, define more space. Typically, if the database is running out of space, this situation is indicated by other issued server messages.

To determine if a storage pool is out of space, issue the **QUERY STGPOOL** command. If the percent utilized is at or near 100%, make more storage space available. To add more space to a DISK storage pool, allocate one or more new storage pools and define them to the server using the **DEFINE VOLUME** command. You can configure Tivoli Storage Manager to automatically allocate storage pool DISK and FILE space by using the **DEFINE SPACETRIGGER** command.

To add more space to a sequential media storage pool, evaluate the tape library and determine if more scratch tapes can be added. If so, add the additional scratch volumes to the library and update the **MAXSCR** parameter for the storage pool by issuing the **UPDATE STGPOOL** command.

Allocating additional server memory

Allocate more memory on the server if there are indications that your server is low on memory resources. Refer to your operating system's documentation for information about adding memory.

Tip: The amount of memory that DB2 uses might contribute to reports that show that the operating system is out of memory. You can limit the amount of memory that DB2 uses by including the `DBMEMPERCENT` option. The `DBMEMPERCENT` option specifies the percentage of virtual address space that is dedicated to the database manager processes.

Complete the following actions to allocate additional storage resources for the server:

- **AIX** Ensure that there is sufficient paging space. You can also use SMIT (System Management Interface Tool) to determine if the number of applications is causing a memory shortage.
- **HP-UX** Ensure that there is sufficient paging space and that a sufficient amount of shared memory is available. For information about shared memory, refer to the *IBM Tivoli Storage Manager Installation Guide*.
- **Solaris** Ensure that there is sufficient paging space. Consult your Solaris system documentation for details.
- **Windows** The preferred method of solving a low memory condition is to add physical memory to the system. Otherwise, from the control panel, increase the amount of the virtual storage by running the system applet and increasing the total paging file size.

Changing the copy frequency

IBM Tivoli Storage Manager server policy demands that an incremental copy frequency be a non-zero value.

The copy frequency attribute of the current *copygroup* management class for the file that is specified dictates the minimum number of days that must elapse between successive incremental backups. If you are trying to perform an incremental backup on a file and this number is set higher than 0 days, then the file will not be sent to the Tivoli Storage Manager server, even if it has changed.

A number of steps can be taken to correct this problem:

- Contact the Tivoli Storage Manager server administrator to change the copy frequency attribute.
- Issue a selective backup of the file. For example, `DSMC SELECTIVE C:\FILE.TXT`

You can issue the **QUERY COPYGROUP** command to determine the setting of the copy frequency parameter:

```
tsm: WINBETA>q copygroup standard active f=d
Policy Domain Name: STANDARD
...
Copy Frequency: 1
...
```

Resolving RELABEL operation errors

If you run a RELABEL operation when all drives are busy, the target volume cannot be relabeled because it cannot obtain a drive. Busy drives are drives that are in use for regular operations such as backup, restore, migration, and reclamation.

When a RELABEL error occurs, the following example information is produced:

```
ANR0984I Process 25 for RELABEL started in the BACKGROUND at 22:10:36.
ANR8799I RELABEL: Operation for library IBMVTL started as process 25.
ANR1341I Scratch volume 007403 has been deleted from storage pool VTLPOOL.
ANR8847E No LTO-type drives are currently available in library IBMVTL.
ANR8801I LABEL LIBVOLUME process 25 for library IBMVTL completed; 0 volume(s)
labeled, 0 volume(s) checked-in.
ANR0985I Process 25 for RELABEL running in the BACKGROUND completed with
completion state SUCCESS at 22:10:36.
```

To resolve a RELABEL error, complete the following steps:

1. Ensure that one drive is kept available for the RELABEL operation and relabel a target volume.
2. Update the device classes that point to the library. Update the device classes with a **MOUNTLIMIT** parameter value that is less than the total number of drives available.

If a RELABEL operation cannot obtain a drive or fails to relabel a volume, IBM Tivoli Storage Manager tries to relabel the volume during each future RELABEL operation.

If the RELABEL operation fails, issue the **LABEL LIBVOLUME** command for all volumes that are checked out of Tivoli Storage Manager but that are not relabeled. Include the following parameters with the **LABEL LIBVOLUME** command:

```
SEARCH=YES LABELSOURCE=BARCODE OVERWRITE=YES CHECKIN=SCRATCH
```

Resolving installation and upgrade problems

Resolving installation problems with IBM Tivoli Storage Manager server can involve reviewing log files, reinstalling the server, or several other possible options.

Installation log files

If you experience errors during installation, these errors are recorded in several log files that are distributed in various locations.

Depending on the operating system, files stored in the OS temporary directory might not be saved after the system is shut down.

Linux Define set SELINUX=disable or set SELINUX=permissive in the `/etc/sysconfig/selinux` file if you receive the following message:

The installer cannot run on your configuration. It will now stop.

Table 10 on page 54 describes the files that are created when you install or uninstall IBM Tivoli Storage Manager. Recommendations on the files to check when you are looking for information that might help troubleshoot problems are available.

Table 10. Contents of the log.zip file

File name	Description	Location
<p>The log.txt file contains information about the following environment checks:</p> <ul style="list-style-type: none"> • Platform • Version • Architecture • Prerequisites 	<p>Contains installation log files.</p> <p>Review this log file when any installation failures occur.</p>	<p>For Windows, this file is located in the InstallAnywhere location that you specified. For AIX, HP-UX, Linux, and Solaris, the file is located in /var/tivoli/tsm.</p> <p>The InstallAnywhere exit codes show you the state of the installation. The codes are in the log.txt file. You can retrieve the exit codes after the installation is completed. The exit codes are for both the Tivoli Storage Manager installer and the Tivoli Monitoring for Tivoli Storage Manager installer.</p> <p>AIX HP-UX Linux Solaris To use the command line, issue the following command: echo \$?</p> <p>Windows Issue the following command: echo %ERRORLEVEL%</p> <p>See Table 11 on page 56 for all of the InstallAnywhere exit codes.</p>

Table 10. Contents of the log.zip file (continued)

File name	Description	Location
logs.zip	<p>The logs.zip file might contain the following directories:</p> <p>ac Contains setup log files</p> <p>coi Contains installation log files</p> <p>de Contains setup log files for device</p> <p>ia Contains the log.txt file</p> <p>tcr Contains the Tivoli Common Reporting setup log files</p> <p>tip Contains Tivoli Integrated Portal setup log files</p> <p>profiles Contains TIPProfile setup log files</p>	<p>Windows This file is located in the InstallAnywhere location that you specified.</p> <p>AIX HP-UX Linux</p> <p>Solaris The file is located in /var/tivoli/tsm.</p>
DE_Install.log	<p>Contains information about the Deployment Engine (DE) installation.</p> <p>Review this log file if the DE installation fails.</p>	de\root
<p>AIX HP-UX</p> <p>Linux Solaris</p> db2setup.log	<p>Contains information about the DB2 installation.</p>	<p>coi\plan\tmp if the DB2 installation fails</p> <p>coi\plan\logs if the DB2 installation succeeds</p>
<p>Windows</p> db2setup.log	<p>Contains information about the DB2 installation</p>	<p>coi\plan\tmp if the DB2 installation fails</p> <p>coi\plan\logs if the DB2 installation succeeds</p>
db2_uninst.log	<p>Contains information about the DB2 uninstallation</p>	coi\plan\logs

Table 10. Contents of the log.zip file (continued)

File name	Description	Location
DB2.log	Contains information about the installation and uninstallation commands. Return codes can be retrieved from this log file, but not for DB2. If installation or uninstallation completed, the executePackage or remove-package scripts for a component are available.	coi\plan\install or coi\plan\uninstall
<div style="display: flex; justify-content: space-between; margin-bottom: 5px;"> AIX Linux </div> <div style="display: flex; justify-content: space-between; margin-bottom: 5px;"> Solaris Windows </div> Administration Center installation log files	Installation log files. Review these log files if the Administration Center installation or uninstallation fails.	<ul style="list-style-type: none"> • coi\plan\install\logs • coi\plan\install\MachinePlan_host_name\00001_eWAS • coi\plan\install\MachinePlan_host_name\00002_TIP • coi\plan\install\MachinePlan_host_name\00003_TSM_AdminCenter or <ul style="list-style-type: none"> • install_root_uninst\plan\install\MachinePlan_host_name\00001_eWAS • install_root_uninst\plan\install\MachinePlan_host_name\00002_TIP • install_root_uninst\plan\install\MachinePlan_host_name\00003_TSM_AdminCenter
<div style="display: flex; justify-content: space-between; margin-bottom: 5px;"> AIX Linux </div> <div style="background-color: #800000; color: white; padding: 2px 5px; width: fit-content;">Solaris</div> Files with the following file extensions are available for troubleshooting: .log .out .err	Contains installation log files. Review this log file when any installation failures occur.	Files are in the <i>Adm_Cntr_Root</i> /logs directory <i>Adm_Cntr_Root</i> is the directory where the Administration Center is installed.

Table 11. InstallAnywhere exit codes

Code	Description
0	Success: The installation completed successfully without any warnings or errors.
1	The installation completed successfully, but one or more of the actions from the installation sequence caused a warning or a disrupting error.

Table 11. InstallAnywhere exit codes (continued)

Code	Description
-1	One or more of the actions from the installation sequence caused an unrecoverable error.
1000	The installation was canceled by the user.
1001	The installation includes an invalid command-line option.
2000	Unhandled error.
2001	The installation failed the authorization check, might indicate an expired version.
2002	The installation failed a rules check. A rule placed on the installer itself failed.
2003	An unresolved dependency in silent mode caused the installer to exit.
2004	The installation failed because not enough disk space was detected during the execution of the installation action.
2005	The installation failed while trying to install on a Windows 64 - bit system, but installation did not include support for Windows 64 - bit systems.
2006	The installation failed because it was started in a UI mode that is not supported by this installer.
3000	Unhandled error specific to a launcher.
3001	The installation failed due to an error specific to the lax.main.class property.
3002	The installation failed due to an error specific to the lax.main.method property.
3003	The installation was unable to access the method specified in the lax.main.method property.
3004	The installation failed due to an exception error caused by the lax.main.method property.
3005	The installation failed because no value was assigned to the lax.application.name property.
3006	The installation was unable to access the value assigned to the lax.nl.java.launcher.main.class property.
3007	The installation failed due to an error specific to the lax.nl.java.launcher.main.class property.
3008	The installation failed due to an error specific to the lax.nl.java.launcher.main.method property.
3009	The installation was unable to access the method specified in the lax.nl.launcher.java.main.method property.
4000	A Java executable file was not found at the directory specified by the java.home system property.
4001	An incorrect path to the installer compressed (.jar) file caused the relauncher to start incorrectly.

Resolving a stopped uninstallation process

A stoppage during an IBM Tivoli Storage Manager uninstallation process might be caused by an expired DB2 password.

If the DB2 administrator's userid password is expired, the uninstallation process cannot complete. You must log in using the DB2 administrator's ID and reset the password, then uninstall Tivoli Storage Manager.

Server installation failure: deployment engine initialization

The IBM Tivoli Storage Manager installation might be stopped due to a deployment engine initialization failure due to .lock files.

If the deployment engine interferes with the Tivoli Storage Manager installation, the following error message is produced:

```
Tivoli Storage Manager
Deployment Engine failed to initialize.
The installer will now shutdown. Please check with the log files for a more
complete description of the failure.
PRESS ENTER TO CONTINUE:
```

The cause might be deployment engine .lock files that come from a concurrent installation that is running or from an installation that stopped before completing. If another installation is running, wait until that installation finishes before installing the Tivoli Storage Manager. If there are no other installations running and you encounter this problem, delete any .lock files.

Important: Do not delete any .lock files if there are other Tivoli Storage Manager installations running.

Windows To delete .lock files on 32-bit platforms, issue the following command:

```
cd C:\Program Files\IBM\Common\acsi\logs
del .lock*
```

AIX **HP-UX** **Linux** **Solaris** To delete .lock files, issue the following command:

```
cd /usr/ibm/common/acsi/logs
rm .lock*
```

Manually recreating server instances due to connectivity problems

When a connection cannot be established, the installer cannot recreate your IBM Tivoli Storage Manager V6.3 server instances. You must manually recreate your server instances.

The installation wizard uses the following methods to establish a connection to the system to recreate the server instances:

- Secure shell (SSH)
- Remote shell (RSH)
- Remote Execution Protocol (REXEC)
- **Windows** Windows server message block (SMB)

When you use one of these methods on the default port, the port cannot be blocked by a firewall. If it is blocked, complete the following steps to manually upgrade the server instance:

AIX

HP-UX

Linux

Solaris

1. Close the installation wizard.
2. You can use the contents of the `instance.info` file in the commands that you issue to recreate the server instance.
3. Complete the following steps for each Tivoli Storage Manager server instance:

- a. After the upgrade completes, issue the following command to recreate the instance:

```
/opt/tivoli/tsm/db2/instance/db2icrt -u instance_user instance_name
```

- b. Recreate the variables in the instance file. Issue the **db2set -i** command for each variable in the instance file. For example:

```
/opt/tivoli/tsm/db2/instance/db2set -i instance_name "variable"
```

Be sure to include the quotation marks in your command.

- c. Issue the **db2stop** command to stop the database instance.
- d. Use the user ID that owns the server instance to issue the **db2start** command to start the database instance.
- e. Catalog and upgrade each database by issuing the following commands:

```
db2 catalog db TSMDB1 on "database_path"  
db2 upgrade db TSMDB1
```

- f. Issue the **db2stop** command.
- g. Start the Tivoli Storage Manager server.

Resolving a server upgrade failure

AIX

HP-UX

Linux

Solaris

If you are using the **DB2CKUPGRADE** command to upgrade the IBM Tivoli Storage Manager server, the command might fail from an incorrect DB2 authentication. The `log.txt` file records the errors after a failed upgrade.

The following is a sample `log.txt` file:

```
Wed Apr 21 15:52:53.956  
CDT 2010 :  
INFO :  
Issuing cmd: . ~/sqlllib/db2profile;  
/code_path/COI/PackageSteps/DB2/FILES  
/ese/db2/aix/utilities/db2ckupgrade/bin/db2ckupgrade  
TSMDB1 -u tsmint1 -p xxxxxxx -l /tmp/db2ckupgrade_tsmint1_TSMDB1.log  
(from com.tivoli.dsm.install.RedefineInstances.checkDatabase)  
Wed Apr 21 15:52:53.964  
CDT 2010 :  
INFO : signalEvent (from com.tivoli.dsm.install.Db2ckupgradeConsole.signalEvent)  
Wed Apr 21 15:52:55.444  
CDT 2010 :  
SEVERE : Cmd failed (from com.tivoli.dsm.install.RedefineInstances.checkDatabase)
```

This example indicates that the failing command is the following entry:

```
/code_path/COI/PackageSteps/DB2/FILES  
/ese/db2/aix/utilities/db2ckupgrade/bin/db2ckupgrade  
TSMDB1 -u tsmint1 -p xxxxxxx -l /tmp/db2ckupgrade_tsmint1_TSMDB1.log
```

Verify the cause of the upgrade failure by looking in the `/tmp/db2ckupgrade_tsmint1_TSMDB1.log` file. For example:

```

cat /tmp/db2ckupgrade tsminst1_TSMDB1.log
Version of DB2CKUPGRADE being run: VERSION 9.7.
Database: 'TSMDB1'
** ERROR **
Database cannot be connected.

```

The previous example shows that the following command failed:

```

$ db2 connect to TSMDB1 user tsminst1 using password
SQL30082N Security processing failed with reason "42"
("ROOT CAPABILITY REQUIRED"). SQLSTATE=08001

```

Complete the following steps to resolve the upgrade problem:

1. Confirm that you have the correct password. Log in to the computer system with the server instance system user ID and password.
2. Issue the **db2stop force** command to stop the server instance.
3. As root user, issue the **db2iupdt** command. For example: `/opt/tivoli/tsm/db2/instance/db2iupdt tsminst1`
4. Run the failed **db2ckupgrade** command. For example:

```

/code_path/COI/PackageSteps/DB2/FILES/ese
/db2/aix/utilities/db2ckupgrade/bin/db2ckupgrade
TSMDB1 -u tsminst1 -p xxxxxxx

```
5. Run the `install.bin` program to upgrade Tivoli Storage Manager.

Resolving GSKit installation problems

When you use the Tivoli Storage Manager installation software, the correct Global Security Kit (GSKit) version is installed automatically.

If the Tivoli Storage Manager server instance environment is not configured properly, the Tivoli Storage Manager server might not load the appropriate GSKit libraries. The Tivoli Storage Manager server instance configuration wizard helps you to avoid many issues that might be prevalent when you manually configure the Tivoli Storage Manager instance.

Windows Issue the following command:

```
set PATH=X:\Program Files\IBM\gsk8\bin;X:\Program Files\IBM\gsk8\lib64;%PATH%
```

where *X* is the system drive. The PATH environmental variable is modified to point to the correct directory.

HP-UX **Linux** **Solaris** Update the LD_LIBRARY_PATH or shell by issuing the following command:

```
export LD_LIBRARY_PATH=platform-specific-gskit-library-directory:$LD_LIBRARY_PATH
```

where *platform-specific-gskit-library-directory* is one of these directories, according to your platform:

- **HP-UX** **Solaris** /opt/ibm/gsk8_64/lib64
- **Linux** /usr/local/ibm/gsk8_64/lib64

AIX For AIX, issue the following command:

```
export LIBPATH=/usr/opt/ibm/gsk8_64/lib64:$LIBPATH
```

AIX **HP-UX** **Linux** **Solaris** You must update the following files to set the library path when DB2 or the Tivoli Storage Manager server are started:

- `instance_directory/sqlllib/usercshrc`
- `instance_directory/sqlllib/userprofile`

For the `instance_directory/sqlllib/usercshrc` file, add the following lines:

- **AIX**
`setenv LIBPATH /usr/opt/ibm/gsk8_64/lib64:$LIBPATH`
- **HP-UX** **Solaris**
`setenv LD_LIBRARY_PATH /opt/ibm/gsk8_64/lib64:$LD_LIBRARY_PATH`
- **Linux**
`setenv LD_LIBRARY_PATH /usr/local/ibm/gsk8_64/lib64:$LD_LIBRARY_PATH`

For the `instance_directory/sqlllib/userprofile` file, add the following lines:

- **AIX**
`LIBPATH=/usr/opt/ibm/gsk8_64/lib64:$LIBPATH`
`export LIBPATH`
- **HP-UX** **Solaris**
`LD_LIBRARY_PATH=/opt/ibm/gsk8_64/lib64:$LD_LIBRARY_PATH`
`export LD_LIBRARY_PATH`
- **Linux**
`LD_LIBRARY_PATH=/usr/local/ibm/gsk8_64/lib64:$LD_LIBRARY_PATH`
`export LD_LIBRARY_PATH`

Verify the library path settings and the GSKit version by issuing the following commands:

- **AIX**
`echo $LIBPATH`
`gsk8capicmd_64 -version`
- **HP-UX** **Linux** **Solaris**
`echo $LD_LIBRARY_PATH`
`gsk8ver_64`

If your GSKit version is not 8.0.14.14 or later, you must reinstall the Tivoli Storage Manager server. The reinstallation ensures that the correct GSKit version is available.

Resolving server stoppages

Server stoppages can occur from processing errors, the system trap handler, or other errors. When you determine the source of your server stoppage, the reason might resolve other known issues.

The server might stop for one of the following reasons:

- A processing error causes memory to be overwritten or some other event triggers the system trap handler to terminate the server process.
- The server processing has validation algorithms throughout the application that check various conditions prior to continuing running. As part of this validation checking, there are cases where if the validation check fails, the server will actually terminate itself instead of allowing processing to continue. These catastrophic validations are referred to as an assert. If the server terminates due to an assert, the following message is issued:

```
ANR7837S Internal error XXXNNN detected.
```

where *XXXNNN* is an identifier assigned to the assertion failure.

Other server messages that are indicative of a stoppage are ANR7836S and ANR7838S.

Whether the server stopped as a result of an assert or the system trap handler, the *tsmdiag* utility can collect the following information and package it for submission for IBM service so the situation can be diagnosed:

- Server error file (*dsmserv.err*)
- System image (core file)
- Libraries and other files
- System logs
- Activity log

Package all the data (files) collected and contact IBM service to report this problem.

Related reference:

Appendix B, "Installing and running the *tsmdiag* utility," on page 255

Resolving a stoppage or loop

A stoppage is a situation where the server does not start or complete a function and is not using any microprocessor power.

A stoppage might be just one session or process that is not processing, or it could be the entire IBM Tivoli Storage Manager server not responding. A loop is a situation where no progress is being made, but the server is using a high amount of microprocessor power. A loop can affect just one session or process, or it could affect the entire Tivoli Storage Manager server.

You might collect documentation to resolve this type of problem, depending on whether the server is able to respond to commands. A Perl script is available for you to collect Tivoli Storage Manager server data. It is helpful to schedule the **SHOW** command list to run intermittently so that you can then see the behavior that precedes the stoppage situation.

- For a stoppage or a loop where the server can respond to commands, issue the following commands to help determine the cause of the stoppage:
 - **QUERY SESSION f=d**
 - **QUERY PROCESS**
 - **SHOW RESQ**
 - **SHOW THREADS**
 - **SHOW DEADLOCK**
 - **SHOW TXNT**
 - **SHOW DBTXNT**
 - **SHOW LOCKS**
 - **SHOW LIBR**
 - **SHOW MP**
 - **SHOW SESS**
 - **SHOW ASQ**
 - **SHOW ASVOL**
 - **SHOW DBV**
 - **SHOW SSS**

- **SHOW CSV** (do this only if the problem appears to be related to scheduling)
- When a server hangs or loops, issue the following commands to provide a detailed diagnostic snapshot of the Tivoli Storage Manager environment:


```
db2fodc -hang -alldbs
db2support . -d database -s
```

You can use the `db2support.zip` file that is generated for troubleshooting.

- In addition to the output from the listed commands, or in the case of a server that cannot respond to commands, collect a dump. The way that you collect a dump depends on the operating system.
 - **AIX** **HP-UX** **Linux** **Solaris** Issue the **KILL -11** command on the `dsmserv` process to create a core file. You can obtain the process ID to perform the “kill” by issuing the **PS** command.
 - **Windows** Refer to the Microsoft Knowledge Base item 241245 at microsoft.com for instructions on installing and using the `userdump.exe` program to obtain a dump.

Resolving wait state problems with external user repository servers

AIX **Linux** **Solaris**

If the IBM Tivoli Storage Manager server seems unresponsive, it might be related to the operating system and the operating system’s use of an external user repository.

Slow Tivoli Storage Manager server performance can be attributed to an operating system that uses an external user repository that has too many user groups defined. NIS (Network Information Service) and LDAP (Lightweight Directory Access Protocol) servers are two types of external user repository servers.

An example of the unresponsive behavior is when Tivoli Storage Manager takes a long time to connect to the IBM DB2 server. Another example is when the Tivoli Storage Manager server seems unresponsive to administrative requests.

Complete the following steps to resolve a wait state problem that occurs with AIX, Linux, or Solaris servers when using an LDAP server:

1. Stop the Tivoli Storage Manager server.
2. Issue the following commands for the AIX servers:
 - a. `db2set DB2_ALTERNATE_GROUP_LOOKUP=GETGRSET`
 - b. `db2stop force`
 - c. `db2start`

Issue the following commands for the Solaris servers:

- a. `db2set FASTGROUPS=TRUE`
- b. `db2stop force`
- c. `db2start`

Issue the following commands for the Linux servers:

- a. `db2set DB2_ALTERNATE_GROUP_LOOKUP=GETGROUPLIST`
- b. `db2stop force`
- c. `db2start`

3. Restart the Tivoli Storage Manager server.

Finding the server error file (dsmserv.err)

When the server stops, it appends information to the `dsmserv.err` file which is located in the same directory as the server.

AIX **Linux** The trap handler is disabled to prevent the function traceback from printing on the console and in the `dsmserv.err` file. This change is required in order to ensure that we will get a more complete core file. As part of disabling the trap handler, a new script, `getcoreinfo`, is in the Linux packages. The `getcoreinfo` script gets the function traceback for the failing thread and registers values and function traceback for all other threads. The amount of information available in the core for other threads is still incomplete on some Linux platforms/distributions. See the `getcoreinfo` script (in the server bin directory) for more details.

Windows If the server is running as a service, the file is named `dsmsvc.err`.

HP-UX It is possible to use the following script to get basic thread information from the core on the customer system without sending it to IBM.

Perform the following steps to capture the server error file:

1. Make sure that the GNU debugger (gdb) is installed on the customer system.
2. Copy the `gt` shell script to the server bin directory (where the server executable `[.exe]` file and core file are located).
3. Make sure the script is an executable file (`chmod a+x gt`).
4. Invoke the script with the paths/names of the executable file (default is `./dsmserv`) and the core (default is `./dsmcore`). The output is in the `dsm_gdb.info` file (which should be sent to IBM).

Finding the system image (core file)

Typically, a core file or other system image of the memory is in use by IBM Tivoli Storage Manager at the time of the failure.

In each case, rename the core file to prevent it from being overwritten by a later stoppage. For example a file should be renamed to “`core.Aug29`” instead of just “`core`.” The type and name of the core file varies depending upon the platform:

- **AIX** **HP-UX** **Linux** **Solaris** A file is typically created called `core`. Be sure that there is enough space in the server directory to accommodate a dump operation. It is common to have a dump file as large as 2 GB for the 32-bit Tivoli Storage Manager server. Additionally, make sure the `ulimit` for core files is set to unlimited to prevent the dump file from being truncated.
- **Windows** The contents of the system are dumped automatically through a system application programming interface (API) call. If the server is running as a service, the dump file is called `dsmsvc.dmp`. Otherwise, the dump file is called `dsmserv.dmp`

If the system was not configured to capture a core file or the system did not have sufficient space to create a complete core file, it might be of limited use in determining the cause of the problem.

Retrieving library files for core analysis

AIX

HP-UX

Linux

Solaris

Core files are specific to the application, libraries, and other system resources in use by the application on the system where it was running.

To accurately read the core file on your computer system, you need all of the following files which are located in the directory where the server is installed:

- dmserv
- dsmlicense
- ndmpspi
- dsmcored
- dsmaio
- centera

The library files that are needed vary between all the platforms:

- **AIX** Collect the following files:
 - /usr/ccs/lib/libpthreads.a
 - /usr/ccs/lib/libc.a
 - Collect any other loaded libraries such as message exits. To see what libraries are loaded, invoke dbx by issuing the **dbx dmserv core_file** command. Then, from the dbx prompt, issue the **map** command to show all of the libraries that are loaded and needed for core analysis.
- **HP-UX** Issue the **chatr dmserv** command and send in all the dynamic shared libraries. For example:
 - /usr/lib/libpthread.1
 - /usr/lib/libm.2
 - /usr/lib/libstd.2
 - /usr/lib/libstream.2
 - /usr/lib/libCsup.2
 - /usr/lib/libcl.2
 - /usr/lib/libc.2
 - /usr/lib/libdl.2
- **Linux** Issue the **ldd dmserv** command and send in all the dynamic shared libraries. For example:
 - libm.so.6 =>/lib64/libm.so.6
 - libnsl.so.1 =>/lib64/libnsl.so.1
 - libpthread.so.0 =>/lib64/libpthread.so.0
 - libdl.so.2 =>/lib64/libdl.so.2
 - libc.so.6 =>/lib64/libc.so.6
 - /lib64/ld64.so.1 =>/lib64/ld64.so.1
- **Solaris** Issue the following commands to collect the needed libraries:
 - **sh**
 - **cd /usr**
 - **find . -name "ld.so" -print ; **
 - **find . -name "ld.so.?" -print ; **

```

- find . -name "libm.so.?" -print ; \
- find . -name "libsocket.so.?" -print ; \
- find . -name "libnsl.so.?" -print ; \
- find . -name "libthread.so.?" -print ; \
- find . -name "libthread_db.so.?" -print ; \
- find . -name "libdl.so.?" -print ; \
- find . -name "libw.so.?" -print ; \
- find . -name "libgen.so.?" -print ; \
- find . -name "libCrun.so.?" -print ; \
- find . -name "libc.so.?" -print ; \
- find . -name "libmp.so.?" -print ; \
- find . -name "libc_psr.so.?" -print ; \
- find . -name "librtld_db.so.?" -print) > runliblist
- tar cfh runliblist.tar -I runliblist

```

Retrieving system log files

You can retrieve system log files to help resolve the causes of server stoppages.

Retrieve the following log files to give to IBM service:

- For the AIX platform redirect the output from the command **errpt -a** into a file:
errpt -a >errpt.txt.
- For the HP-UX platform, copy the /var/adm/syslog/syslog.log file.
- For the Linux platform, copy the /var/log/messages file.
- For the Solaris platform, copy the /var/adm/messages file.
- For the Windows platform, save a copy of the Event Logs, as seen from the Event Viewer.

Retrieving the activity log

Activity log files can be retrieved to help resolve problems with server stoppages.

View and collect the activity log entries that start at least two hours prior to the stoppage and 30 minutes after the stoppage by issuing the **QUERY ACTLOG** command. See the *Tivoli Storage Manager Administrator's Reference* for more details.

Detecting errors after a server service starts and stops

Windows

If a server service unexpectedly starts and stops, you can determine the cause of the error by requesting an error log file.

A service can be started from either the Tivoli Storage Manager Management Console or from the Windows Services applet. After you start the service, the service might indicate that it is started, but after refreshing, the service indicates that it is stopped. In the following steps, "Server1," is used as the name of the server that started and stopped. To determine the cause of the error for Server1, finish the following steps:

1. Issue the following command to start the Management Console:

```
Tivoli_Storage_Manager_installation_path\console\tsmnet64.msc
```

where *Tivoli_Storage_Manager_installation_path* is the location where the Tivoli Storage Manager server is installed.

2. Expand **Tivoli Storage Manager > [Hostname] (Windows - Local) > Server1 > Reports > Service Information** to show the server service.
3. In the right pane, right-click **Server1** service, and select **Properties**.
4. Select the option **Log output to file** and click **OK**.
5. Start the Server1 service.
6. If the service stops again, open a text editor to read the contents of the following file:
C:\Program Files\Tivoli\TSM\Server1\console.log
7. Determine the cause of the error by reviewing the error messages that are generated.

sqllib/db2dump directory causes shutdown

Tivoli Storage Manager V6 servers might shut down unexpectedly if the `sqllib/db2dump` directory overfills. The most common time for a shutdown is when the DB2 first occurrence data capture (FODC) files are written to the directory.

The `sqllib/db2dump` directory is a diagnostic data directory path that DB2 uses to write diagnostic information for FODC. Over time, DB2 can write many FODC files to the directory related to the health of the database. When files are not removed or deleted, the file system can become full. The location of the DB2 first occurrence data capture (FODC) files depends on your DB2 configuration settings or the DB2 environmental variable settings.

Locate the diagnostic data directory by checking DB2 configuration settings or DB2 environment variable settings. If the files in the diagnostic directory path cause the file system to become full, take one of the following actions:

- Add space to the file system.
- Move the files to another file system. See Table 12.
- Use the Tivoli Storage Manager server to archive the files, and then delete them by using the following steps:
 1. Run the `db2support` utility to collect the DB2 system diagnostic information.
 2. Archive the `db2support.zip` file and diagnostic files that are listed in Table 12 to Tivoli Storage Manager server with the Tivoli Storage Manager client.
 3. Delete the files that are listed in Table 12.

Table 12. Files that can be deleted after archiving

File name	Description
instance_name.nfy instance_name.n.nfy (where <i>n</i> is a number)	Administration notification logs
db2dasdiag.log	DB2 administration server (DAS) diagnostic log
db2eventlog.xxx (where <i>xxx</i> is the database partition number)	DB2 event log
nnnnnn.nnnn.nnn.dump.bin (where <i>n</i> is a number)	Binary dump files of key in-memory structures
nnnnnn.n.nnn.trap.txt (where <i>n</i> is a number)	Trap files

Table 12. Files that can be deleted after archiving (continued)

File name	Description
<i>nnnnnnn.nnnnnn.nnn.apm.bin</i> (where <i>n</i> is a number)	Access the plan-manager binary dump files
<i>nnnnnnn.nnnnnn.nnn.stack.txt</i> (where <i>n</i> is a number)	Stack traces
F0DC_XXX/core<pid>	Core files These F0DC_XXX directories contain the timestamp in the directory name. Keep the most recent directories and its files. The history can be useful for diagnosing possible future problems related to the database. A guideline is to keep at least 1 weeks worth.
events/db2optstats. <i>n</i> .log (where <i>n</i> is a number)	Statistics log file

Tip: Do not delete the db2diag.log file and files within the stmmlog directory. The history contained within them can be useful for diagnosing server problems that are related to the database.

Related reference:

“Locating DB2 diagnostic log files” on page 71

Resolving database errors

Server errors might be caused by database irregularities. Some more common issues are running out of space and errors caused by insert, update, or delete operations.

Tivoli Storage Manager V6.3 is installed with the IBM DB2 database application. Users who are experienced DB2 administrators can choose to perform advanced SQL queries and use DB2 tools to monitor the database. However, do not use DB2 tools to change DB2 configuration settings from those settings that are preset by Tivoli Storage Manager. Do not alter the DB2 environment for Tivoli Storage Manager in other ways, such as with other products. The Tivoli Storage Manager Version 6.3 server must be used with the data definition language (DDL) and database configuration that Tivoli Storage Manager deploys.

You can find information about DB2 in the DB2 information center (<http://pic.dhe.ibm.com/infocenter/db2luw/v9r7>).

HP-UX and Solaris users must run the **db2osconf** utility after installing and configuring Tivoli Storage Manager to ensure that the system kernel parameters are set properly.

See the *Installation Guide* for information on the **db2osconf** utility.

Related information:

 <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp>

Resolving database manager starting problems

Tivoli Storage Manager server might not start if the DB2 database manager is configured to use the dsmdb2pw plug-in. When the server cannot load the plug-in, the database manager does not start and, in turn, the server fails to start.

Due to the plug-in problem, the server issues an error message similar to this example:

```
db2start
SQL1365N db2start or db2stop failed in processing the plugin "dsmdb2pw".
Reason code = "10".
04/26/2011 16:04:11 0 0 SQL1365N
db2start or db2stop failed in processing the plugin "". Reason code = "".
```

You might also receive this error:

```
SQL1032N No start database manager command was issued
```

Review the db2diag.log file for diagnostic information regarding this type of error.

An example from the db2diag.log:

```
2011-04-26-16.04.11.820963-420 I2345542E1168 LEVEL: Error
PID : 25178 TID : 47207843621184PROC : db2sysc 0
INSTANCE: hannigan NODE : 000
EDUID : 1 EDUNAME: db2sysc 0
FUNCTION: DB2 Common, OSSe, OSSHLlibrary::load, probe:80
MESSAGE : ECF=0x90000076=-1879048074=ECF_LIB_CANNOT_LOAD
          Cannot load the specified library
DATA #1 : Hex integer, 4 bytes
0x00000002
DATA #2 : String, 58 bytes
/home/hannigan/sql/lib/security64/plugin/server/dsmdb2pw.so
CALLSTCK:
[0] 0x00002AEF63DD267E pdOSSeLoggingCallback + 0x20C
[1] 0x00002AEF68486A42 /home/hannigan/sql/lib/lib64/libdb2osse.so.1 + 0x1C4A42
[2] 0x00002AEF6848825E ossLog + 0xA6
[3] 0x00002AEF684928E9 _ZN11OSSHLlibrary4loadEPKcm + 0x1D3
[4] 0x00002AEF63F63BDC _Z20secLoadPluginGenericP19SEC_PLUGIN_HANDLE_Tpc + 0x68
[5] 0x00002AEF63F62FBB _Z23secLoadServerAuthPluginP19SEC_PLUGIN_HANDLE + 0x57
[6] 0x00002AEF63F6C833 _Z25sqllexLoadAllPluginsServerP5sqlca + 0x3B5
[7] 0x00002AEF6431737C /home/hannigan/sql/lib/lib64/libdb2e.so.1 + 0x123637C
[8] 0x00002AEF643164C5 sqloRunInstance + 0x191
[9] 0x00000000040D31D DB2main + 0xD41

2011-04-26-16.04.11.825930-420 I2346711E1178 LEVEL: Error
PID : 25178 TID : 47207843621184PROC : db2sysc 0
INSTANCE: hannigan NODE : 000
EDUID : 1 EDUNAME: db2sysc 0
FUNCTION: DB2 Common, OSSe, OSSHLlibrary::load, probe:90
MESSAGE : ECF=0x90000076=-1879048074=ECF_LIB_CANNOT_LOAD
          Cannot load the specified library
DATA #1 : String, 109 bytes
../shared/gskit8/lib/linux64_x86/libgsk8iccs_64.so: cannot open shared object
file: No such file or directory
CALLSTCK:
[0] 0x00002AEF63DD267E pdOSSeLoggingCallback + 0x20C
[1] 0x00002AEF68486A42 /home/hannigan/sql/lib/lib64/libdb2osse.so.1 + 0x1C4A42
[2] 0x00002AEF6848825E ossLog + 0xA6
[3] 0x00002AEF6849294D _ZN11OSSHLlibrary4loadEPKcm + 0x237
[4] 0x00002AEF63F63BDC _Z20secLoadPluginGenericP19SEC_PLUGIN_HANDLE_Tpc + 0x68
[5] 0x00002AEF63F62FBB _Z23secLoadServerAuthPluginP19SEC_PLUGIN_HANDLE + 0x57
[6] 0x00002AEF63F6C833 _Z25sqllexLoadAllPluginsServerP5sqlca + 0x3B5
[7] 0x00002AEF6431737C /home/hannigan/sql/lib/lib64/libdb2e.so.1 + 0x123637C
[8] 0x00002AEF643164C5 sqloRunInstance + 0x191
[9] 0x00000000040D31D DB2main + 0xD41
```

At startup, the Tivoli Storage Manager server detects these types of errors and tries to remove the plug-in from the configuration. If the server cannot remove the plug-in, you must remove it from the database manager configuration. This command removes the plug-in from the database manager configuration:

```
db2 get database manager configuration | grep SRVCON_PW_PLUGIN
db2 update database manager configuration using SRVCON_PW_PLUGIN \"\"
```

Tracing the User ID/Password plug-in

If set up correctly, the Tivoli Storage Manager server can automatically trace the user ID and password plug-in (dsmdb2pw).

AIX HP-UX Linux Solaris

Add the following text to the `~instance/sqllib/userprofile` file and restart DB2:

```
EXPORT DB2_TSM_DEBUG_PLUGIN=YES
```

After DB2 restarts, trace output is stored in the `dsmdb2pw.trc` file in the `~/sqllib/db2dump/` directory.

Windows

If the `DB2_VENDOR_INI` db2set variable is set, find the file that it points to and add the following line: `set DB2_TSM_DEBUG_PLUGIN=YES`. If the `DB2_VENDOR_INI` db2set variable is not set, create a file, add `set DB2_TSM_DEBUG_PLUGIN=YES` to the file. Then issue the following command:

```
db2set -i instance DB2_VENDOR_INI=path to file just created.
```

After you restart DB2, the trace file is stored in the `dsmdb2pw.trc` file in the directory pointed to by the `DB2_VENDOR_INI` db2set variable.

Resolving DB2 automatic database backup errors

Error messages ANR9999D and ANR4581W might occur when you back up the IBM Tivoli Storage Manager (version 6.1.2 and later) database directly from DB2.

If you initiate a backup directly from DB2 either manually or automatically by using DB2 automatic maintenance, the backup fails.

The Tivoli Storage Manager server database must be backed up from a Tivoli Storage Manager administrative interface. For example, you can issue the **BACKUP DB** command from the command-line interface or include the **BACKUP DB** command in a Tivoli Storage Manager administrative schedule.

Refer to the wiki article ([http://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli Storage Manager/page/Tivoli Storage Manager V6.1.2 automatic DB backup](http://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli%20Storage%20Manager/page/Tivoli%20Storage%20Manager%20V6.1.2%20automatic%20DB%20backup)) for technical details about automatic database backups initiated from Tivoli Storage Manager.

Limiting DB2 memory allocation

You might discover that DB2 is using a large amount of memory. You can limit the amount of memory that DB2 uses by issuing the **db2 update** command.

By default, DB2 is installed and configured to use automatic memory management, which causes DB2 to use a large percentage of the physical memory. To restrict the amount of memory, use the **DB2 UPDATE** command to specify the memory limit. For example,

```
db2 update dbm cfg using instance_memory number_of_4_KB_blocks
```

In the preceding example, using 4 KB blocks, the following command limits the DB2 instance to use only 3,200,000 KB of memory:

```
db2 update dbm cfg using instance_memory 800000
```

For more information about DB2 instance memory configuration, see the DB2 information center.

Retrieving DB2 version information

The version of DB2 that is installed with the IBM Tivoli Storage Manager server is updated periodically.

AIX **HP-UX** **Linux** **Solaris** Issue the **DB2LS** command to show where DB2 products are installed on your server, and to list the DB2 product level. See the following example command results in the output that is displayed:

```
>/opt/tivoli/tsm/db2/install/db2ls
```

Install Path	Level	Fix Pack	Special Install	Install Date	Installer UID
/opt/IBM/db2/V9.7	9.7.0.1	1	1	Wed Jan 27 20:22:09 2010 MST	0
/opt/tivoli/tsm/db2	9.7.0.1	1	3	Wed Mar 31 15:29:57 2010 MST	0

Windows Issue the **DB2LEVEL** command to show where DB2 products are installed on your server, and to list the DB2 product level. See the following example command results in the output that is displayed for Windows operating systems:

```
C:\>db2level
DB21085I Instance "SERVER1" uses "64" bits and DB2 code release "SQL09054"
with level identifier "06050107".
Informational tokens are "DB2 v9.5.401.3564", "special_22833", "WR21456_22833",
and Fix Pack "4".
Product is installed at "C:\PROGRA~1\Tivoli\TSM\db2" with DB2 Copy Name "DB2TSM1".
```

Locating DB2 diagnostic log files

The `db2diag.log` file contains diagnostic information that can help you to resolve problems that might occur with your database.

The location of the `db2diag.log` file and the DB2 first occurrence data capture (FODC) files depends on your DB2 configuration settings or the DB2 environmental variable settings. DB2 writes messages about internal operations, events, or status in the administration notification log file (`db2SID.nfy`).

AIX **HP-UX** **Linux** **Solaris** Complete the following steps to determine where the diagnostic data directory path is located:

1. Log in as the server user instance.

2. Issue the following command:

```
db2 get dbm cfg | grep DIAGPATH
```

If no path is specified in the **DIAGPATH** configuration parameter, the diagnostic data directory is in the `sqllib/db2dump` subdirectory of your instance directory. For example, `/home/tsminst1/sqllib/db2dump` where `/home/tsminst1` is the instance home directory.

Windows Complete the following steps to determine where the diagnostic data directory path is located:

1. Stop the DB2 interactive mode. Start a DB2 command-line prompt and issue the `quit` command.
2. Find the path by using the **DIAGPATH** configuration parameter. Issue command

```
db2 get dbm cfg | findstr /s /i diagpath
```

If no path is specified in the **DIAGPATH** configuration parameter, the `DB2INSTPROF` directory path is used.

3. Find the path that was set in the `DB2INSTPROF` environment variable. Issue the following command from the DB2 command-line prompt:

```
db2set db2instprof
```

The first section of output from this command shows the diagnostic data directory path:

```
[DB2INSTPROF_path]\[DB2_copy_name]\[Instance_name]
```

The output translates into the following directory path:

```
C:\Documents and Settings\All Users\Application Data\IBM\DB2\DB2COPY1\V620
```

where

- `DB2INSTPROF_path` is `C:\Documents and Settings\All Users\Application Data\IBM\DB2`
- `DB2_copy_name` is `DB2COPY1`
- `Instance_name` is `V630`

If no path is specified in the `DB2INSTPROF` environment variable, the `DB2PATH` path set is used. Issue the **db2set db2path** command from a command prompt to find the path that was set in the `DB2PATH` environment variable. The diagnostic data directory path is the first part of the output from the **db2set db2path** command, with `DB2DAS00\dump` added. For example:

```
[DB2PATH]\[DB2DAS00]\dump
```

The output translates into the following directory path:

```
C:\Program Files\IBM\SQLLIB
```

where

- `DB2PATH` is `C:\Program Files\IBM\SQLLIB`
- The diagnostic data directory path is `C:\Program Files\IBM\SQLLIB\DB2DAS00\dump`

Related concepts:

“sqllib/db2dump directory causes shutdown” on page 67

Related reference:

“Installation log files” on page 53

Locating DB2 log files after an upgrade

When upgrading the server from IBM Tivoli Storage Manager V6.1 to V6.2 or later, a DB2 script, `DB2CKUPGRADE`, runs to conduct checks and create log files for server databases. The log files contain the results from the `DB2CKUPGRADE` command for each database.

Purpose

The following log files are created during an upgrade:

- `AIX` `HP-UX` `Linux` `Solaris` `/tmp/db2ckupgrade_instance_name_db_name.log`
- `Windows` `installation_directory\db2ckupgrade_instance_name_db_name.log`

The wizard automatically corrects some errors in a database during the upgrade to Version 6.2 or later and DB2 V9.7 or later. Other errors might need to be corrected manually.

Check the `DB2CKUPGRADE` log files if the following are true:

- You receive an error message about the database when the wizard is running the `DB2CKUPGRADE` script.
- You have to cancel or close the wizard or check the `log.text` file during a silent installation.

See the DB2 information center for details about possible error messages found in the listed log files: <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp>

Upgrade again after fixing any errors.

Resolving a missing or incorrect database ID file problem

If you restore a database to a different server after a disaster, the database ID file (`dsmserv.dbid`) might not be restored. The IBM Tivoli Storage Manager server, therefore, cannot find the file after the restore operation and cannot start.

After upgrading from Tivoli Storage Manager version 6.1 to 6.2, you might have difficulty in restoring Tivoli Storage Manager version 6.1 databases. You must start the Tivoli Storage Manager version 6.2 server to generate a new backup image in DB2. After Tivoli Storage Manager version 6.2 server initializes, a database backup is started automatically. When the backup completes, stop the server and issue the `RESTORE DB` command. If the automatic database backup does not complete successfully, resolve the problem and issue the `BACKUP DB` command. Ensure that it completes before issuing the `RESTORE DB` command.

Important: You must have a successful database backup image generated by the Tivoli Storage Manager version 6.2 server for incremental database backups or database restores to be successful.

If you started the upgraded Tivoli Storage Manager version 6.2 server and the automatic database backup completed successfully, you can drop the database before restoring it. You must not drop the database immediately after upgrading to version 6.2. If you drop the database before a backup image is generated, you must reinstall Tivoli Storage Manager version 6.1 server and then restore the database.

If you must restore a Tivoli Storage Manager version 6.1 database and the database does not exist, you must restore it through Tivoli Storage Manager version 6.1. You can then upgrade to Tivoli Storage Manager version 6.2.

A lost or incorrect dbid file can affect starting the Tivoli Storage Manager server after a database restore operation.

When a database is restored, the database ID file must stay in sync with the database. With Tivoli Storage Manager version 6.2, if you format the database before you restore it, the database ID file changes. This change causes a mismatch of the date and time in the database and keeps the server from starting.

If your database ID file is causing errors during a restore operation, you might have to use the `-S` (skip DB ID check) parameter. The `dsmserv.dbid` file must be absent from your server when you use the `-S` parameter. The following situations describe where the `-S` parameter is useful:

- If you reformat the server after backing it up, you will have mismatched the date and time that is stored in the new `dsmserv.dbid` file. Use the `-S` parameter when you start the server after restoring.
- When the `dsmserv.dbid` file gets damaged or lost.

After the initial use of the `-S` parameter in a restore scenario, the server creates a `dsmserv.dbid` file in the instance directory.

Resolving problems with the BACKUP DB and the RESTORE DB commands

The IBM Tivoli Storage Manager server **BACKUP DB** and **RESTORE DB** commands request the IBM DB2 database application to back up the Tivoli Storage Manager database to the server.

Backup data is then sent to the Tivoli Storage Manager server through the Tivoli Storage Manager client application programming interface (API).

When a **BACKUP DB** or **RESTORE DB** command fails with a DB2 SQLCODE or a SQLERRMC message with return codes, get a description of the DB2 SQLCODE by completing the following procedures:

1. Open a DB2 command-line interface:

Windows For Windows, click **Start > All Programs > IBM DB2** and click **Command Line Tools > Command Line Processor**.

AIX **Linux** **HP-UX** **Solaris** For all other supported platforms, log on to the DB2 instance ID and open a shell window, then issue the command `DB2`.

2. Enter the SQLCODE. For example, if the DB2 SQLCODE is -2033, issue the following command:

```
? sql2033
```

You can use the details of the error condition to debug the problem with the **BACKUP DB** or **RESTORE DB** command. If the SQLERRMC code is also displayed, it is

explained in the SQLCODE description that you are provided. You can find more information about the API return codes through the following files:

- **Windows** tsm\api\include\dsmrc.h
- **AIX** **HP-UX** **Linux** **Solaris** tsm/client/api/bin64/sample/dsmrc.h

Resolving incorrect environment variables for BACKUP DB and RESTORE DB

Many of the **BACKUP DB** or **RESTORE DB** processing problems are as a result of incorrectly set DSMI_CONFIG or DSMI_DIR,DSMI_LOG environment variables.

Requirement:

The environment variables are used by the Tivoli Storage Manager client API to locate API codes and the options files. The DB2 instance must be running in a shell with correctly set environmental variables.

AIX **HP-UX** **Linux** **Solaris** For AIX, HP-UX, Linux, or Solaris, the DSMI_* variables are set in the instance's userprofile file. For example: /home/tsminst1/sql1lib/userprofile

Windows On Windows systems, the DSMI_* variables are set in the file that the DB2 instance registry variable, DB2_VENDOR_INI, points to. For example, this file might be c:\tsminst1\tsmdbmgr.env. You can verify the file name and location by issuing the db2set -i tsminst1 DB2_VENDOR_INI command, where tsminst1 is the DB2 instance.

The DSMI_* variables are initially set up automatically by the Tivoli Storage Manager instance configuration wizard. They can be set up manually, though, as described in the *IBM Tivoli Storage Manager Installation Guide*.

AIX **HP-UX** **Linux** **Solaris** Use the following example procedure to debug DSMI environment variable errors. The SQL error code (SQLERRMC) 409 is shown when you issue the **DSMAPIPW** command.

1. Open the /home/tsminst1/sql1lib/userprofile file and review the statements. If you change this file, stop and restart the DB2 instance so that the changes are recognized. For example, consider the following scenario. The userprofile file has statements like the following example text:

```
export DSMI_CONFIG=/home/tsminst1/tsminst1/tsmdbmgr.opt
export DSMI_DIR=/usr/tivoli/tsm/client/api/bin64
export DSMI_LOG=/home/tsminst1/tsminst1
```

The /home/tsminst1/tsminst1/tsmdbmgr.opt file has the following text:
SERVERNAME TSMDBMGR_TSMINST1

The /usr/tivoli/tsm/client/api/bin64/dsm.sys file has the following text:

```
SERVERNAME TSMDBMGR_TSMINST1
commethod tcpip
tcpserveraddr localhost
errorlogname /home/tsminst1/tsminst1/tsmdbmgr.log
```

Verify that the SERVERNAME entry in the tsmdbmgr.opt file matches the SERVERNAME entry in the dsm.sys file.

2. Run the **DSMAPIPW** command. You must be logged on using the root user ID.

3. If you can run the **DSMAPIPW** command, remove the `/home/tsminst1/tsminst1/tsmdbmgr.log` file. Remove the file while still using the root user ID to eliminate permission problems between the root user ID and the `tsminst1` user ID.

Resolving error message ANR2968E

Error message ANR2968E is surfaced during the **BACKUP DB** command.

There are two causes for this error message:

- If the Tivoli Storage Manager error log file is owned by the root user ID rather than the server instance user ID.
- **Windows** If you use quotation marks to surround the paths that are in the `tsmdbmgr.env` file. Use a path that does not contain spaces or use the Windows short name for the path.

To correct the error caused by the root user ID, complete the following steps:

1. Log on using a Tivoli Storage Manager server instance ID and verify the name of the error log file. For example:

```
$ grep -i "ERRORLOGNAME" $DSMI_DIR/dsm.sys
ERRORLOGNAME /home/db2inst1/tsminst1/tsmdbmgr.log
```

where `db2inst1` is the server instance user ID and `/home/db2inst1/tsminst1/` is the server instance directory.

2. Issue the following example command to verify the current owner of the error log file:

```
$ ls -la /home/db2inst1/tsminst1/tsmdbmgr.log
-rw-r--r-- 1 root system 834 May 05 09:43 /home/db2inst1/tsminst1/tsmdbmgr.log
```

3. If the error log file is not owned by the Tivoli Storage Manager instance user ID, remove it. You must have root authority to remove the file. Issue the following example command to remove the log file:

```
$ su root password
# rm /home/db2inst1/tsminst1/tsmdbmgr.log
# exit
```

4. Issue the **BACKUP DB** command and verify that the command completed successfully. Verify that the log file is owned by the server instance ID. For example:

```
$ ls -la /home/db2inst1/tsminst1/tsmdbmgr.log
-rw-r--r-- 1 db2inst1 db2iadml 834 May 05 09:50
/home/db2inst1/tsminst1/tsmdbmgr.log
```

Troubleshooting error message ANR2971E using the SQL code

Error message ANR2971E might show when you are restoring or backing up a database operation, and the process stops. Use the SQL code attached to the error to help you resolve this problem.

If you are restoring a database because the server stopped during normal operation, review the `db2diag.log` file *before* backing up or restoring the database.

The following message can be issued when you are restoring or backing up data:
ANR2971E Database backup/restore/rollforward terminated - DB2 sqlcode -2581 error

In the following scenario, the **DSMSERV RESTORE DB** process failed with a DB2 SQL 2581 code. This following scenario does not pertain to problems with the DSMI environment variables.

1. Issue the following command from the DB2 command-line interface:

? SQL2581

An explanation is generated about the SQL code.

SQL2581N Restore is unable to extract log files or restore a log directory from the backup image to the specified path. Reason code 2581

2. Review the db2diag.log file where you can find status and error messages. A portion of the db2diag.log file is displayed in the following example:

```
2009-02-10-09.49.00.660000-300 E8120712F500      LEVEL: Info
PID      : 4608                TID   : 3956      PROC  : db2syscs.exe
INSTANCE: SERVER1            NODE  : 000        DB    : TSMDB1
APPHDL  : 0-7                APPID : *LOCAL.SERVER1.090210144859
AUTHID   : B1JRP01
EDUID    : 3956                EDUNAME: db2agent (TSMDB1)
FUNCTION: DB2 UDB, database utilities, sqludPrintStartingMsg, probe:1292
DATA #1 : <preformatted>
Starting a full database restore.
Agent EDU ID: 3956
```

```
2009-02-10-09.50.21.051000-300 E8123213F483      LEVEL: Severe
PID      : 4608                TID   : 5080      PROC  : db2syscs.exe
INSTANCE: SERVER1            NODE  : 000
EDUID    : 5080                EDUNAME: db2bm.3956.1 (TSMDB1)
FUNCTION: DB2 UDB, database utilities, sqluWriteLogFile, probe:1498
MESSAGE  : ZRC=0x850F000C=-2062614516=SQL0_DISK "Disk full."
          DIA8312C Disk was full.
DATA #1 : String, 46 bytes
F:\tivoli\tsm\Beta\sarch\RstDbLog\S0000262.LOG
```

```
2009-02-10-09.50.21.051000-300 E8124165F912      LEVEL: Severe
PID      : 4608                TID   : 5080      PROC  : db2syscs.exe
INSTANCE: SERVER1            NODE  : 000
EDUID    : 5080                EDUNAME: db2bm.3956.1 (TSMDB1)
FUNCTION: DB2 UDB, database utilities, sqluWriteLogFile, probe:1500
MESSAGE  : SQL2581N Restore is unable to extract log files or restore a log
          directory from the backup image to the specified path. Reason code "".
DATA #1 : SQLCA, PD_DB2_TYPE_SQLCA, 136 bytes
sqlcaid : SQLCA      sqlcab: 136      sqlcode: -2581      sqlerrml: 1
sqlerrmc: 4
sqlerrp : sqluWrit
sqlerrd : (1) 0x00000000      (2) 0x00000000      (3) 0x00000000
          (4) 0x00000000      (5) 0x00000000      (6) 0x00000000
sqlwarn : (1)      (2)      (3)      (4)      (5)      (6)
          (7)      (8)      (9)      (10)     (11)
sqlstate:
```

The preceding example shows from the “Disk Full” message that there was not enough disk space to place the needed log files from the backup operation.

3. Add disk space and run the operation again.

Common BACKUP DB and RESTORE DB errors

Common errors that are derived from **BACKUP DB** or **RESTORE DB** commands might include SQL return or error codes.

The following errors are some of the more common errors that are displayed when issuing the **BACKUP DB** or **RESTORE DB** commands:

ANR2968E - Database backup terminated DB2 SQLCODE -2033 SQLERRMC 406

SQL error message code 406 requires that the following issues are resolved:

- The DSMI_CONFIG environment variable points to a valid IBM Tivoli Storage Manager options file.
- The instance owner has read access to the dsm.opt file.
- The DSMI_CONFIG environment variable is set in the db2profile.

DB2 SQLCODE: -2033, DB2 SQLERRMC: 106

SQL error message code 106 can mean that there is a permissions problem with the log file that the Tivoli Storage Manager client application programming interface (API) writes.

AIX **HP-UX** **Linux** **Solaris** On AIX, HP-UX, Linux, or Solaris servers, this error is flagged because the **DSMAIPW** command is using the root user ID but the DB2 instance is using a different ID. To resolve the problem, find the log file with the permissions problem and delete it using the root user ID.

DB2 SQLCODE: -2033, DB2 SQLERRMC: 168

Verify that the DSMI_DIR environment variable is pointing to the client API executable directory that contains the trusted communication agent (dsmtca). Verify that the TSM.PWD file exists in the Tivoli Storage Manager server instance directory. Typically, this file is pointed to by the passworddir option in the tsmdbmgr.opt file.

ANR2971E - Database backup/restore/rollforward terminated DB2 SQLCODE - 2071 error.

The library could not be loaded because it (or a library required by it) does not exist or does not have a valid format. This typically means that a 32-bit library is being loaded in a 64-bit instance, or vice versa. This also indicates that the DSMI_DIR environment variable is pointing to the wrong Tivoli Storage Manager client API executable files. More information about the error is available by opening a DB2 command-line processor window and issuing the following command:

```
db2 => ? sql2071
```

Verify that if any changes were made to the tsmdbmgr.opt, dsm.sys, or sqllib/userprofile files, the DB2 instance is recycled so that it picks up the new values. To recycle the DB2 instance, stop and restart the Tivoli Storage Manager server. Also, verify that the **EXPORT** command precedes the DSMI_*= entries in the file sqllib/userprofile. DB2 might run under a different shell than Telnet, so while the **DSMAIPW** command works from a Telnet shell, the DB2 instance does not.

Error message indicates that the node is locked

You might experience an error when DB2 contacts the server and a particular node, but gets an error stating that the node is locked.

To correct the error, specify the tcpserveraddress localhost in the SERVERNAME TSMDBMGR_TSMINST1 stanza of the dsm.sys file.

Do not specify the dotted IP address of the computer, for example 127.0.0.1.

Characteristics of the \$\$_TSMDBMGR_\$\$ user ID

The Tivoli Storage Manager server generates the \$\$_TSMDBMGR_\$\$ user ID at startup.

You can view the \$\$_TSMDBMGR_\$\$ user ID in the results of a **QUERY SESSION** command. This ID is also present in the activity log file and other server log files.

The Tivoli Storage Manager server uses the \$\$_TSMDBMGR_\$\$ user ID to back up the server database. By using the \$\$_TSMDBMGR_\$\$ user ID, you can make the database accessible for processing if the server is unavailable. Changing this ID jeopardizes the ability to recover or restore a server if a disaster occurs.

Restriction: You cannot change the `dsm.sys` or `dsm.opt` file to set up or use a different client node name. The local Tivoli Storage Manager server database uses the `dsm.sys` or `dsm.opt` file to back up its own database.

Resolving database reorganization problems

Database table reorganization and index reorganization require a significant amount of system resources. To avoid occupying system resources that can be used elsewhere, run your reorganization routines on off times.

Unexpected database growth and unexpected active and archive log space requirements can occur if tables or the indexes associated with tables are not reorganized. Tivoli Storage Manager reorganizes tables by default. If automatic reorganization is affecting server performance, you can schedule reorganization manually.

The following suggestions might help when you configure your reorganization:

- Turn on index reorganization if you are running deduplication on your server. See the server option `ALLOWREORGINDEX`.
- By default, table reorganization is turned on 24-hours-a-day. Run reorganization during a relatively idle time during the day. See the following server options for defining an idle time when reorganization can run:
 - `REORGBEGINTIME`
 - `REORGDURATION`

Analyzing the process symptoms to resolve problems

You can sometimes determine the cause of errors by observing the process symptoms.

You might encounter one of the following process symptoms:

- Insufficient space in a target copy storage pool
- Damaged file found on volume
- Files are not expired after reducing the number of versions that need to be kept
- Migration does not run for sequential media storage pool
- Migration only uses one process
- Process running slow

Reviewing process messages to determine the state of server operations

Server processes, whether run in the foreground or background, will always issue a “process started” message and a “process ended” message in addition to the general process messages. You can use these messages to determine the status of your server operation.

Processes that run on the server

A server process is a task that is performed on the server. You can assign the task to perform a specific operation, such as migrating data from a storage pool to the next storage pool in the hierarchy. Issue the server processes to resolve problems that you are having with your server.

Server processes are typically initiated as an automated process on the server. The process might or might not be influenced by a server option or other setting. The server process can also be started by a command.

Many server processes can be run in the FOREGROUND or synchronously. Processes that run in the FOREGROUND can be initiated by a command using the WAIT=YES parameter. Commands that start server processes that do not allow the WAIT=YES parameter or commands specified with WAIT=NO are run in the BACKGROUND or asynchronously.

Some server processes can initiate multiple processes simultaneously to accomplish the task. See Table 13 for the descriptions of the server processes.

Table 13. IBM Tivoli Storage Manager server processes

Process or command	Description	Runs in the foreground or as a multiple process
AUDIT VOLUME	Audit the contents of a volume to validate that the data can still be read and that the server database definitions describing the data are correct.	
BACKUP DB	Back up the server database (FULL or INCREMENTAL).	The BACKUP DB can run as a synchronous process by specifying WAIT=YES .
BACKUP STGPOOL	Back up a primary server storage pool to a copy storage pool. The result is that you can make duplicate copies of the data and potentially take duplicate copies to an off-site location.	The BACKUP STGPOOL can run as a synchronous process by specifying WAIT=YES . BACKUP STGPOOL might be run using multiple concurrent processes, which is controlled by the MAXPROCESS parameter specified on the BACKUP STGPOOL command.
CHECKIN LIBVOLUME	Check a tape volume into a tape library.	
CHECKOUT LIBVOLUME	Check a tape volume out from a tape library.	

Table 13. IBM Tivoli Storage Manager server processes (continued)

Process or command	Description	Runs in the foreground or as a multiple process
Expiration	<p>Delete client backup and archive files from the server, based on the policies defined to manage those files.</p> <p>You can run expiration automatically by specifying <code>EXPINTERVAL=<i>n</i></code> in the server options file, where <i>n</i> is any number other than zero. Expiration can also be initiated by issuing the EXPIRE INVENTORY command. It is not possible to have more than one expiration process running at a time, although you can run more than one thread at a time.</p>	The EXPIRATION command can run as a synchronous process by specifying <code>WAIT=YES</code> .
IMPORT	<p>Import data from sequential media volumes or directly from another server using TCP/IP communication connections between the servers.</p> <p>Import processing can be started by any of the following commands:</p> <ul style="list-style-type: none"> • IMPORT ADMIN • IMPORT NODE • IMPORT POLICY • IMPORT SERVER 	
LABEL LIBVOLUME	Label one or more library volumes in a library.	
Migration	<p>Migrate data from one storage pool to the next in the storage hierarchy.</p> <p>Migration starts and stops, based on the HighMig and LowMig thresholds defined for the storage pool. Whenever UPDATE STGPOOL is issued, these values are reexamined and, if appropriate, MIGRATION is started. Otherwise, the server monitors the percentage utilization for non-migrated data in a storage pool. As the server needs, it starts migration processing for that storage pool when the HighMig threshold is exceeded. You can also issue the MIGRATE STGPOOL command to manually start migration processing.</p>	Migration might be configured to run multiple concurrent processes. The multiple processes are controlled by the <code>MIGPROCESS</code> attribute of the storage pool and might be updated by issuing the UPDATE STGPOOL command.
MOVE DATA	Move data from one volume to other volumes in the same storage pool or to a different storage pool.	The MOVE DATA command can run as a synchronous process by specifying <code>WAIT=YES</code> .

Table 13. IBM Tivoli Storage Manager server processes (continued)

Process or command	Description	Runs in the foreground or as a multiple process
MOVE DRMEDIA	Manage the disaster recovery media by moving on-site volumes off-site, or by bringing back off-site volumes, on-site. Disaster recovery media is the database backup and storage pool backup volumes necessary to protect and recover the server.	The MOVE DRMEDIA command can run as a synchronous process by specifying WAIT=YES .
MOVE MEDIA	Move volumes from a tape library to the overflow location to prevent a library from becoming full.	
MOVE NODEDATA	Move all the data for the node or nodes specified to other volumes in the same storage pool or to a different storage pool.	The MOVE NODEDATA command can run as a synchronous process by specifying WAIT=YES .
PREPARE	Create a recovery plan file.	The PREPARE command can run as a synchronous process by specifying WAIT=YES .
Reclamation	<p>Reclaim space from tape volumes by moving active data to other volumes and returning the volume back to empty and private, or else back to scratch.</p> <p>The server monitors the RECLAMATION THRESHOLD defined for a storage pool. It starts a reclamation process for that storage pool to reclaim any eligible volumes if it determines that one or more eligible volumes exist.</p>	
RESTORE STGPOOL	Restore all files for a given storage pool from a copy storage pool.	The RESTORE STGPOOL can run as a synchronous process by specifying WAIT=YES . RESTORE STGPOOL might be run using multiple concurrent processes, which is controlled by the MAXPROCESS parameter specified on the RESTORE STGPOOL command.
RESTORE VOLUME	Restore all files for a given volume from a copy storage pool.	The RESTORE VOLUME command can run as a synchronous process by specifying WAIT=YES . RESTORE VOLUME might be run using multiple concurrent processes, which is controlled by the MAXPROCESS parameter specified on the RESTORE VOLUME command.

Messages issued when processes start

When the server runs tasks as processes, the processes are assigned an identification message and report that they have started.

The reported start is issued in the following message:

```
ANR0984I Process process_id for process_name started in the process_state at time
```

The following list defines the variables from this message:

process_id

Numeric process identifier.

process_name

The name of the process.

process_state

FOREGROUND or **BACKGROUND**. If the process is running in the foreground, the command was issued with the **WAIT=YES** parameter. Foreground processing causes the administrative session that issued the command to wait until the processing completes. A process running in the background returns immediately to the administrative session that issued the command, indicating that a process was started while the process still runs. Processes running in the background might be monitored with the **QUERY PROCESS** command.

time

The time that the process was started.

Messages issued when processes end

When the server runs tasks as processes, the processes report when they end. The “process ended” messages that are issued vary from process to process. The message depends on whether the process must report about items and bytes processed, no items or bytes processed, items processed, or just bytes processed.

Process ended

When a process completes and it does not have bytes or number of files to report, the following message is issued:

```
ANR0985I Process process_id for process_name running in the process_state completed with the completion_state at time
```

The following list defines the variables from this message:

process_id

Numeric process identifier.

process_name

The name of the process.

process_state

FOREGROUND or **BACKGROUND**. If the process is running in the foreground, the command was issued with the **WAIT=YES** parameter. Foreground processing causes the administrative session that issued the command to wait until the processing completes. A process running in the background returns immediately to the administrative session that issued the command, indicating that a process was started while the process still runs. Processes running in the background can be monitored with the **QUERY PROCESS** command.

completion_state

SUCCESS or FAILURE.

time The time that the process was started.

Process ended with items and bytes

When a process completes and has bytes and items processed to report, the following message is issued:

```
ANR0986I Process process_id for process_name running in the process_state
processed number_of_items items for a total of bytes_processed bytes with a
completion state completion_state at time
```

The following list defines the variables from this message:

process_id

Numeric process identifier.

process_name

The name of the process.

process_state

FOREGROUND or **BACKGROUND**. If the process is running in the foreground, the command was issued with the **WAIT=YES** parameter. Foreground processing causes the administrative session that issued the command to wait until the processing completes. A process running in the background returns immediately to the administrative session that issued the command, indicating that a process was started while the process still runs. Processes running in the background can be monitored with the **QUERY PROCESS** command.

number_of_items

The number of items processed.

bytes_processed

The number of bytes processed.

completion_state

SUCCESS or FAILURE.

time The time that the process was started.

Process ended with items

When a process completes and has items processed to report, the following message is issued:

```
ANR0987I Process process_id for process_name running in the process_state
processed number_of_items items with a completion state completion_state at time
```

The following list defines the variables from this message:

process_id

Numeric process identifier.

process_name

The name of the process.

process_state

FOREGROUND or **BACKGROUND**. If the process is running in the foreground, the command was issued with the **WAIT=YES** parameter. Foreground processing causes the administrative session that issued the command to wait until the processing completes. A process running in the background returns immediately to the administrative session that issued the command,

indicating that a process was started while the process still runs. Processes running in the background can be monitored with the **QUERY PROCESS** command.

completion_state
SUCCESS or FAILURE.

time The time that the process was started.

Process ended with bytes

When a process completes and has bytes processed to report, the following message is issued:

```
ANR0988I Process process_id for process_name running in the process_state
processed bytes_processed bytes with a completion state completion_state at time
```

The following list defines the variables from this message:

process_id
Numeric process identifier.

process_name
The name of the process.

process_state
BACKGROUND or **BACKGROUND**. If the process is running in the foreground, the command was issued with the **WAIT=YES** parameter. Foreground processing causes the administrative session that issued the command to wait until the processing completes. A process running in the background returns immediately to the administrative session that issued the command, indicating that a process was started while the process still runs. Processes running in the background can be monitored with the **QUERY PROCESS** command.

bytes_processed
The number of bytes processed.

completion_state
SUCCESS or FAILURE.

time The time that the process was started.

Analyzing the ANR1221E error message

When you receive error message ANR1221E, the cause is typically due to insufficient space in the target copy storage pool.

Perform the following steps to resolve error message ANR1221E:

1. Issue the **QUERY STGPOOL** *stgpoolName* **F=D** command.
2. Issue the following SQL select statement from an administrative client to this server: "select *stgpool_name*,*devclass_name*,count(*) as 'VOLUMES' from volumes group by *stgpool_name*,*devclass_name*."
3. Compare the number of volumes reported by the select statement to the maximum scratch volumes allowed (as reported by the **QUERY STGPOOL** command). If the number of volumes reported by the **select** is equal to or exceeds the "Maximum Scratch Volumes Allowed," update the storage pool and allow more scratch volumes. If scratch volumes are not used in the storage pool (scratch=0), then ensure that you add more private volumes. Issue the **UPDATE STGPOOL** *stgpoolName* **MAXSCR=nn** command, where *stgpoolName* is the

name of the storage pool to update and *nn* is the increased number of scratch volumes to make available to this copy storage pool.

Important: The tape library should have this additional number of scratch volumes available, or you need to add scratch volumes to the library prior to issuing this command and retrying the **BACKUP STGPOOL** operation.

Analyzing the ANR2317W error message

The ANR2317W error message is issued when a process determines that there is a damaged file.

The message is shown with the following information:

```
ANR2317W Audit Volume found damaged file on volume volumeName: Node nodeName,  
Type fileType, File space fileSpaceName, fsId fileSpaceID,  
File name fileName is number version of totalVersions versions.
```

Perform the following steps to resolve error message ANR2317W:

1. Issue the **QUERY VOLUME** *volumeName* **F=D** command.
2. Issue the following SQL select statement from an administrative client to this server: "select* from VOLHISTORY where VOLUME_NAME='volume_name' AND TYPE='STGNEW.'" The results of the **QUERY VOLUME** command indicate when this volume was last written. The information from the **SELECT** operation reports when this volume was added to the storage pool. Often, **AUDIT VOLUME** might report files as damaged because at the time that the data was written, the hardware was malfunctioning and did not write the data correctly, even though it reported to the IBM Tivoli Storage Manager server that the operation was successful. As a result of this device malfunction, many files on many different volumes might be affected. Perform the following steps to resolve this issue:
 - a. Evaluate the system error logs or other information about this drive to determine if it still reports an error. If errors are still reported, they must first be resolved. To resolve a hardware issue, work with the hardware vendor to correct the problem.
 - b. If this storage pool is a copy of a storage pool volume, simply delete this volume using the **DELETE VOLUME** *volumeName* **DISCARDATA=YES** command. The next time a storage pool backup is run for the primary storage pool or storage pools where this damaged data resides, it will be backed up again to this copy storage pool and no further action is necessary.
 - If this storage pool is a primary storage pool volume and the data was written directly to this volume when the client stored the data, then it is likely that there are no undamaged copies of the data on the server. If possible, back up the files again from the client.
 - If this storage pool is a primary storage pool volume but the data was put on this volume by **MIGRATION**, **MOVE DATA**, or **MOVE NODEDATA** commands, there might be an undamaged copy of the file on the server. If the primary storage pool that contained this file was backed up to a copy storage pool prior to the **MIGRATION**, **MOVE DATA**, or **MOVE NODEDATA** commands, then an undamaged file might exist. If an undamaged file exists, issue the **UPDATE VOLUME** *volumeName* **ACCESS=DESTROYED** command and then issue the **RESTORE VOLUME** *volumeName* command to recover the damaged files for this volume from the copy storage pool.

Analyzing error messages ANR1330E and ANR1331E

You might receive error message ANR1330E or ANR1331E while data is being read from a IBM Tivoli Storage Manager storage pool volume.

When the Tivoli Storage Manager server stores data to a storage pool volume, self-describing information is inserted periodically throughout the data. This information is checked for validity while the server reads the data. Messages ANR1330E and ANR1331E are issued if the check reveals that the information is invalid. Error message ANR1330E displays the actual values that were read, and error message ANR1331E displays the values that were expected. The server issues these messages for the following reasons:

- The hardware (disk subsystem, tape drive) encountered a problem while reading the data
- An error occurred while writing the data and the data is damaged
- A database restore operation was performed and a volume was not appropriately audited so that it is in sync with the point-in-time (PIT) restore time.

You must first determine if data is damaged on the media or whether there was an error when the server read the intact data. Issue the following command for the volume on which the data is stored:

```
AUDIT VOLUME FIX=NO
```

If the audit reports no damaged files, Tivoli Storage Manager successfully read the data that was earlier reported as damaged. In this case, the error was caused by a temporary hardware malfunction when the server read the data. However, if the audit still reports that the data is damaged, determine what might have caused the damage.

You can ignore the error, but do so only if it occurs infrequently. Hardware occasionally encounters an error while reading data. In most cases, the hardware recognizes that an error occurred and recovers without having to report it. But there are times when the data is read in an altered (damaged) state because of a temporary hardware error. The following list defines the results of reading data and receiving an error:

Audit OK, error reading intact data on media

Tivoli Storage Manager checks the self-describing information and reports the data as damaged if it does not match what is expected. In messages ANR1330E and ANR1331E, the data is reported as damaged.

If after auditing the volume, messages ANR1330E and ANR1331E are displayed frequently, determine which hardware device is causing the data to be read incorrectly. Query the activity log to find the date and time that messages ANR1330E and ANR1331E were issued and provide the information to your hardware support team. With this information, they can examine the hardware error logs for any operations that might have completed abnormally. Also, have your hardware support team ensure that the device drivers and microcode maintenance for the hardware is up to date.

A common place for such errors to occur is on a storage area network (SAN). Typically, these errors occur if many link level interrupt (LLI) errors occur on the switch or the network. LLI errors indicate that the system is performing poorly and are known to cause data to be modified during retransmission. Ask your hardware support team to examine the network

error logs for instances of LLI errors. Look for LLI errors that were logged around the time that the ANR1330E and ANR1331E message were issued.

Audit failed, data damaged on media

If the audit reports the data as damaged, an error might have occurred that caused the data to be written incorrectly onto the media. Also, a database restore operation might have a volume that was not appropriately audited to synch-up with the PIT restore time. Determine, from the audit reports, when the data was written and examine message ANR1331E to find out which hardware device damaged the data. See the following example data:

ANR1330E

The server has detected possible corruption in an object being restored or moved. The actual values for the incorrect frame are: magic C6A2D75D hdr version 35134 hdr length 43170 sequence number 160421181 data length 7E53DCD8 server id 348145193 segment id 327643666840426461 crc 06E04914.

ANR1331E

Invalid frame detected. Expected magic 53454652 sequence number 00000023 server id 00000000 segment id 2062.

The segment ID number in message ANR1331E in this example is 2062. To determine the date that the data was inserted into the server, issue the following command:

```
SHOW INVO 0 2062
```

The following example shows the output from the **SHOW INVO** command:

```
OBJECT: 0.2062 (Backup):
Node: NODE1 Filespace: \\node1\c$ (Unicode).
\5400\BF\ BFDEFS.H
Type: 2 (File) CG: 1 Size: 0.89088 HeaderSize: 364
```

BACKUP OBJECTS ENTRY:

```
State: 1 Type: 2 MC: 1 CG: 1
\\node1\c$ (Unicode) : \TESTFILES\ FILE1.TXT (MC: DEFAULT)
Active, Inserted 11/29/2009 13:28:26
```

EXPIRING OBJECTS ENTRY:

Expiring object entry not found.

Find the Inserted field and note the date and time. In this example, the object was inserted on 11/29/2009 at 13:28:26. Provide your hardware support team with the date and time. The support team can examine the hardware error logs for any operations that completed abnormally. Also ask the support team to ensure that the device drivers and microcode maintenance for the hardware is up to date. Your hardware support team must examine the SAN network error logs. Look for errors around the time that the data was inserted into Tivoli Storage Manager.

If the **SHOW INVO** command returns unhelpful output, issue the following command to determine the date of insertion:

```
SHOW BFO 0 xxx
```

where xxx is the segment group ID. The example shows the output from the **SHOW BFO** command:

```
Bitfile Object: 0.xxx
**Super-bitfile 0.xxx contains following aggregated bitfiles
(offset/length)
0.2063 0.75295 0.3071 Active
0.2064 0.78366 0.88780 Active
0.2065 0.167146 0.13831 Active
0.2066 0.180977 0.21254 Active
0.2067 0.202231 0.3808 Active
```

0.2068 0.206039 0.11261 Active

```
**Disk Bitfile Entry
Bitfile Type: PRIMARY
Storage Format: 22
Logical Size: 0.217364
Physical Size: 0.221184
Number of Segments: 1,
Deleted: False
Storage Pool ID: 1
Volume ID: 2
Volume name: TapeVol1
```

Get an aggregated bit file number from the first entry on the list of aggregated bit files. In the preceding example, the first aggregated bit file number is 2063. Issue the **SHOW INVO** command using 2063.

No hardware errors at time of insertion

If the hardware support team discovers that no hardware errors occurred at the time the data was inserted into Tivoli Storage Manager, contact the IBM support team. Provide the team with the activity log at the time that messages ANR1330E and ANR1331E were issued. Also, issue the **AUDIT VOLUME FIX=NO** command with the following trace, and provide the Tivoli Storage Manager support team with the trace:

```
TRACE ENABLE BF AF DF SS AS DS SSFRAME
TRACE DISABLE BFLOCK AFLOCK SSLOCK
TRACE BEGIN filename
```

Fixing damaged files on media

If you find that the data is damaged on a volume, issue the **AUDIT VOLUME FIX=YES** command on the volume. If the following conditions are true, the data remains marked as damaged on the primary pool volume:

- The volume is a primary pool volume
- The data is backed up to a copy storage pool
- The data is damaged

After the **AUDIT VOLUME FIX=YES** command completes, issue the **RESTORE VOLUME** command for the primary pool volume. The damaged data is replaced with a new copy of the data. If the **AUDIT VOLUME FIX=YES** command successfully read the data, the data is no longer marked as damaged in the primary storage pool.

If there is no backup copy, the **AUDIT VOLUME FIX=YES** command deletes the data. If the data that was deleted is backup data, it is placed on the server the next time the client backup runs.

If the data that is being deleted by the **AUDIT VOLUME FIX=YES** command is on a copy-storage-pool volume, the data is deleted from the copy pool volume. The next time that the primary storage pool is backed up, a new copy is added to the copy storage pool.

Resolving error messages CTGTRV009E and CTGTRV011E

The CTGTRV009E and CTGTRV011E error messages are issued from the Tivoli Integrated Portal reporting and monitoring feature.

The messages are generated from the data that is returned for IBM Tivoli Storage Manager server or client reports. The cause for the errors typically comes from the need for increased memory heap size for the Tivoli Integrated Portal.

To determine if memory heap size is causing the error messages, perform the following steps:

1. Log on to Tivoli Integrated Portal and expand **Reporting** in the navigation pane on the left.
2. Click **Common Reporting**.
3. On the Reports panel, click the Navigation tab and expand **Report Sets > Tivoli Products > Tivoli Common Reporting > Tivoli Storage Manager > Client Reports > Server Reports**.
4. Click **Client Reports** and select the “Client Schedule Status” report.
5. When prompted for the report period, select **Last 30 days** and click **Run**. If memory heap size is the reason for the error messages, the following messages are displayed:

CTGTRV009E Processing has ended because of an unexpected error.

CTGTRV011E See the Tivoli Common Reporting log files for more information.

Increasing the maximum heap size

The CTGTRV009E and CTGTRV011E errors are often because of a need for an increased memory heap size for Tivoli Integrated Portal.

To increase the maximum memory heap size, complete the following steps:

1. Open a command-line interface and change directory to the Administration Center default installation directory.

AIX **HP-UX** **Linux** **Solaris** The default installation directory is `TSM_home/opt/IBM/tivoli/tsmac/bin`. Where *TSM_home* represents the Tivoli Storage Manager server installation directory.

Windows The default installation directory is `C:\IBM\tivoli\tsmac\bin`.

2. Run the following utility:

AIX **HP-UX** **Linux** **Solaris** `./supportUtil.sh`.

Windows `supportUtil.bat`.

3. Enter the Tivoli Integrated Portal administrator user ID and password.
4. From main menu of the Administration Center Support utility, complete the following steps to change the memory heap size:
 - a. Select **Manage the maximum memory size the Administration Center can use**.
 - b. Select **Update the maximum memory size the Administration Center can use**.
 - c. Enter **1024** for the value.
 - d. Select the option **99**, and return to the main menu.
 - e. Select **Exit**.
5. Stop and restart Tivoli Integrated Portal to update the memory heap size.

Files are not expired after reducing versions

Errors are sometimes generated after you update the server policies to reduce the number of versions of a file to retain.

Issue the **QUERY COPYGROUP** *domainName policySetName copyGroupName* **F=D** command. If either the “Versions Data Exists” or “Versions Data Deleted” parameters were changed for a **TYPE=BACKUP** copy group, it might affect expiration.

If the “Versions Data Exists” or “Versions Data Deleted” values for a **TYPE=BACKUP** copy group were reduced, the server expiration process might not immediately recognize this fact and expires these files. The server only applies the “Versions Data Exists” and “Versions Data Deleted” values to files at the time they are backed up to the server. When a file is backed up, the server will count the number of versions of that file and if that exceeds the number of versions that should be kept, the server will mark the oldest versions that exceed this value to be expired.

Process symptoms indicate migration errors

You might be faced with process symptoms that point to migration as the cause for errors.

Migration does not run for sequential media storage pool

If migration does not run for sequential media storage pools, issue the **QUERY STGPOOL** *stgpoolName* **F=D** command.

Migration from sequential media storage pools calculates the “Pct. Util” as the number of volumes in use for the storage pool, relative to the total number of volumes that can be used for that storage pool. Similarly, it calculates the “Pct. Migr” as the number of volumes with data that can be migrated, in use for the storage pool, relative to the total number of volumes that can be used for that storage pool. Because it might be considering unused scratch volumes in this calculation, there might not appear to be sufficient data that can be migrated in the storage pool to require migration processing.

Migration uses only one process

Issue the **QUERY STGPOOL** *stgpoolName* **F=D** and **QUERY OCCUPANCY * * STGPOOL=** *stgpoolName* command.

The following are the reasons why only one migration process is running:

- The Migration Processes setting for the storage pool is set to one or is not defined (blank). If true, issue the **UPDATE STGPOOL** *stgpoolName* **MIGPROCESS=*n*** command, where *n* is the number of processes to use for migrating from this pool. Note that this value must be less than or equal to the number of drives (mount limit) for the NEXT storage pool where migration is storing data.
- If the **QUERY OCCUPANCY** command only reports a single client node and file space in this storage pool, migration can only run a single process even if the Migration Processes setting for the storage pool is greater than one. Migration processing is partitioning data, based on client node and file space. For migration to run with multiple processes, data for more than one client node needs to be available in that storage pool.

Resolving storage pool issues

Storage pools are integral to a successful server operation. The IBM Tivoli Storage Manager database contains information in storage pools about registered client nodes, policies, schedules, and the client data.

This information must be available and valid in order for Tivoli Storage Manager to function correctly. Storage pool errors can be related to the following issues:

- Failed transactions
- A storage pool experiencing a high volume usage after increasing the MAXSCRATCH value
- A storage pool having “Collocate?=Yes” but volumes still containing data for many nodes
- Unable to store data in an active data pool by using the simultaneous-write function or by issuing the **COPY ACTIVEDATA** command

“ANR0522W Transaction failed...” message received

The ANR0522W message is displayed when the server is unable to allocate space in the storage pool that is identified to store data for the specified client.

There are a number of possible causes for running out of space in a storage pool. Perform the following procedures to resolve the space allocation error:

1. Issue **QUERY VOLUME volname F=D** for the volumes in the referenced storage pool. For any volumes reported with access other than Read/Write, check that volume. A volume might be marked Read/Only or Unavailable because of a device error. If the device error is resolved, issue the **UPDATE VOLUME volname ACCESS=READWRITE** command to let the server select and try to write data to that volume.
2. Issue **QUERY VOLUME volname** for the volumes in the referenced storage pool. Volumes that report “pending” for the volume status are volumes that are empty but waiting to be reused again by the server. The wait time is controlled by the REUSEDELAY setting for the storage pool and displayed as “Delay Period for Volume Reuse” on the **QUERY STGPOOL** command. Evaluate the REUSEDELAY setting for this storage pool and, if appropriate (based upon your data management criteria), lower this value by issuing the **UPDATE STGPOOL stgpoolname REUSEDELAY=nn** command, where *stgpoolname* is the name of the storage pool and *nn* is the new reuse delay setting. The key to getting the data collocated is to have sufficient space in the target storage pool for the collocation processing to select an appropriate volume. Having sufficient space in the target storage pool is significantly influenced by the number of scratch volumes in a storage pool.
3. Issue the **QUERY STGPOOL F=D** command to verify whether the ACCESS is Read/Write.

Storage pool experiences high volume usage after increasing MAXSCRATCH value

For collocated sequential storage pools, increasing the **MAXSCRATCH** value might cause the server to use more volumes.

The server uses more storage pool volumes in this case because of the collocation processing. Collocation groups user data for a client node onto the same tape. During a client backup or archive operation, if no tapes currently have data for this client node, the server selects a scratch volume to store the data. Then, for other client nodes storing data, the server again selects a scratch volume. The reason that scratch volumes are not selected prior to changing the **MAXSCRATCH** setting is that if there is no scratch volume available and no preferred volume already assigned for this client node, the volume selection processing on the server ignores the collocation request and stores the data on an available volume.

Storage pool is set to use collocation, but volumes contain data that is not collocated

When a storage pool is collocation enabled (the **COLLOCATION** parameter is set to **GROUP**, **NODE**, or **FILESPEC**), many volumes might contain data that is not collocated.

There are two possibilities for this situation:

- The data was stored on volumes in this storage pool before enabling the storage pool for collocation.
- The storage pool ran out of scratch tapes and stored data on the best possible volume, even though the request to collocate was ignored.

If data for multiple nodes ends up on the same volume for a storage pool that is collocation enabled, use one of the following actions:

- Issue the **MOVE DATA** command for the volume or volumes affected. The process reads the data from the specified volume and moves it to a different volume in the same storage pool if:
 - Scratch volumes are available, or
 - Volumes with sufficient space are assigned to this client node for collocating their data
- Allow migration to move all the data from that storage pool by setting the **HIGHMIG** and **LOWMIG** thresholds. By allowing the migration of all data to the **NEXT** storage pool, the collocation requirements are processed if the following are true:
 - The **NEXT** storage pool is collocation enabled
 - The **NEXT** storage pool has sufficient scratch volumes
 - The **NEXT** storage pool is assigned volumes to satisfy the collocation requirements
- Issue the **MOVE NODEDATA** command for the client nodes whose data is in that storage pool. If scratch volumes are available or volumes with sufficient space are assigned to this client node for collocating their data, the following events happen:
 - The **MOVE NODEDATA** process reads the data from the volumes that this node has data on
 - The **MOVE NODEDATA** process moves the data to a different volume or volumes in the same storage pool

The key to getting the data collocated is to have sufficient space in the target storage pool for the collocation processing to select an appropriate volume. There must be enough empty volumes available in the storage pool to allow collocation to select a new volume. Ensure that enough empty volumes are available instead of using a volume that already has data on it from a different node. Empty volumes can be scratch volumes if the storage pool is defined with enough scratch volumes, or define the empty volumes by issuing the **DEFINE VOLUME** command.

Resolving storage problems for active data pools

You might experience difficulty in storing data in an active data pool by using the simultaneous-write function or by issuing the **COPY ACTIVE** command.

Before data can be stored in an active data pool, you must establish a policy to allow the data into the pool. The node that owns the data must be assigned to a domain whose active data pool is listed in the domain **ACTIVEDESTINATION** field. Issue the following command to determine if the node is assigned to a domain that authorizes storing into the active data pool:

```
QUERY NODE node_name F=D
```

The Policy Domain Name field lists the domain to which the node is assigned. Issue the following command to determine if the active data pool is listed in the domain **ACTIVEDESTINATION** field:

```
QUERY DOMAIN domain_name F=D
```

If the active data pool is not listed, issue the following command to add the active data pool to the list:

```
UPDATE DOMAIN domain_name ACTIVEDESTINATION=active-data_pool_name
```

Tip: After you issue the **UPDATE DOMAIN *domain_name* ACTIVEDESTINATION=*active-data_pool_name*** command, all nodes assigned to the domain are authorized to store data in the active data pool. If having the nodes assigned to the domain authorized to store data is not acceptable, you must create a new domain for those nodes whose data you want stored in the active data pool and assign those nodes to the newly created domain. See the *IBM Tivoli Storage Manager Administrator's Guide* to learn how to establish a new policy domain.

Chapter 4. Resolving communication problems

The need for connectivity in IBM Tivoli Storage Manager means that any communication error can render your application useless. Communication errors might be attributed to TCP/IP configuration, client and server connections, and other causes.

Resolving errors created when connecting to the server

Problems that are generated while you are connecting to the server might be related to your communication options.

To correct the error, perform any or all of the following options:

- Review the changes in the client communication options in the client option file (if they exist) and try to revert back to the previous values. Retry the connection.
- If the server communication settings were changed, either update the client communication options to reflect the changed server values or revert the server back to its original values.
- If any network settings were changed, such as the TCP/IP address for the client or server (or a firewall), work with the network administrator to update the client, server, or both for these network changes.

Resolving failed connections by clients or administrators

The two main cases for connection failures are general failure, where no connections at all are allowed, or an isolated failure where some connections are allowed but others fail.

If no connections at all are possible, it might be necessary to run the server in the foreground so that a server console is available, and additional diagnostic steps can be taken. Check the settings to verify the proper configuration for communicating with the server:

- Ensure that the server is able to bind to a port when it is started. If it is unable to bind to a port, then it is likely that some other application is using that port. The server can not bind (use) a given TCP/IP port if another application is already bound to that port. If the server is configured for TCP/IP communications and successfully binds to a port on startup for client sessions, the following message is issued:

```
ANR8200I TCP/IP driver ready for connection with clients on port 1500.
```

If a given communication method is configured in the server options file, but a successful bind message is not issued during server startup, then there is a problem initializing for that communication method.

- Verify that the code **TCPPORT** setting in the server options file is correct. If the code setting is inadvertently changed, the clients fail to connect. That is because the clients are trying to connect to a different TCP/IP port than the one the server is listening on.
- If multiple servers are using the same TCP/IP address, ensure that the **TCPPORT** and **TCPADMINPORT** for each server are unique. For example, there are two servers at the same TCP/IP address. The first server has a **TCPPORT** of 1500 and a **TCPADMINPORT** of 1500. The second server has a **TCPPORT** of 1501 and a

TCPADMINPORT of 1500. The first server to grab port 1500 locks out the other server from port 1500 and clients can no longer access the first server. Administrative clients always connect to the second server. A better choice of ports for each server would be 1500 and 1501 for **TCPSPORT**; 1510 and 1511 for **TCPADMINPORT**.

- Check that the server is enabled for sessions. Issue the **QUERY STATUS** command and verify that “Availability: Enabled” is set. If the result states “Availability: Disabled,” issue the **ENABLE SESSIONS** command.
- If specific clients are unable to connect to the server, check the communication settings for those clients. For TCP/IP, check the **TCPSERVERADDRESS** and **TCPSERVERPORT** options in the client options file.
- If only a specific node is rejected by the server, verify that the node is not locked on the server. Issue the **QUERY NODE** *nodeName* command, where *nodeName* is the name of the node to check. If the result states “Locked?: Yes,” then evaluate why this node is locked. Nodes can only be locked by using the **LOCK NODE** administrative command. If it is appropriate to unlock this node, issue the **UNLOCK NODE** *nodeName* command, where *nodeName* is the name of the node to unlock.
- If the computer on which the server is running is having memory or resource allocation problems, it might not be possible to start new connections to the server. The memory or resource allocation problem might be cleared up temporarily if you either halt and restart the server, or if you halt and restart the computer itself. This action is a temporary solution, and diagnosis should be continued for either the operating system or the IBM Tivoli Storage Manager server because the memory and resource allocation problem might indicate an error in either.

Resolving Secure Sockets Layer errors

Secure Sockets Layer (SSL) errors can be attributed to an incorrect environment setup, a bad server certificate, connection problems, out-of-sync conditions, or other causes.

The common SSL client-to-server and server-to-server problems are shown here:

Missing key database file (**cert.kdb**)

The server creates the **cert.kdb** key database file if it does not exist. One or both of the **SSLTCPPORT** and **SSLTCPADMINPORT** options must be in the server options file when the IBM Tivoli Storage Manager server is started. The server generates a changeable password and also generates a self-signed certificate that can be extracted for clients and IBM Business Partner’s servers to use. If the **cert.kdb** file exists and the server did not create it, an out-of-sync condition occurs, preventing the server from setting up SSL communications.

Not connecting to the server after using a vendor-acquired certificate authority (CA) certificate

If you are using a vendor-acquired certificate and it was not added to the server, specify the root certificate as trusted in the server key database. To add the root certificate to the database, issue this command:

```
gsk8capicmd -cert -add -db cert.kdb -pw password  
-label name -file .der_file -format ascii
```

The CA root certificate was not added to the client

Add the root certificate as trusted into the client key database:

```
gsk8capicmd -cert -add -db dsmcert.kdb -pw password
-label my CA -file ca.arm -format ascii
```

Unable to run gsk8capicmd.exe (IBM Global Security Kit [GSKit])

In most cases, this Windows error is generated by an incorrect environment setup. See the *IBM Tivoli Storage Manager Administrator's Guide* and set up the PATH variable as directed, before running the gsk8capicmd utility.

ANS1595E Bad server certificate

This error is reported when the server certificate is not known to the client or server. The "bad server certificate" error can occur under these conditions:

- The certificate was never imported
- The cert256.arm certificate file was corrupted before being imported
- The command to import the certificate was entered incorrectly
- The DSM_DIR variable points to the wrong directory, which contains an incorrect client key database (dsmcert.kdb)
- The server is set up for Transport Layer Security (TLS) 1.2 but the client is not at a sufficient level (6.3 is required).
- The server is set up for TLS 1.2 but the client imported the cert.arm file instead of the cert256.arm file.
- The server is set up for TLS 1.2 but the client imported the cert256.arm file instead of the cert.arm file.

Repeat all the steps necessary for importing the server certificate and check the DSM_DIR variable. See the dsmererror.log file for more information about the failure. The client error log might also contain information about specific IBM GSKit failure.

ANS1592E Failed to initialize SSL protocol

This error occurs on the client and indicates that the SSL connection was not established. See the client error log for more information about the failure. The server does not accept SSL sessions on the port to which the client or server is trying to connect. Determine whether the client or server points to the correct server port (TCPPort), which can be a port number that is different than the default 1500.

ANR8583E and GSKit return code 406

This error might indicate that a non-SSL-enabled client is trying to contact an SSL port. When a client contacts a Tivoli Storage Manager server at a port defined by SSLTCPPORT or SSLTCPADMINPORT, the server establishes a session and initiates an SSL "handshake." If the client is not SSL-enabled, it cannot complete the SSL handshake process. The session then seems to stop, but times out through the server IDLEWAIT option or end when the server administrator issues the **CANCEL SESSION** command to manually cancel it. The example illustrates a session in this state, from the server:

```
TSM:SERVER1>query session
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY SESSION
```

Sess Number	Comm. Method	Sess State	Wait Time	Bytes Sent	Bytes Recvd	Sess Type	Platform	Client Name
1	SSL	IdleW	17 S	0	0	Node		

Important: Because the computing environment might cause a valid handshake process to take some time, do not assume that the result always indicates a non-SSL client.

ANR8583E and GSKit return code 420, and ANR8581E with GSKit return code 406 occur for the same Tivoli Storage Manager client session

When server messages ANR8583E and ANR8581E occur for the same client session, it is likely that the client has generated an ANS1595E message. Message ANS1595E is typically issued while Tivoli Storage Manager attempts to establish a session with the server. If true, follow the guidance in the Tivoli Storage Manager message manual for ANS1595E to eliminate these errors.

Cross-defining servers without SSL=YES causes a server hang

If you plan to use SSL communication, the SSL infrastructure must be in place on the source and target replication servers. Required SSL certificates must be in the key database file that belongs to each server. The SSL function is active if the server options file contains the SSLTCPPORT or SSLTCPADMINPORT option or if a server is defined with **SSL=YES** at startup. When using server-to-server communication with SSL on the Administration Center, you must select the **SSL=YES** check box. No error message is issued if you do not select **SSL=YES**.

An entry occurs when a vendor-acquired certificate in use was not added to the server, or the CA certificate was not added to the client. When an SSL session is started, the session startup message includes the serial number from the server certificate. Therefore, the certificate being used can be uniquely identified.

Related reference:

Appendix C, "IBM Global Security Kit return codes," on page 259

Recovering the key database file password

If you forgot the current key database file password, IBM Tivoli Storage Manager can help you to recover it.

You must have system privileges to administer the key database file password recovery.

Complete the following steps to recover and update the key database file password:

1. Issue the **QUERY SSLKEYRINGPW** command to display the current key database password.
2. Issue the following command to use the server record of the key database password to update the password:

```
SET SSLKEYRINGPW password UPDATE=Y
```

where *password* is the password retrieved by the **QUERY SSLKEYRINGPW** command.

Tip: If the cert.kdb file does not exist, you can create a new file by restarting the server. The Tivoli Storage Manager server creates a database file with the old password and generates a new self-signed certificate at startup. If you use self-signed certificates, you must extract the certificate and install it on a client system. If you use a vendor-acquired certificate, you must add it back in the server key database file and restart the server.

Chapter 5. Resolving Administration Center problems

The Administration Center is a web-based interface that can be used to centrally configure and manage IBM Tivoli Storage Manager servers.

You can resolve Administration Center errors through several methods, such as establishing a connection to a Tivoli Storage Manager server, reinstalling the Tivoli Integrated Portal, or by reviewing the log files that are generated during an error.

Re-establishing a connection between the Administration Center and a Tivoli Storage Manager server

If you have a problem establishing a connection to an IBM Tivoli Storage Manager server, use the Administration Center to isolate or resolve the problem.

Perform the following steps to establish a connection between the Administration Center and a Tivoli Storage Manager server:

1. Determine if the computer is accessible over the network. Issue a **PING** command in a command prompt or shell, or try connecting using Telnet or FTP. If the computer is not accessible over the network, the computer is down or there is a network communication error. Use the same TCP/IP address entered in the Administration Center.
2. Determine if the Tivoli Storage Manager server is running. If the server is not running, the Administration Center will not be able to establish a connection to it.
3. Try connecting to the Tivoli Storage Manager server by using the administrative client. Use the same connection settings as in the Administration Center (TCP/IP address, TCP/IP port, administrator ID, administrator password). This method is the quickest way to determine if the problem is on the Tivoli Storage Manager server computer or on the computer running the Administration Center. If both the Tivoli Storage Manager server and the Administration Center are running on the same computer, skip this step and go to the directory where the administrative client is installed. Issue the **DSMADMC** command with parameters matching those in the connection settings for the Administration Center.

```
dsmadmc -id=<admin id> -pass=<admin password> -tcpport=<TCP/IP port>  
-tcps=<TCP/IP address>
```

4. Determine what port the Tivoli Storage Manager server is using. You can issue the **QUERY OPT TCPPOINT** command on the server to determine the server port. Make sure the port used in the Administration Center is the same port on which the server is running. Go to the directory where the administrative client is installed. Issue **DSMADMC** with parameters matching those in the connection settings for the Administration Center. This command can also be issued from the server console.

```
dsmadmc -id=<admin id> -pass=<admin password> -tcpport=<TCP/IP port>
```

```
-tcps=<TCP/IP address>
```

```
tsm:SERVER1>QUERY OPT TCPPOINT  
Server Option      Option Setting
```

```
-----
TCPPort          1500
-----
```

```
tsm: SERVER1>
```

5. Determine if the computer you are attempting to access is behind a firewall. If the computer running the Tivoli Storage Manager server is behind a firewall, then the computer running the Administration Center must be authenticated to the firewall. Try using the computer running the Administration Center to connect to the computer running the Tivoli Storage Manager server. Use ping, Telnet, or FTP. If these methods do not work, try connecting to another computer under the same firewall as the server computer.
6. Determine if the TCP/IP address for the Tivoli Storage Manager server is entered correctly in the Administration Center. In the navigation tree, click **Manage Servers** and select **Modify Server Connection** from the drop-down list. This action lets you view the current connection settings for the Tivoli Storage Manager server connection. Verify that the server connection address is correct.
 - a. Issue the **PING** command in a command prompt or shell. Verify that the computer, that is indicated by the address, is accessible.
 - b. Issue the **QUERY OPT TCPPORT** command in a server console to determine on which port the Tivoli Storage Manager server is running. To determine which port the Tivoli Storage Manager server is using, go to the directory where the administrative client is installed. Issue the **DSMADMC** command with parameters matching those in the connection settings for the Administration Center. This command can also be issued from the server console. Verify that the server connection port is correct.

```
dsmadmc -id=<admin id> -pass=<admin password> -tcpport=<TCP/IP port>
-tcps=<TCP/IP address>
```

```
tsm: SERVER1>QUERY OPT TCPPORT
```

```
Server Option      Option Setting
-----
TCPPort           1500
```

```
tsm: SERVER1>
```

7. Determine if the administrator ID for the Tivoli Storage Manager server is entered correctly in the Administration Center. In the navigation tree, click **Manage Servers** and select **Modify Server Connection** from the drop-down list. This action lets you view the current connection settings for the Tivoli Storage Manager server connection. Verify that the administrative ID is correct.
8. Determine if the administrator password for the Tivoli Storage Manager server is entered correctly in the Administration Center. In the navigation tree, click **Manage Servers** and select **Modify Server Connection** from the drop-down list. This action lets you view the current connection settings for the Tivoli Storage Manager server connection. Enter the correct administrative password to connect to the Tivoli Storage Manager server. Retry the operation to see if changing the password resolved the problem.
9. Determine if any messages are issued on the server when you try to connect using the Administration Center. The activity log contains useful information if you happen to experience any trouble connecting to the server. You can see sessions starting and stopping when the Administration Center successfully communicates with the server. If the network is down or the TCP/IP address or port is incorrect, no information shows up in the activity log when you make connection attempts. To check the server activity log, go to the directory

where the administrative client is installed. Then issue the **DSMADMC** command with parameters matching those in the connection settings for the Administration Center. This command can also be issued from the server console.

```
dsmadmc -id=<admin id> -pass=<admin password> -tcpport=<TCP/IP port>  
-tcps=<TCP/IP address>
```

```
tsm: SERVER1>QUERY ACTLOG
```

10. Check the database file on the server for corruption. The database file is located in the Tivoli Integrated Portal installation directory under *tip_installation_location\products\tsm\tsmserver.xml*. If this file is corrupted, problems might occur when you try to connect to the servers. You will, however, typically see other error messages indicating that this file is corrupt. Open the file in an editor (preferably an XML editor), and verify that the end tag has a closing end tag.

HP-UX

Administration Center not supported

The Administration Center is a web-based interface for centrally configuring and managing Tivoli Storage Manager servers. The Administration Center provides wizards to help guide you through common configuration tasks. Properties notebooks allow you to modify settings and perform advanced management tasks.

The Administration Center cannot be installed on HP-UX, but it can be used to manage HP-UX servers. For Administration Center system requirements, see the following website: <http://www.ibm.com/support/docview.wss?uid=swg21515628>.

Resolving Tivoli Integrated Portal user authority problems

Tivoli Integrated Portal users can access only the pages for which they are authorized.

You can use the pages on the **Users and Groups** tab of the Tivoli Integrated Portal to create, delete, and view information for users and groups.

User administration for Tivoli Storage Manager Administration Center is based on the user and group management features of the Tivoli Integrated Portal. Individual users or groups can be granted role-based authorization to access Administration Center resources.

Roles and groups

User permissions for the Administration Center are managed with roles and groups.

A role represents a set of authorizations for a particular resource. When a user ID is assigned a role for a resource, that user ID is then authorized to perform specific actions on that resource.

You can assign a role to an individual user or to a group. If assigning a role to a group, the role applies to all members of the group. A user ID can have multiple roles, either assigned directly to the user ID or assigned to groups to which that user ID is a member. If a user ID has multiple roles, the authorizations associated with all of the roles apply (a role can only grant authorization, not deny it).

The Administration Center defines the tsmAdministrator role at the time of installation. The tsmAdministrator role has access to all Administration Center components.

When you install the Administration Center, a new user group is created for the Administration Center called TSM_AdminCenter. The group is given the tsmAdministrator role. Any new Tivoli Integrated Portal user IDs are created in the group. These user IDs receive access to the Administration Center.

For more information about users, roles, and groups, see the Tivoli Integrated Portal online help (access this information by clicking **Help** at the top of any page).

HP-UX

Administration Center not supported

The Administration Center is a web-based interface for centrally configuring and managing Tivoli Storage Manager servers. The Administration Center provides wizards to help guide you through common configuration tasks. You can modify settings and perform advanced management tasks in the properties notebooks.

The Administration Center cannot be installed on HP-UX, but it can be used to manage HP-UX servers. For Administration Center system requirements, see the following website: <http://www.ibm.com/support/docview.wss?uid=swg21515628>.

Adding a user ID with access to the Administration Center

AIX

Linux

Solaris

Windows

If you are having difficulty with user authority, you might want to create a user with the required permissions.

Perform the following steps to create a Tivoli Integrated Portal user ID with access to the IBM Tivoli Storage Manager Administration Center:

1. In the navigation tree, select **Users and Groups**.
2. Click **Manage Users**.
3. Click **Create**.
4. Click **Group Membership**.
5. Select **Group Name**, then click **Search**.
6. Add **TSM_AdminCenter** to the Current Groups list.
7. Click **Close**.
8. Complete the form and click **Create**.

Resolving a user ID access problem with the Administration Center

If you are having difficulty with user authority, you might want to look into the user ID permissions.

Ensure that the group to which the user ID belongs has the tsmAdministrator role. Typically, the group is TSM_AdminCenter, but that is not a requirement. Complete the following steps to determine the group to which the user ID belongs and to verify the role that is assigned to the group:

1. In the navigation tree, expand **Users and Groups**.
2. Click **Manage Users**.
3. Click **Search**.
4. Click the user ID in the table.
5. Click the **Groups** tab.
6. Perform the following steps to determine the group to which the user ID belongs:
 - a. In the navigation tree, expand **Users and Groups**.
 - b. Click **Group Roles**.
 - c. Click **Search**.
 - d. Click the Group Name in the table.
 - e. Verify that the group has the tsmAdministrator role.

Resolving a stopped Tivoli Integrated Portal server

The Tivoli Integrated Portal is built on top of a WebSphere Application Server. If the WebSphere server were to stop running, the Tivoli Integrated Portal would no longer be accessible from a browser.

Before restarting the Tivoli Integrated Portal server, gather information about the stoppage so that you can report the problem.

WebSphere includes several troubleshooting tools that are designed to help you isolate the source of problems. Many of these tools are designed to generate information that IBM Support can use.

Related reference:

“Resolving server access problems” on page 107

Running the collector tool to obtain problem-analysis information

The collector tool gathers information about your WebSphere server installation. The tool then packages the information in a Java archive (JAR) file which you can send to the WebSphere Customer Center. The center helps you to analyze your problem.

Information in the JAR file includes log files, property files, configuration files, operating system and Java data, and the presence and level of each software prerequisite.

There are two ways to run the collector tool. You can run the collector tool to collect summary data or to traverse the system to gather relevant files and command results. The collector tool produces a JAR file with information needed

to determine and solve a problem. The collector summary option produces a lightweight collection of the version and other information that is useful when first reporting the problem to the WebSphere Customer Center.

The collector program runs to completion as it creates the JAR file, despite any errors that it might find. Errors might include missing files or commands. The collector tool collects as much data in the JAR file as possible.

There are two phases of using the collector tool. The first phase involves running the collector tool on your WebSphere server and producing a JAR file. The Support team performs the second phase, which involves analyzing the JAR file.

Perform the following steps to run the collector tool:

1. Log on to the system as root (or Administrator on Windows).
2. Verify that Java 1.5.0 or higher is available in the path.
3. With multiple Java Development Kits (JDKs) on the system, verify that the JDK that the WebSphere server uses is the one in the path for the collector program. The collector program requires Java to run. It also collects data about the JDK in which it runs.
4. Verify that all necessary information is in the path being used by the collector program. Ensure that you are not running the program from within the Tivoli Integrated Portal installation root directory.
 - If this system is an AIX, Linux, or Solaris platform, verify that the path contains the following entries:
 - /bin
 - /sbin
 - /usr/bin
 - /usr/sbin

HP-UX

Administration Center not supported

The Administration Center is a web-based interface for centrally configuring and managing IBM Tivoli Storage Manager servers. The Administration Center provides wizards to help guide you through common configuration tasks. Properties notebooks allow you to modify settings and perform advanced management tasks.

The Administration Center cannot be installed on HP-UX, but it can be used to manage HP-UX servers. For Administration Center system requirements, see the following website: <http://www.ibm.com/support/docview.wss?uid=swg21515628>.

- If this system is a Windows platform, include `regedit` in the path. The collector tool is located in the `<TIP_HOME>\bin` directory.
5. Create a work directory where you can start the collector program and go to the work directory.

The collector program writes its output JAR file to the current directory and also creates and deletes a number of temporary files in the current directory. By creating a work directory to run the collector program, the collector program avoids naming collisions and makes cleanup easier. You cannot run the collector tool in a directory under the Tivoli Integrated Portal installation directory.

6. Run the collector program by issuing the **COLLECTOR** command from the command line.

For example: `c:\work\dir>collector`

One copy of the node data and data from each server in the node is collected when you issue the **COLLECTOR** command. The collector program stores the data in a single JAR output file. To gather data from a specific server in the node, issue the **COLLECTOR** *servername* command, where *servername* is the name of the problem server. The name of the Tivoli Integrated Portal server is `tsmServer`.

The collector program creates a log file, `Collector.log`, and an output JAR file in the current directory.

The name of the JAR file is composed of the host name, cell name, node name, and profile name: `host_name-cell_name- node_name-profile_name.jar`.

For example, if you run the collector tool on server `bohml`, the file name is `bohml-TIPCell-TIPNode-TIPProfile-WASenv.jar`.

The `Collector.log` file is one of the files collected in the `host_name--cell_name- node_name-profile_name.jar` file.

Tip: Contact Support for assistance in deciphering the collector tool output.

Diagnosing log-entry problems by using the log analyzer tool (showlog)

The log analyzer takes one or more service or activity logs, merges all of the data, and displays the entries.

Based on its symptom database, the tool analyzes and interprets the event or error conditions in the log entries to help you diagnose problems. The log analyzer has a special feature, which enables it to download the latest symptom database from the IBM website.

About the Service or Activity log

The WebSphere server creates the service or activity log file from the activity of the various WebSphere server components. The log analyzer is used to view the service or activity log file and can merge service or activity log files into one log file. The service or activity log file (`activity.log`) is a binary file in the following directory:

`tip_install_root\profiles\TIPProfile\logs`, where *tip_install_root* is the root directory for your installation.

HP-UX

Administration Center not supported

The Administration Center is a web-based interface for centrally configuring and managing IBM Tivoli Storage Manager servers. The Administration Center provides wizards to help guide you through common configuration tasks. Properties notebooks allow you to modify settings and perform advanced management tasks.

The Administration Center cannot be installed on HP-UX, but it can be used to manage HP-UX servers. For Administration Center system requirements, see the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21515628>.

Using the log analyzer

You cannot view the service or activity log with a text editor. Use the log analyzer tool to view the file.

The Websphere server on which the log analyzer tool runs does not include the Java administrative console, so there is no graphical interface available for viewing a service or activity log file. The alternate viewing tool, `showlog`, must be used to view the service or activity log file:

1. Change to the directory `tip_install_root\bin` where `tip_install_root` is the root directory for your installation.
2. Run the showlog tool with no parameters to display usage instructions:
 - **Windows** On Windows systems, run `showlog.bat`
 - For all other supported systems, run `showlog.sh`

To direct the service or activity log (`activity.log`) contents to stdout, issue the **SHOWLOG ACTIVITY.LOG** command.

To save the service or activity log to a text file for viewing with a text editor, issue the **SHOWLOG ACTIVITY.LOG SOME_TEXT_FILE_NAME** command.

Tip: Contact IBM Support for assistance in deciphering the output of the collector tool.

Resolving excessive memory consumption problems with the Tivoli Integrated Portal server

The Tivoli Integrated Portal is built on top of a WebSphere application server. If the Websphere server is using a large amount of memory, you can take action to resolve this issue.

To determine if a larger amount of memory is in play, use the operating system tools that provide such information. For example, on Windows, the Task Manager shows memory usage.

Logging out at the end of each day should help minimize the memory use. Logging out instead of only closing the browser is recommended. The default value of Session Timeout is 30 minutes. Reduce this value to a lower number to help reduce memory consumption requirements. You can change the `SESSIONTIMEOUT` value with the Administration Center support utility.

Related reference:

“Administration Center Support utility” on page 111

Configuring the IP address to align with the Administration Center

The Administration Center assumes that the host system uses a static IP address instead of a dynamically-assigned IP address. A static IP address is necessary because the Administration Center server must be listed on the domain name servers, which map the host name to the physical address of the system.

On a system that is not connected to the network, you must configure your system so that the IP loopback port is mapped to the fully-qualified host name. To enable that mapping, perform the following steps:

1. Go to the workstation where the Administration Center will be installed.

HP-UX

Administration Center not supported

The Administration Center is a web-based interface for centrally configuring and managing IBM Tivoli Storage Manager servers. The Administration Center provides wizards to help guide you through common configuration tasks. Properties notebooks allow you to modify settings and perform advanced management tasks.

The Administration Center cannot be installed on HP-UX, but it can be used to manage HP-UX servers. For Administration Center system requirements, see the following website: <http://www.ibm.com/support/docview.wss?uid=swg21515628>.

2. Locate the TCP/IP hosts file on your system.
For Windows systems, look in `WINNT\system32\drivers\etc`.
For all other supported systems, look in `path/etc`.
3. Use a text editor to open the file named "hosts."
4. At the end of the hosts file, add lines similar to the following lines:
`127.0.0.1 localhost`
`127.0.0.1 your.server.name`

where *your.server.name* is the fully-qualified host name for the Administration Center system.

5. Save the hosts file.

Resolving server access problems

You might encounter a problem when you try to access the Tivoli Integrated Portal and the Tivoli Storage Manager Administration Center from a web browser.

The following examples are of the messages that you could receive:

- The page cannot be displayed.
- The page you are looking for is currently unavailable.
- Error: Cannot find server or DNS Error.
- The connection was refused when attempting to contact host:port.

These types of errors can occur for a number of reasons. The following list gives you different ways in which you can correct these errors:

- Ensure that the system on which the browser is running is connected to the network and the target system on which Tivoli Integrated Portal is installed.
- Ensure that the browser settings are correct.
- If you are behind a firewall and Internet security or proxy software is active, try disabling it temporarily. If Tivoli Integrated Portal is now accessible, the firewall settings should be investigated.
- Ensure that Tivoli Integrated Portal is installed on the target system.
- Ensure that Tivoli Integrated Portal is running on the target system. Tivoli Integrated Portal might have been terminated.
- Ensure that the Tivoli Storage Manager Administration Center was deployed into Tivoli Integrated Portal. HP-UX

Administration Center not supported

The Administration Center is a web-based interface for centrally configuring and managing Tivoli Storage Manager servers. The Administration Center provides wizards to help guide you through common configuration tasks. Properties notebooks allow you to modify settings and perform advanced management tasks.

The Administration Center cannot be installed on HP-UX, but it can be used to manage HP-UX servers. For Administration Center system requirements, see the following website: <http://www.ibm.com/support/docview.wss?uid=swg21515628>.

Tivoli Integrated Portal is actually composed of two servers: the console server and the console help server. If the system on which Tivoli Integrated Portal is installed is shut down, the Tivoli Integrated Portal server that was started will stop.

If the system is configured to automatically start the server when the system restarts, you should not have to take any action.

To start the server from a command-line interface, go to the *TIP_HOME*\TIPProfile\bin subdirectory of the Tivoli Storage Manager installation directory and use the appropriate command for your operating system:

- Windows - startServer.bat server1
- All other supported systems - startServer.sh server1

tip_home is the root directory for your Tivoli Storage Manager installation.

To stop the server from a command-line interface, go to the *tip_home*\profiles\TIPProfile\bin subdirectory of the Tivoli Storage Manager installation directory and use the appropriate command for your operating system:

- Windows - stopServer.bat server1 (-user tipadmin -password tippass)
- All other supported systems - stopServer.sh server1 (-user tipadmin -password tippass)

where *tipadmin* is the administrator user ID for Administration Center and *tippass* is the password for the administrator.

To stop the server you must specify a user ID and the password for that user ID.

Resolving Administration Center health monitor problems

A health monitor presents a view of the overall status of multiple servers and their storage devices.

From the health monitor you can link to details for a server, including a summary of the results of client schedules and a summary of the availability of storage devices. Conditions are evaluated on a point system, where the points indicate the kind of problems occurring with the server.

- Scoring between 0 - 4, the status is normal
- Scoring between 5 - 9 indicate a warning status
- Scoring 10 and higher indicate a critical status

An exception occurs if a server is stopped. If the server is stopped, server status cannot be obtained and the status is reported as critical.

The point calculations are provided by the *dshealthmon.xml* and *dshealthmon_pre61.xml* report files. In addition to calculating the overall condition of the server, these files count the activity log and event numbers that are displayed on the health monitor details page.

Restriction: If you make changes to the *dshealthmon.xml* or *dshealthmon_pre61.xml* files, those changes are not supported. Save a copy of these files before you change them to ensure that you can restore the original files if an error with the health monitor occurs.

Health monitor conditions that can cause an unknown server status

The Administration Center offers the functions of most administrative commands, as well as unique functions, such as the health monitor and wizards to help you perform complex tasks.

Review the following symptoms of server status problems:

- The ADMIN_CENTER ID not defined: If the ADMIN_CENTER administrative ID was deleted from the server, the health monitor is unable to obtain status. You must resynchronize the health monitor password.
- The ADMIN_CENTER password for the server does not match the password for the health monitor. You must resynchronize the health monitor password.
- The ADMIN_CENTER ID is locked. The ID must be unlocked or the health monitor is unable to obtain status. You must resynchronize the health monitor password.
- The health monitor worker is no longer running. Tracing is the only way to determine if the worker thread is running. Follow the Administration Center tracing instruction to activate tracing for the HEALTH trace class. Search the `tsmServer/logs/trace.log` file for strings like "Time for a nap" or "Back to work." If you find messages like that in the trace file, the health monitor is running. You might need to wait to see those messages, depending on if the health monitor worker is sleeping for the given interval. When you do not see any activity for a period of time longer than the refresh interval, try restarting the Tivoli Integrated Portal, which restarts the health monitor worker.

Related tasks:

"Enabling Administration Center trace" on page 127

Resolving health monitor conditions that can cause a warning or critical storage status

The health monitor determines the server storage status by evaluating a set of rules. No single condition can cause the change in status, but rather, the health monitor looks at a variety of items to determine the overall status.

Perform the following steps to determine the overall health of the storage hardware:

1. Determine the percentage of drives that are offline in a given library.
2. Determine the percentage of paths that are offline for each drive in a given library.
3. Find the average of the numbers in the above step, which represents the score for the drive path status.
4. Determine the percentage of library paths that are offline for the given library. The score for the library is the maximum of the three scores. The number should be somewhere in the range of 0 - 10.
5. Repeat the previous steps for every library on the server.

The overall health of the storage hardware is the average of the sum of all library scores, which is the overall health of the storage hardware on a scale from zero to ten. A score of four or less represents Normal status. A score greater than four but less than eight represents Warning status. A score greater than seven represents a Critical status.

The library status becomes Critical in the following instances:

- If you take offline, all of the drives in the library
- If you take offline, all of the paths for all of the drives in a single library
- If you take offline, the only path to the library

Depending on the number of libraries on the server, the Critical status for a single library is diluted by other healthy libraries. For example, if the status of one library is Critical and two libraries are Normal, then the storage status is Warning. To view the individual status of the library, use the Storage Devices work item.

Important: If the server is stopped, the health monitor is unable to obtain status. The status for a stopped server is reported as critical.

Health monitor conditions that can cause a warning or critical database status

The health monitor determines the server database status by evaluating a set of rules. No single condition can cause the change in status, but rather, the health monitor looks at a variety of items to determine the overall status.

If the server is down, however, the health monitor is unable to obtain status. The status for a down server is reported as critical.

If the score is between 0 - 4, the status is normal. Scores between 5 - 9 cause a warning, and a score of ten and higher indicate a critical status.

The databaseScore is calculated by a summation of the following rules:

- Free space less than 10% on file systems where the database resides: 10 points
- Free space less than 20% on file systems where the database resides: 5 points
- No database backup in the last 24 hours: 3 points
- Buffer pool hit ratio is less than 98%: 3 points

Determining when to resynchronize the ADMIN_CENTER administrator ID password

When you configure the health monitor, the password for the ADMIN_CENTER administrator ID is saved in the health monitor configuration. The same password is used to contact all the servers on the list.

If the password is being updated through the Configure Health Monitor action from the Health Monitor Server table, the password is synchronized for all the servers. When the ADMIN_CENTER password gets changed from the command-line interface or the Server Properties notebook, however, the health monitor knows the new password in order to update the configuration file. The following actions occur:

1. The health monitor contacts the server and the server returns a false credential message.
2. After the third try, the ADMIN_CENTER ID is locked.

You cannot unlock the ADMIN_CENTER ID through the command-line interface or the Server Properties notebook because the health monitor does not have a valid password. This action unlocks the ID; then after three tries, the ID is locked again.

If you reconfigure the health monitor to update the password, you are unable to do so because the health monitor attempts to log on using the password that is no longer valid on that server.

Resynchronize the ADMIN_CENTER password on all of the servers to resolve the problem. The health monitor performs the following procedures on every server with a connection to the current Tivoli Integrated Portal user ID:

- Creates the ADMIN_CENTER ID if it does not exist
- Unlocks the ADMIN_CENTER ID if it is locked
- Updates the ADMIN_CENTER password to the current health monitor password for all of the servers

Administration Center Support utility

The Administration Center Support utility is a command-line driven utility designed to assist in performing basic support tasks with the Administration Center.

This utility is located in the IBM Tivoli Storage Manager installation directory under *install_root\tsmac\bin*.

Windows On Windows, start the tool by issuing `supportUtil.bat`.

AIX **Linux** **Solaris** On all other supported systems, start the tool by issuing `supportUtil.sh`.

Windows The following example demonstrates what you will see on your Windows system:

```
C:\IBM\AC\products\tsm\bin>supportUtil
Administration Center Support Utility - Main Menu
=====
1. Manage Administration Center tracing
2. Manage the maximum memory size the Administration Center can use
3. Manage the Administration Center session timeout setting
4. Collect trace files, logs and system information to send to support
5. Generate a heap dump of the Java virtual machine
6. Generate a Java core dump of the Java virtual machine
7. View the log file for this utility.

9. Exit.
```

Related reference:

“Trace classes for the Administration Center” on page 125

Responding to Administration Center task failure messages

The Administration Center provides both informational and error messages when a failure occurs. Most failure messages typically provide information about what went wrong.

The Administration Center messages often explicitly diagnose the problem and include suggestions about how to manage it. The messages can contain information about Administration Center internal errors or server errors. In some cases, errors can be resolved by retrying a failed task. Before retrying the task, close any error messages to ensure that only relevant error messages are displayed if the task fails again.

Some Administration Center failure messages include one or more error messages returned from the server. The server messages are displayed in the language that was enabled for the server, which might not match the language enabled for the web browser. In such cases, consult the Message Manual for the message number identified for the specific server platform.

To better understand the problem, check the Activity Log for the server chosen when the failure message occurred. To view the activity log, select a server and select the Server Properties table action. In the properties notebook, click the Activity Log tab. The Activity Log tab is very useful for the following cases:

- Where the failure message does not seem related to the behavior of the task
- Where the failure message does not include server messages at all
- Where only the last server message is shown when multiple tasks were performed

In these cases, multiple commands are typically issued to the server. Look carefully in the Activity Log for the activity at the time frame of the failure. Keep in mind that the server and the computer that was used to access the Administration Center might be running in a different time zone. Activity log information is displayed using the time on the server.

If errors are encountered during command routing, check the activity log of the source server first. If the activity log does not provide sufficient information to diagnose the problem, check the activity log of the target server. Make sure that the two servers are running, that they are set up correctly to communicate with each other, and that the source server has appropriate authority level for the commands performed on the target server at the time of the failure.

If you need to contact IBM Software Support to resolve an Administration Center issue, you can use the tracing features provided by the Administration Center Support Utility to obtain additional information.

Related tasks:

“Enabling Administration Center trace” on page 127

Related reference:

“Responding to Administration Center messages about unexpected results”

“Determining the source of a message” on page 114

“Administration Center Support utility” on page 111

Responding to Administration Center messages about unexpected results

There are several ways to determine the source of a problem when a task fails on the Administration Center with unexpected results.

Related reference:

“Responding to Administration Center task failure messages” on page 111

Checking the server activity log to resolve Administration Center problems

The Administration Center processes its work by issuing the IBM Tivoli Storage Manager query, select, and other commands to the server.

Check the activity log of the server that was selected for the failing task. To view the activity log, select a server and select the Server Properties table action. In the properties notebook, click the Activity Log tab. Ensure that commands related to the task appear in the activity log and are being issued to the correct server. If there is no evidence of activity related to the Administration Center task, there might be a connection problem or the incorrect server might have been selected.

Look for failures related to commands that were issued in the time frame of the failing Administration Center task. These failures could be due to the incorrect formatting of the command or false parameter values. In some cases, a parameter value comes directly from the Administration Center, but in many others it is specified by the user through a text box, radio button, or other input method. Correct the false value if possible, but if a false parameter or incorrectly constructed command cannot be changed by the user, an internal error might be the cause.

Resolving errors caused by starting or stopping a wizard or portlet

Problems can occur because of the manner in which wizards or portlets are started or restarted. Some of the wizards commit objects to the server database as you proceed through the wizard, instead of only when you finish the wizard.

When the wizard defines objects before you click **Finish**, a summary panel listing the committed objects is displayed. The **Back** button is disabled because you cannot return to the previous wizard panel to undo the committed changes. If the wizard is then canceled, the objects that were created are not deleted. You cannot restart the wizard from the point at which it was canceled.

If you cancel and then restart a wizard, you cannot duplicate the work done earlier. For example, if you use the **Add a Storage Device** wizard to add a tape device, you will encounter several such summary panels that show the objects that are created before the wizard is finished. The first summary panel is for the creation of a device class, a library, and paths from the server to the library. If you cancel this wizard following the first summary panel and restart it, you cannot create another library with the same name because it now already exists.

After an object is created with a wizard, any changes to that object must be made through a properties notebook. In the **Add a Storage Device** example, changes to the device class, library, or storage pools must be made through the device class, library, or storage pool properties notebooks. Typically, you can click an object's name to open its properties notebook. The properties notebook contains tabbed pages for modifying all of the attributes created by the wizard. Attributes that are set to the default by the wizard can be changed only in these properties notebooks.

Working with multiple portlets

In the V6.3 Administration Center, one portlet is automatically opened when you select a notebook item in the navigation tree. After you select a properties notebook, the notebook is generated in the portlet from which it was started, and not in a new portlet. With V6.3, if you want to retain a portlet, you can select **Open Page in New Tab** from the Select Action menu. You can then move to the properties notebook from the new tab.

Updates to a server made on one tabbed page do not necessarily propagate to all other open tabbed pages. In particular, the settings in a properties notebook are gathered when the properties notebook is first started. If server objects are created in a different tab while the properties notebook is active, newer objects are not automatically added to the properties notebook. Likewise, a wizard that offers a selection of objects lists those objects that existed at the start of the wizard. However, in some cases, object lists can be manually refreshed.

To ascertain whether changes made in another table, notebook, or wizard affected objects in a particular tab, it might be necessary to restart the current task. Issued commands can add objects that are not displayed in notebooks or wizards when they were started. In general, it is best to not use the command-line client or server console while using the Administration Center.

Resolving problems caused by internal errors

An internal error is an error in the Administration Center itself, which can be more difficult to diagnose.

In most cases an error message is displayed. Another symptom is when an **OK** or **Cancel** button does not close the current panel. Finally, the server activity log file might contain commands issued by the Administration Center that are improperly formatted or contain false parameter values.

These errors can only be accurately diagnosed by support personnel. In addition to the server activity log file, the Tivoli Integrated Portal trace log can often provide information relating to the error. If the problem is repeatable, the tracing level of the trace log can be increased. The information in the trace log file will likely be useful to support personnel.

Related tasks:

“Enabling Administration Center trace” on page 127

Determining the source of a message

Messages shown in the Administration Center are typically either IBM Tivoli Storage Manager server messages or messages specific to the Administration Center.

Tivoli Storage Manager server messages are displayed, in some cases, when the Administration Center performed a command on a Tivoli Storage Manager server. In some cases, both types of messages are displayed. To distinguish between Administration Center messages and Tivoli Storage Manager server messages, observe the prefix of the message number. Messages that do not show any message number are Administration Center informational messages specific to the Tivoli Integrated Portal.

Administration Center messages have the prefix ANRW. Informational and error messages are distinguished by the suffix of the message, I and E, respectively.

For example:

```
ANRW0022I The operation completed successfully.
```

```
ANRW0023E An internal error occurred during validation.
```

Informational messages, for example ANRW0022I, provide some general information or identify limitations or requirements for the task that is being performed.

Error messages, for example ANRW0023E, identify that there is a problem completing the requested task. The message can prompt the user to perform additional setup (such as server-to-server communication), identify that the requested action is false in the current work item or task, or report an internal error. Resolving an internal error, such as ANRW0023E, generally requires

assistance from IBM Software Support. You are typically asked to provide tracing information, as well as information about the sequence of actions performed before the failure occurred.

Messages issued by the Administration Center are displayed in the language selected for the web browser, or its closest default. Administration Center message numbers are not server dependent, unlike Tivoli Storage Manager server messages. Administration Center messages are generally based on the return code specified when a Tivoli Storage Manager server command was performed. In some cases the server return code that is the basis for the Administration Center message is shown, along with the server message. In rare cases, the return code and message displayed will be inconsistent. If you use tracing to capture information for a server command implementation problem, the return codes identified with the server messages can help IBM Software Support diagnose the issue.

Some Tivoli Storage Manager server messages are dependent on the server operating system. When the Administration Center displays a Tivoli Storage Manager server message, the message number can be associated with different messages for servers running on different operating systems. A single message number might have different variations for different platforms.

Tivoli Storage Manager messages use different message prefixes, typically ANR or ANE.

Typically, each Administration Center message contains a link to the help system message explanation. Click the link to display help for the message. If the message does not have a link, use the Administration Center command-line interface or a command-line interface on the server where the failure occurred. To request help for a specific Tivoli Storage Manager server message number, issue the `HELP message_number` command.

Related reference:

“Responding to Administration Center task failure messages” on page 111

Tivoli Storage Manager messages

The full catalog of IBM Tivoli Storage Manager messages can be found in the PDF or in the information center.

Visit the information center or the IBM Publication Center for a copy of the PDF.

Locating Tivoli Integrated Portal error messages

If you must find more information about a Tivoli Integrated Portal error message, that information is available on the server.

Messages issued in the Tivoli Integrated Portal are available only in the Tivoli Integrated Portal online help.

Help text for messages that are issued through the Tivoli Integrated Portal can be viewed in the Tivoli Integrated Portal online help. For example, to view information about the ANRW0557E message, use the following URL:
`https://xxx:16316/ibm/help/topic/com.tivoli.dsm.admincenter.message.help/msghelp/ANRW0557E.html?noframes=true`

Where *xxx* is the TCPIP host name of the workstation running the Tivoli Integrated Portal.

Resolving Tivoli Storage Manager server command-definition file problems

The Administration Center uses command-definition files that specify the commands supported for each server operating system.

There is a different command-definition file for each server operating system. If this file was corrupted on the Tivoli Integrated Portal system or on the server system, parts of the interface might not be able to validate creation of new objects and might be unable to correctly construct server commands. Administration Center messages that identify server command syntax errors or validation errors (for text entry or the selection of interface controls) can sometimes be caused by problems with the command-definition files.

The Administration Center acquires a command-definition file the first time it establishes a connection with a server. The next time the Administration Center contacts a server running on the same operating system, it uses the command file that was previously obtained for that operating system. If the command file is not available in the appropriate location (due to renaming or deletion), then any request following the disappearance of the file triggers another transfer of the command file for that specific operating system.

Do not update, modify, or delete the command-definition file. If problems occur that might be related to the command-definition file, verify that the appropriate command-definition file is available to the Administration Center for the server operating system. If this file does not exist, verify that the IBM Tivoli Storage Manager server can communicate with the Administration Center application programming interface (API) and that it has the platform-specific command file and `dsmcmd.xml` file available in the server directory hierarchy. Verify that the appropriate platform-specific command file is available to the Administration Center for the platform of the server of interest. If not, verify that the Tivoli Storage Manager server can communicate with the Administration Center API and it has the platform-specific command file and `dsmcmd.xml` file available and valid in the server directory hierarchy.

One scenario of a missing or corrupted command-definition file is, for example, when a failure occurs as a task is performed on an AIX 6.1.0 server. When a connection is established to that server or a command must be generated or verified, the Administration Center obtains the command-definition file. The platform-specific command file is copied from the server to the following location in the Tivoli Storage Manager installation directory:

```
install_root\tsmac\props\CmdFileCache\aix_6_1_0_0.xml
```

If you identify a corrupted command-definition file and cannot recover the original command-definition file, you can use a connection to another server running the same operating system to obtain a valid command-definition file, after renaming the original corrupted file.

If other Administration Center `.xml` or `.xsd` files in the `WEB-INF` directory are corrupted, similar issues can occur. For example, the Administration Center might not be able to construct tables or launch portlets.

Resolving backup-archive client deployment problems

Administration Center backup-archive client deployment feature problems can sometimes be attributed to missing prerequisites.

Ensure that the following items are completed:

- Install the V6.3 IBM Tivoli Storage Manager server.
- Install the V6.3 Tivoli Storage Manager Administration Center.
- Set the client PASSWORDACCESS option to *generate*.
- Verify that the client acceptor (CAD) or backup-archive client scheduler is running when you deploy the backup-archive client.

HP-UX

Administration Center not supported

The Administration Center is a web-based interface for centrally configuring and managing Tivoli Storage Manager servers. The Administration Center provides wizards to help guide you through common configuration tasks. You can modify settings and perform advanced management tasks using the properties notebooks.

In Tivoli Storage Manager Version 6.2 and later, the Administration Center cannot be installed on HP-UX servers, but it can be used to manage HP-UX servers. For Administration Center system requirements, see the following website: <http://www.ibm.com/support/docview.wss?uid=swg21515628>.

Configuring the server for automatic backup-archive client deployments

After you install IBM Tivoli Storage Manager and the Administration Center, ensure that the prerequisites are fulfilled.

HP-UX

Administration Center not supported

The Administration Center is a web-based interface for centrally configuring and managing Tivoli Storage Manager servers. The Administration Center provides wizards to help guide you through common configuration tasks. Properties notebooks allow you to modify settings and perform advanced management tasks.

In Tivoli Storage Manager Version 6.2, the Administration Center cannot be installed on HP-UX, but it can be used to manage HP-UX servers. For Administration Center system requirements, see the following website: <http://www.ibm.com/support/docview.wss?uid=swg21515628>.

You might experience problems with the client deployment that can be resolved by configuring your server. Before deploying a client, configure the server by completing the following steps:

1. Issue the **SET SERVERHLADDRESS** command to identify the server's host or IP address.
2. Issue the **QUERY STATUS** command on the deployment server to view the IP address that the client can use to connect to the server. If the client cannot connect to the server by using the server's **SERVERHLADDRESS** as displayed in the **QUERY STATUS** results, you might see this error message:

ANE4237W Query system information was completed with warning or error

3. Enable any applicable client events and assess their appropriate size and pruning duration. By enabling the events, the deployment manager can propagate the deployment messages to the server activity log. Determine an appropriate activity log size and pruning duration that can give you enough time to observe and react to the deployment results.
4. Configure the server for client automatic deployments. If you are using the Administration Center, access the Configure Client Auto Deployment wizard. Click **Tivoli Storage Manager > Manage Servers**. Select a server from the table and then select **Configure Client Auto Deployment** from the table actions. If you use the command-line interface to deploy client packages, issue the following example commands to configure the server:

```
define devclass ibm_client_deploy_import devtype=file
directory=import_directory

define stgpool stgpool_name storage_dc_name maxscratch=20

define domain ibm_client_deploy

define policyset ibm_client_deploy ibm_client_deploy

define mgmtclass ibm_client_deploy ibm_client_deploy ibm_client_deploy

define copygroup ibm_client_deploy ibm_client_deploy ibm_client_deploy
standard type=archive destination=stgpool_name retver=retention_value

assign defmgmtclass ibm_client_deploy ibm_client_deploy ibm_client_deploy

activate policyset ibm_client_deploy ibm_client_deploy

set serverhladdress=server.serveraddress.com
```

where:

- *ibm_client_deploy_import* is the temporary location from where the deployment packages are imported. This parameter is defined by the deployment manager.
 - *import_directory* is a previously defined directory that is accessible from the server.
 - *stgpool_name* is the name of a storage pool of your choosing, where the deployment packages are stored on the server. The storage pool name is based on a previously-defined device class. That device class is different from the one that is used to import data.
 - *storage_dc_name* represents the device class where the deployment packages are stored on the server.
 - *retention_value* sets the retention time for the package. You can set it to `NOLimit` or to a number of days. The default for the Administration Center is 5 years. If you specify a value other than the default value, take into account the number of days that the package is on the FTP site. The days spent on the FTP site are included in the number.
 - *server.serveraddress.com* is the server IP address or host name from which you scheduled the client automatic deployment.
5. Verify that the server objects are defined. All server objects that are created are named, unless otherwise specified, with a common prefix of `IBM_CLIENT_DEPLOY`. To view and modify server objects, access the **Properties** notebooks by clicking **Tivoli Storage Manager > Manage Servers**. Select a server and select **Manage Client Auto Deployments** from the action menu.

Resolving automatic deployment scheduling problems

If you define a schedule for your automatic backup-archive client deployment outside of the Administration Center, you must remove any blank spaces in the `-POSTNSCHEDULE` command.

If you are upgrading a V6.1 or later backup-archive client that is not using the Windows operating system, this limitation applies to you. A post-schedule command that contains spaces causes the deployment process to fail. No error messages are posted to indicate a problem and backup-archive clients are not upgraded. Replace the spaces with the number sign (#) to avoid a problem deployment.

Restarting the client operating system during a deployment

With the backup-archive client automatic deployment feature, you can specify that the client operating system is to restart after the deployment completes. As you schedule a deployment, you must decide if a restart is appropriate.

Restarting the client operating system can impact critical applications that are running on the client operating system. Applications that are not IBM Tivoli Storage Manager applications running on the client operating system are not restarted automatically. In most cases, the installation can complete without restarting the client operating system. If you do not check the box to allow a restart and one is required, most deployments are stopped and the original backup-archive client is not impacted. Stopping the deployment before removing the backup-archive client ensures that you have a working backup-archive client.

There are rare cases where the deployment manager cannot detect the restart. For example, if client processes are started from a script. In these cases, the new backup-archive client installation continues but a manual restart of the client computer is required. When you specify `AUTODEPLOY=NOREBOOT` on the command-line interface, the client operating system does not restart after the deployment completes.

Restart the client operating system when the client installation program detects that files that are in use must be removed or replaced.

HP-UX

Administration Center not supported

The Administration Center is a web-based interface for centrally configuring and managing Tivoli Storage Manager servers. The Administration Center provides wizards to help guide you through common configuration tasks. Properties notebooks allow you to modify settings and perform advanced management tasks.

In Tivoli Storage Manager Version 6.2, the Administration Center cannot be installed on HP-UX, but it can be used to manage HP-UX servers. For Administration Center system requirements, see the following website: <http://www.ibm.com/support/docview.wss?uid=swg21515628>.

Backup-archive client deployment status is unknown

There are a number of operating systems that can use the automatic backup-archive client deployment feature. Some operating systems that are not supported for the deployment feature can be used, but with limitations.

If you use an operating system that is not one of the supported operating systems, the deployment status is not reported and registers as "unknown." Also, error messages that are typically sent to the server during a deployment are not generated.

See the *IBM Tivoli Storage Manager Backup-Archive Clients Installation and User's Guide* for the supported operating systems.

Chapter 6. Resolving storage agent problems

Tivoli Storage Manager can back up and restore client data directly to and from SAN-attached storage by using the storage agent.

Checking the server activity log for storage agent information

Check the server activity log file and look at the reports occurring 30 minutes before and 30 minutes after the time of the error.

Storage agents start and manage many sessions to the server. Review the server activity log file for messages from the storage agent. To review the activity log messages, issue the **QUERY ACTLOG** command. See the *IBM Tivoli Storage Manager Administrator's Reference* for more details.

If no messages in the server activity log file are for this storage agent, verify the communication settings:

- Issue **QUERY SERVER F=D** on the server and verify that the high-level address (HLA) and low-level address (LLA) set for the server entry representing this storage agent are correct.
- In the device configuration file specified in the `dsmsta.opt` file, verify that the **SERVERNAME** as well as the HLA and LLA are set correctly in the **DEFINE SERVER** line.

Check for any error messages on the server for this storage agent.

Resolving an error caused by reading or writing to a device

If the problem is an error involving the reading or writing of data from a device, many systems and devices record information in a system error log file.

The system error log file for AIX is `errpt`, and for Windows it is the Event Log.

If a device or volume that is used by IBM Tivoli Storage Manager is reporting an error to the system error log file, it is likely a device issue. The error messages recorded in the system error log file might provide enough information to resolve the problem.

Storage agents are particularly vulnerable if path information is changed or not correct. Issue the **QUERY PATH F=D** command on the server. For each of the storage agent's paths, verify that the settings are correct. In particular, verify that the device listed matches the system device name. If the path information is not correct, update the path information by issuing the **UPDATE PATH** command.

Resolving problems caused by changing storage agent options

Changes to options in the storage agent option file might cause operations to fail, even though they had previously succeeded.

Review any changes to the storage agent option file. Try reverting the settings to their original values and retrying the operation. If the storage agent now works correctly, try reintroducing changes to the storage agent option file one-at-a-time and retry storage agent operations until the option file change that caused the failure is identified.

Resolving problems caused by changing server options or settings

Changes to options in the server option file or changes to server settings using the **SET** commands might affect the storage agent.

Review any changes to server option settings. Try reverting the settings to their original values and retrying the operation. If the storage agent now works correctly, try reintroducing changes to the storage agent option file one-at-a-time and retry storage agent operations until the option file change that caused the failure is identified.

Review server settings by issuing the **QUERY STATUS** command. If any settings reported by this query have changed, review the reason for the change and, if possible, revert it to the original value and retry the storage agent operation.

Storage agent LAN-free setup

LAN-free data movement is the direct movement of client data between a client computer and a storage device on a SAN, rather than on a LAN. You might be experiencing problems with the storage agent that are related to your LAN-free setup.

AIX **Linux** **Solaris** **Windows** To manage your LAN-free data movement, access the LAN-free configuration wizard in the Administration Center.

Resolving the issue of data being sent directly to the server

The client summary statistics do not report any bytes transferred LAN-free.

The client reports the bytes sent LAN-free by issuing the “**ANE4971I LAN-free Data Bytes: xx KB**” command. Similarly, the server does not report any instance of “**ANR0415I Session SESS_NUM proxied by STORAGE_AGENT started for node NODE_NAME**” for this node and storage agent, indicating that the LAN-free proxy operation was done for this client node.

The client will only attempt to send data LAN-free with the storage agent if the primary storage pool destination in the server storage hierarchy is LAN-free. A server storage pool is LAN-free-enabled for a given storage agent if one or more paths are defined from that storage agent to a SAN device.

To determine if the storage pool destination is configured correctly, perform the following procedures:

1. Issue the **QUERY NODE** *nodeName* command to report the policy domain to which this node is assigned.

2. Issue the **QUERY COPYGROUP** *domainName policySetName mgmtclassName F=D* command for the management classes that this node would use from their assigned policy domain. Note that this command reports information for backup files. To query copy-group information for archive files, issue the **QUERY COPYGROUP** *domainName policySetName mgmtclassName TYPE=ARCHIVE F=D* command.
3. Issue the **QUERY STGPOOL** *stgpoolName* command, where *stgpoolName* is the destination reported from the previous **QUERY COPYGROUP** queries.
4. Issue the **QUERY DEVCLASS** *deviceClassName* command for the device class used by the destination storage pool.
5. Issue the **QUERY LIBRARY** *libraryName* command for the library reported for the device class used by the destination storage pool.
6. Issue the **QUERY DRIVE** *libraryName F=D* command for the library specified for the device class used by the destination storage pool. If no drives are defined to this library, review the library and drive configuration for this server and issue the **DEFINE DRIVE** command to define the needed drives. If one or more of the drives report "ONLINE=No" evaluate why the drive is offline and, if possible, update it to online by issuing the **UPDATE DRIVE** *libraryName driveName ONLINE=YES* command.
7. Issue the **QUERY SERVER** command to determine the name of the storage agent as defined to this server.
8. Issue the **QUERY PATH** *stgAgentName* command, where *stgAgentName* is the name of the storage agent defined to this server and reported in the **QUERY SERVER** command. Review this output and verify that one or more paths are defined for drives defined for the device class used by the destination storage pool. If no paths are defined for this storage pool, issue the **DEFINE PATH** command to define the needed paths. Also, review this output and verify that the path is online. If paths are defined but no paths are online, update the path to online by issuing the **UPDATE PATH** *srcName destName SRCTYPE=SERVER DESTTYPE=DRIVE ONLINE=YES* command.

Resolving a disqualified LAN-free-enabled storage pool

The server disqualifies a storage pool from being a LAN-free-enabled storage pool if it was configured for simultaneous-write operations.

In this case, data from the client is sent directly to a server that will not be using a LAN-free storage pool.

Issue the **QUERY STGPOOL** *stgpoolName F=D* command for the destination storage pool for this client. If the storage pool is set for simultaneous-write operations, the "Copy Storage Pool(s):" value references one or more other storage pool names and IBM Tivoli Storage Manager interprets the simultaneous-write operation to be a higher priority than the LAN-free data transfer. Because simultaneous-write operations are considered a higher priority operation, this storage pool is not reported as LAN-free-enabled and as such, the client sends the data directly to the server. The storage agent does not support simultaneous-write operations.

Ensuring that data is transferred using a LAN-free environment

The storage agent and client are both able to manage failover directly to the server, depending upon the LAN-free configuration and the type of error encountered.

Because of this failover capability, it might not be apparent that data is being transferred over the LAN when it was intended to be transferred LAN-free. It is possible to set the LAN-free environment to limit data transfer to only LAN-free.

To test a LAN-free configuration, issue the **UPDATE NODE** *nodeName* **DATAWRITEPATH=LAN-FREE** command for the client node whose LAN-free configuration you want to test. Next, try a data storage operation such as backup or restore. If the client and storage agent attempt to send the data directly to the server using the LAN, the following error message is received:

```
ANR0416W Session sessionNumber for node nodeName not allowed to operation
using path data transfer path
```

The *operation* reported indicates either READ or WRITE, depending upon the operation attempted. The path is reported as LAN-free.

If this message is received when you are trying a LAN-free operation, evaluate and verify the LAN-free settings. Generally, if data is not sent LAN-free when the client is configured to use LAN-free, the storage pool destination for the policy assigned to this node is not a LAN-free enabled storage pool, or the paths are not defined correctly.

Chapter 7. Using trace to resolve problems

IBM Tivoli Storage Manager can, at times, experience problems that you can resolve through trace.

Trace classes for the Administration Center

AIX Linux Solaris Windows

The Administration Center provides individual and aggregate trace classes. Aggregate trace classes give you a way to enable many related trace classes by specifying the aggregate trace class name.

The trace classes documented in Table 14 are those that are typically requested or used for diagnosing problems. The trace classes that are included in the table do not comprise all possible trace classes that are available.

Use the trace class name when activating tracing using the Administration Center Support utility.

Table 14. Trace classes for the Administration Center

Trace class	Description	Uses
ADMCNTR	Displays processing information for the portlet and servlet classes. ADMCNTR is an aggregate trace class that includes most of the Administration Center trace classes.	Use this trace class for diagnosing problems with the Administration Center when the health monitor and portlet services are working correctly. This trace class is an aggregate of all trace classes except SERVICES and CONNS.
ADMINAPI	Displays information in the processing of the application programming interface (API). The API is used for communications between the Administration Center and the IBM Tivoli Storage Manager server.	Use this trace class for problems related to command implementation on the Tivoli Storage Manager server.
COMPS	Displays all processing information for the components of the Administration Center. COMPS is an aggregate trace class that includes all the component trace classes, such as DRM, MPLAN, and DOMAIN.	Use this trace class when the nature of the problem is unknown. If you use the COMPS trace flag, it is not necessary to specify any other component trace flags because this trace class already includes those trace classes.
CONNS	Displays all processing information for the Tivoli Storage Manager server connections portlet service of the Administration Center.	Use this trace class for problems related to connections portlet service of the Administration Center.

Table 14. Trace classes for the Administration Center (continued)

Trace class	Description	Uses
CONTROL	Displays processing information for the connection manager, the portlet factory, and other key control components.	Use this trace class for problems related to control objects. You can also use this trace class to debug problems that occur when you perform actions that use internal control objects, such as the connection manager.
DEPLOY	Displays processing information for the classes related to the deployment of backup-archive client maintenance packages.	Use this trace class for problems related to definitions and operations for deploying backup-archive client maintenance packages.
DOMAIN	Displays processing information for the domain-related classes.	Use this trace class for problems related to Tivoli Storage Manager domain-related operations.
DRM	Displays processing information for the disaster recovery manager (DRM) related classes.	Use this trace class for problems related to Tivoli Storage Manager DRM operations.
ENTMGMT	Displays processing information for the enterprise management related classes.	Use this trace class for problems related to Tivoli Storage Manager enterprise management operations.
FASTBACK	Displays processing information for the Tivoli Storage Manager FastBack server-related classes.	Use this trace class for problems related to FastBack servers that are defined to the Administration Center.
HEALTH	Displays all information about the health monitor.	Use this trace class for problems related to the health monitor.
MPLAN	Displays processing information for the maintenance-plan-related classes.	Use this trace class for problems related to the Tivoli Storage Manager disaster maintenance plan operations.
NAS	Displays processing information for the network-attached storage (NAS) device-related classes.	Use this trace class for problems related to the NAS device definitions and operations.
NODES	Displays processing information for the client node-related classes.	Use this trace class for problems related to the Tivoli Storage Manager client node definitions and operations.
PS	Displays general information in the processing of presentation services object extensions. This traces all listener, validator, and JavaServer Page (JSP) bean classes.	Use this trace class for problems related to the graphical user interface objects, such as tables and buttons.

Table 14. Trace classes for the Administration Center (continued)

Trace class	Description	Uses
REPORTS	Displays processing information for the Administration Center reports-related classes.	Use this trace class for problems related to the reports generated and displayed by the Tivoli Storage Manager Administration Center.
SERVERS	Displays processing information for the Tivoli Storage Manager server-related classes.	Use this trace class for problems related to Tivoli Storage Manager servers that are defined to the Administration Center.
SERVICE	Displays all processing information for the portlet services of the Administration Center. SERVICE is an aggregate trace class that includes the trace class CONNS.	Use this trace class for problems related to the portlet services of the Administration Center. You can also use this trace class when the nature of the problem is unknown. If you use the SERVICE trace flag, it is not necessary to specify any other service trace flags because they are included.
STGDEVS	Displays processing information for the storage device-related classes.	Use this trace class for problems related to storage devices for the Tivoli Storage Manager server definitions and operations.
UTIL	Displays processing information for the utility classes.	Use this trace class for problems related to utility objects. Use this trace class to debug problems that occur in various internal utility routines.

Related reference:

“Administration Center Support utility” on page 111

Enabling Administration Center trace

Tracing is available for the IBM Tivoli Storage Manager Administration Center, but you must first enable it.

HP-UX

Administration Center not supported

The Administration Center is a web-based interface for centrally configuring and managing Tivoli Storage Manager servers. The Administration Center provides wizards to help guide you through common configuration tasks. Properties notebooks allow you to modify settings and perform advanced management tasks.

The Administration Center cannot be installed on HP-UX, but it can be used to manage HP-UX servers. For Administration Center system requirements, see the following website: <http://www.ibm.com/support/docview.wss?uid=swg21515628>.

Perform the following steps to enable tracing for the Administration Center:

1. Start the Administration Center support utility. This utility is located in the Tivoli Storage Manager installation directory in *TSM_HOME\IBM\tivoli\tsmac*. Open a command prompt window, navigate to the directory, and start the utility by typing *supportUtil.bat* for Windows or *supportUtil.sh* for AIX, Linux, or Oracle Solaris.
2. Choose which trace classes to enable or simply turn on all tracing. See the trace classes for the Administration Center for more information about the different trace classes. Follow the on-screen instructions to enable tracing.

```
C:TSM_HOME\IBM\tsmac>supportUtil
Administration Center Support Utility - Main Menu
=====
1. Manage Administration Center tracing
2. Manage the maximum memory size the Administration Center can use
3. Manage the Administration Center session timeout setting
4. Collect trace files, logs and system information to send to support
5. Generate a heap dump of the Java virtual machine
6. Generate a Java core dump of the Java virtual machine
7. View the log file for this utility

9. Exit
Enter Selection: 1

Administration Center Support Utility - Manage Tracing
=====
1. Turn all tracing on
2. Turn all tracing off
3. Turn a single trace class on
4. View the current trace specification

99. Return to main menu

Enter Selection: 3

Administration Center Support Utility - Turn Single Class On
=====
1. ADMCNTR - traces the Java classes in com.tivoli.dsm.admcntr*
2. CONTROL - traces the Java classes in com.tivoli.dsm.admcntr.cntl*
3. UTIL - traces the Java classes in com.tivoli.dsm.admcntr.util*
4. PS - traces the Java classes in com.tivoli.dsm.admcntr.ps*
5. COMPS - traces the Java classes in com.tivoli.dsm.admcntr.comp*
6. SERVERS - traces the Java classes in com.tivoli.dsm.admcntr.comp.servers*
7. ENTMGMT - traces the Java classes in com.tivoli.dsm.admcntr.comp.entmgmt*
8. DRM - traces the Java classes in com.tivoli.dsm.admcntr.comp.drm*
9. MPLAN - traces the Java classes in com.tivoli.dsm.admcntr.comp.mplan*
10. STGDEVS - traces the Java classes in com.tivoli.dsm.admcntr.comp.stgdevs*
11. REPORTS - traces the Java classes in com.tivoli.dsm.admcntr.comp.reports*
12. DOMAIN - traces the Java classes in com.tivoli.dsm.admcntr.comp.domain*
13. NAS - traces the Java classes in com.tivoli.dsm.admcntr.comp.naswiz*
14. NODES - traces the Java classes in com.tivoli.dsm.admcntr.comp.nodes*
15. FASTBACK - traces the Java classes in com.tivoli.dsm.admcntr.comp.fastback*
16. DEPLOY - traces the Java classes in com.tivoli.dsm.admcntr.comp.deploy*
17. ADMINAPI - traces the Java classes in com.tivoli.dsm.adminapi*
18. SERVICE - traces the Java classes in com.tivoli.dsm.service*
19. HEALTH - traces the Java classes in com.tivoli.dsm.service.health*
20. CONNS - traces the Java classes in com.tivoli.dsm.service.conns*
```

- 98. Return to manage tracing menu
- 99. Return to main menu

Enter Selection:

3. Repeat the task that is causing the error. A trace of the error is created and written to the trace file. For Windows operating systems, the trace file is located in the Tivoli Integrated Portal directory at `C:\TSM_HOME\IBM\tivoli\tipv2\profiles\TIPProfile\logs\server1`. *TSM_HOME* represents the installation directory. For operating systems other than Windows, the trace file is in the `TSM_HOME/opt/IBM/tivoli/tipv2/profiles/TIPProfile/logs/server1` directory.
4. Disable tracing. Leaving tracing enabled might cause performance problems.
5. Collect trace files, log files, and system information to send to IBM Support. This procedure packages the trace information collected and the log files into a single file called `AdminCenterSupport.zip`. This file can then be sent to IBM Support.

Related reference:

“Administration Center Support utility” on page 111

“Trace classes for the Administration Center” on page 125

Enabling a trace for the server or storage agent

You can issue trace commands from the following places: the server console, storage agent console, administrative client connected to either the server or storage agent, server options file (`dsmserv.opt`), or the storage agent options file (`dsmsta.opt`).

Trace commands apply to the server or storage agent to which the command was submitted. Trace commands in the options files are used to trace the applications during startup and initialization or to provide a default set of trace classes. There is one trace class (**ADDMSG**) that is always enabled by default, whether it appears on the options file or not. It is best to trace to a file. Typically, the tracing for the server or storage agent will generate a large amount of output.

Perform the following steps to enable trace classes for the server or storage agent:

1. Determine the trace classes to enable. To have trace messages issued for a given trace class, that trace class needs to be enabled either prior to beginning the trace or after the tracing has begun.
2. Issue the **TRACE ENABLE** *traceClassName* command to enable one or more trace classes. Note that *traceClassName* might be a space-delimited list of trace classes. For example, this command could be entered as **TRACE ENABLE TM SESSION**. The **TRACE ENABLE** command is cumulative, such that extra trace classes can be enabled by issuing **TRACE ENABLE** numerous times. For example, if you wanted to add the PVR trace class in addition to those that are already enabled, issue: **TRACE ENABLE PVR**. To stop having trace messages issued for a given trace class, that trace class needs to be disabled either prior to beginning the trace or after the tracing begins.
3. Issue the **TRACE DISABLE**<*traceClassName*> command to disable one or more trace classes. Note that *trace class name* might be a space delimited list of trace classes. For example, this command could be entered as **TRACE DISABLE TM SESSION**. Additional trace classes can also be disabled by issuing **TRACE DISABLE**. For example, if you wanted to remove the PVR trace class in addition to those that were already disabled, issue: **TRACE DISABLE PVR**. By issuing **TRACE DISABLE** without specifying any trace classes, all currently enabled trace classes are disabled.

4. Tracing can occur to the console or to a file. Perform the following tasks to begin tracing:
 - For tracing to the console, issue: **TRACE BEGIN**
 - For tracing to a file with no size limitation, issue: **TRACE BEGIN *fileName***
 - For tracing to a file with a size limitation, issue: **TRACE BEGIN *fileName* MAXSIZE= *maximum size in megabytes***

Note: The *fileName* can be a fully-qualified path such as /opt/tmp or c:\temp. If a full path is not given, the trace file will be located in the same directory as the running executable file.

5. Perform the operation that is causing the problem.
6. Issue the **TRACE END** command to stop trace messages from being issued. If tracing is being done to a file, ending the trace writes any remaining trace messages to the file and closes the file.

It is possible to enable tracing and begin it using the server or storage agent options file. The commands and syntax discussed are the exact same for the server or storage agent options file, and they are generally used to trace startup and initialization of the server. For example, if the following lines were added to the server's option file, tracing would be started for the DB, TM, and LOG trace classes, and the trace messages written to the file MYTRACE.OUT.

```
TRACE ENABLE DB TM LOG
TRACE BEGIN MYTRACE.OUT BUFSIZE=4096
```

Remember: If you are conducting a trace due to a server crash, do not set the **BUFSIZE** parameter.

Related reference:

“Trace classes for a server or storage agent” on page 131

Enabling a stack trace for specific messages for the server or storage agent

A stack trace reveals information about an application that IBM Software Support can use to help you diagnose your problems faster.

Remember: Stack trace can be extremely robust and might flood the activity log file, depending on the frequency of the failure. You might not be able to view the activity log file, therefore you might want to disable stack trace after completion.

IBM Software Support might find it helpful to enable stack trace on specific messages issued by the server or storage agent. The types of messages on which a stack trace can be enabled are server console, storage agent console, and the administrative client connected to either the server or storage agent.

To get a stack trace when a specific message is issued by the server or storage agent, enable the message for stack trace. Issue the **MSGSTACKTRACE ENABLE <messageNumber>** command to enable one or more messages for stack trace.

Remember: *<messageNumber>* might be a space-delimited list of message numbers.

This command could be entered as **MS ENABLE 2017**. The **MSGSTACKTRACE ENABLE** command is cumulative, such that extra messages are enabled by issuing the **MSGSTACKTRACE ENABLE** command additional times. If you want to add message 985,

in addition to those that are already enabled, issue **MS ENABLE 985**. Notice that only the number part of the message is allowed in the **MSGSTACKTRACE** command. To stop getting stack trace for messages issued by the server or storage agent, the stack trace for those messages needs to be disabled. Issue the **MSGSTACKTRACE DISABLE <messageNumber>** command to disable one or more messages.

Note that *<messageNumber>* might be a space-delimited list of message numbers. For example, this command could be entered as **MSGSTACKTRACE DISABLE 2017 985**. Additional messages can also be disabled by issuing **MS DISABLE**. For example, if you want to remove message number 7837 in addition to those that are already disabled, issue **MSGSTACKTRACE DISABLE 7837**.

The following messages are enabled for stack trace by default:

435	437	486	661	685	727	728	780	781	782
784	785	786	790	793	794	860	881	882	883
884	1032	1078	1092	1117	1156	1227	5010	5015	5019
5021	5093	5099	5100	5267	6753	7823	7837	9600	9601
9602	9604	9605	9606	9607	9608	9999			

Trace classes for a server or storage agent

The server and storage agent provide aggregate trace classes. These trace classes are a shortcut for using many related trace classes by specifying the aggregate trace class name for the **TRACE ENABLE** command.

The trace classes listed in Table 15 are those trace classes that are most typically requested or used for diagnosing problems. This table does not include all possible trace classes that are available. The trace class name is used with the **TRACE ENABLE** and **TRACE DISABLE** commands.

Table 15. Server or storage agent trace classes

Trace Classes	Description	Uses
ADDMSG	Issues console messages (ANR, ANE,...etc.) to the trace file.	This trace class is valuable for correlating server messages to trace messages and for preserving the timing for when each was issued.
ADMCMMD	Traces related to command processing.	Use this trace class to debug the command interpreter, including the PARALLEL and SERIAL command handling.
AF	This trace class displays information about user data stored on sequential media devices. AF is an aggregate trace class that uses AFCREATE, AFMOVE, AFLOCK, AFTXN, and AFCOPY. Typically recommended to issue TRACE DISABLE AFLOCK , unless the locking information is explicitly requested or needed.	Use this trace class to diagnose problems reading or writing user files to sequential media volumes.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
AFCREATE	This trace class displays information about storing user data on sequential media volumes.	Use this trace class to diagnose writing user data on sequential media volumes.
AFMOVE	This trace class displays operations that move user data with sequential media volumes. Move operations are performed by MIGRATION, RECLAMATION, MOVE DATA, and MOVE NODEDATA server processes.	Use this trace class to diagnose problems with the data movement server processes.
AS	This trace class displays information volume selection and assignment, coordination of drives (mount points), and management of data placement on volumes. This aggregate trace class uses ASALLOC, ASRTRV, ASDEALLOC, ASMOUNT, ASVOL, ASTXN, and ASSD. The typical method is to issue TRACE DISABLE ASTXN, unless the locking information is explicitly requested or needed.	Use this trace class to diagnose many different problems relating to volumes, mount points, or data read and write operations.
ASALLOC	This trace class displays information about reserving and allocating space on sequential media volumes for storing data. This is for storing data on behalf of a client session or for server data movement operations such as MIGRATION, RECLAMATION, MOVE DATA, or MOVE NODEDATA.	Diagnose problems where the server or storage agent report no space available but space is supposed to be available in the storage hierarchy.
ASDEALLOC	This trace class displays information about releasing and de-allocating space on sequential media volumes for storing data. Typical deallocation operations on the server are EXPIRATION, MIGRATION, RECLAMATION, MOVE DATA, MOVE NODEDATA, AUDIT VOLUME, DELETE VOLUME, and DELETE FILESPACE.	Use this trace class to diagnose during the deletion of data.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
ASMOUNT	This trace class displays information about drive (mount point) selection and assignment for sequential media devices.	Diagnose situations where sessions or processes are waiting for mount points or cases where an operation fails because no mount point is available. Also useful in cases where a mount point is pre-empted.
ASRTRV	This trace class displays information about reading data from sequential media volumes.	Use this trace class to diagnose problems reading data such as RESTORE or RETRIEVE client by the client, or MIGRATION, RECLAMATION, STORAGE POOL BACKUP, AUDIT VOLUME, GENERATE BACKUPSET, EXPORT, MOVE DATA, or MOVE NODEDATA by the server.
ASTXN	This trace class displays information about transactions used to make database updates relating to information for sequential media volumes, storage pools, device classes, and other attributes.	Use this trace class to diagnose stoppages, database operations, failures reported for sequential media operations, or general data storage problems.
ASVOL	This trace class displays information about volume selection and assignment for sequential media volumes.	Use this trace class to diagnose situations where sessions or processes are waiting for volumes, or cases where an operation fails because no volume is available. Also useful in cases where volume access is pre-empted.
ASSD	This trace class displays information about sequential stream data operations. These are operations that use sequential media device classes, volumes, or mount points but do not store data in the storage hierarchy. Server processes that perform sequential stream data operations are BACKUP DB, EXPORT/IMPORT, and GENERATE BACKUPSET .	Use this trace class to diagnose server processes that perform sequential stream data operations.
BF	Information about user data (files) stored in the storage hierarchy. This aggregate trace class uses BFCREATE, BFRTRV, BFSALVAGE, BFLOCK, BFLAGGR, BFREMOTE, and BFTRG .	Use this trace class to diagnose general data read-or-write problems for client operations and server processes.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
BFAGGR	This trace class displays information about server aggregation of user data. The server aggregates many smaller user files into a larger file in the storage hierarchy to optimize performance for data movement operations such as MIGRATION , MOVE DATA , and MOVE NODEDATA .	Use this trace class to diagnose general data read-or-write problems for client operations and server processes, or both.
BFCREATE	This trace class displays information about client operations that store data in the storage hierarchy. Typically, these client operations are BACKUP , ARCHIVE , or SPACE MANAGE operations by the client.	Use this trace class to diagnose failures or other problems while a client is storing data.
BFREMOTE	Traces the first stage of NDMP (Network Data Management Protocol) backup and restore processes.	This trace class is used to identify NDMP-related backup or restore operations. These trace classes are specific to the functions which implement the NDMP protocol. The SPID trace class provides more detailed tracing, including tracing all NDMP file history records sent by the NDMP file server.
BFRTVR	This trace class displays information about client operations that read data from the storage hierarchy.	Use this trace class to diagnose failures or other problems while a client is reading data.
BITVECTOR	Diagnoses problems where the server reports problems with disk storage pools.	Use this trace class to display information about reserving and allocating space on volumes in disk storage pools.
BKSET/OBJSET	Trace class for backup set functions. The BKSET and OBJSET trace classes are synonymous.	Use this trace class to debug problems in the GENERATE BACKUPSET command or during a client restore operation from a backup set.
BLKDISK	Trace class for viewing disk I/O activity to storage pool, database, and log volumes.	Use this trace class to view I/O activity to disk to diagnose performance and disk I/O errors.
BRNODE	Trace class for the BACKUP and RESTORE NODE commands, used during NDMP operations.	Use this trace class to debug problems in the BACKUP and RESTORE NODE commands.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
COLLOCATE	This trace class displays information about collocation processing on storage pools. COLLOCATEDetail trace class can also be used to get more detailed information about the collocation processing, such as files being processed for a collocation group. Files being processed for a collocation group can cause a large number of output trace statements.	Use this trace class to diagnose problems with collocation processing.
CRC	This trace class displays information about generating and managing cyclic redundancy checks (CRCs) on the server or storage agent. CRC is an aggregate trace class that uses CRCDATA , CRCPROTO , and CRCVAL .	Use this trace class to diagnose data corruption issues where CRC processing did not report data corruption.
CRCDATA	This trace class displays information about generating and managing CRCs for data stored in storage pools with CRCDATA=YES set.	Use this trace class to diagnose data corruption issues where CRC processing did not report data corruption.
CRCPROTO	This trace class displays information about generating and managing CRCs for data exchanged between the client and either the server or storage agent where this node is configured with VALIDATEPROTOCOL=ALL or VALIDATEPROTOCOL=DATAOnly on the server.	Use this trace class to diagnose data corruption issues where CRC processing did not report data corruption.
CRCVAL	This trace class displays information about generating and comparing CRC values.	Informational for showing CRC values during processing.
CRYPTO	This trace class displays information about advanced encryption standard (AES) operations and some general encryption settings.	Use this trace class to isolate and identify encryption-related problems.
DBCLI	Traces the general set of interactions.	Use this trace class to trace the general set of DB2 interactions and the DB2 command-line interface.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
DBCONN	Traces connection activities.	Use this trace class to trace IBM Tivoli Storage Manager connections to DB2 connections. This trace class shows such things as the creation of connection handles and the assignment of connections to transactions.
DBDBG	Traces debugging processes. Consider using this trace class first when debugging a database issue.	Use this trace class to show function entry or exit, exit return codes, and the statements that are built and are being run.
DBITXN	Traces database transaction-related activities. Transaction-related activities concern transaction latch acquisition and release, dbTxnDesc allocation and release, and transaction commit processing from the prepare and commit phase functions.	Use this trace class to trace transaction-related activities for the database interface.
DBNETDB	This trace class displays information about LAN-free operations and the negotiation and management of information between the server and storage agent. Typically applies to server and storage agent prior to Tivoli Storage Manager Version 5.2. After Tivoli Storage Manager Version 5.2, DBNETDB still displays information but is superseded by the LANFREE trace class. The usual practice is to use this trace class on both the server and storage agent.	Use this trace class to diagnose problems with LAN-free data movement.
DBRC	Traces the return codes from functions in the database component.	Use this trace class to trace the return codes.
DEDUP	Traces the general logic path tracing for data deduplication processing. Does not typically include error paths.	Use DEDUP to trace general logic paths for data deduplication processing.
DEDUP1	Traces error paths for data deduplication processing.	Use DEDUP1 to trace error paths for data deduplication processing.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
DEDUP2	Traces the fingerprinting and digital signatures path.	Use DEDUP2 to trace fingerprinting and digital signature paths.
DELTA	Trace class for logical group functions. The DELTA and GROUP trace classes are synonymous.	Use this trace class to debug problems with logical groups, whether delta-base groups (subfile backup) or peer groups (Windows SYSTEM OBJECT or image backups). Group processing is relevant during just about any operation that references backup objects. The backup objects can include client backup and restore, expiration, deletion (DELETE FILESPACE , DELETE VOLUME), export/import, backup set generation and restore, no-query restore, db audit, and others.
DF	This trace class displays information about user data stored on disk volumes. DF is an aggregate trace class that enables DFCREATE , DFRTRV , DFMOVE , DFLOCK , DFTXN , and DFCOPY . Issue the TRACE DISABLE DFLOCK command unless the locking information is explicitly requested or needed.	Use this trace class to diagnose problems reading or writing user files to disk volumes.
DFCREATE	This trace class displays information about storing user data on disk volumes.	Use this trace class to diagnose writing user data on disk volumes.
DFMOVE	This trace class displays operations that move user data by using disk volumes. Move operations are performed by the MIGRATION , MOVE DATA , and MOVE NODEDATA server processes.	Use this trace class to diagnose problems with the data movement server processes.
DFRTRV	This trace class displays information about reading user data from disk volumes.	Use this trace class to diagnose reading user data from disk volumes.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
DS	This trace class displays information about volume selection, space reservation, assignment, and management of data placement on disk volumes. DS is an aggregate trace class that enables DSALLOC, DSRTRV, DSDEALLOC, and DSVOL. Issue TRACE DISABLE DSTXN unless the locking information is explicitly requested or needed.	Use this trace class to diagnose many different problems relating to disk volume data read-and-write operations.
DSALLOC	This trace class displays information about reserving and allocating space on disk volumes for storing data. The data storing might be done on behalf of a client session or for server data movement operations such as MIGRATION, MOVE DATA, or MOVE NODEDATA .	Diagnose problems where the server or storage agent report that no space is available, but there appears to be space available in the storage hierarchy.
DSDEALLOC	This trace class displays information about releasing and de-allocating space on disk volumes. Typical deallocation operations on the server are EXPIRATION, MIGRATION, MOVE DATA, MOVE NODEDATA, AUDIT VOLUME, DELETE VOLUME, and DELETE FILESPACE .	Use this trace class to diagnose during the deletion of data.
DSRTRV	This trace class displays information about reading data from disk volumes.	Use this trace class to diagnose problems reading data such as RESTORE or RETRIEVE client by the client, or MIGRATION, STORAGE POOL BACKUP, AUDIT VOLUME, GENERATE BACKUPSET, EXPORT, MOVE DATA, or MOVE NODEDATA by the server.
DSVOL	This trace class displays information about volume selection and assignment for disk volumes.	Use this trace class to diagnose situations where sessions or processes are waiting for volumes, or cases where an operation fails because no volume is available.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
ICVOLHST	Trace class for volume history functions.	Use this trace class to debug problems with creating volume history entries (for example: during EXPORT , BACKUP DB , or GENERATE BACKUPSET) or deleting volume history entries (for example: during DELETE VOLHISTORY).
IMFS	Trace class for file space functions.	Use this trace class to debug problems related to inventory file spaces (for example, during DELETE FILESPACE).
LANFREE	This trace class displays general information about LAN-free operations on either the server or storage agent. Also shows error information for LAN-free-related operations. LANFREE is an aggregate trace class that enables LNFVERB, LNFMEM, LNFENTRY, and LNFDATA.	Any LAN-free failure.
MMS	This trace class displays information about tape libraries and the server or storage agent use of these. MMS is an aggregate trace class that enables MMSBASE, MMSTXN, MMSLIB, MMSDRIVE, MMSOP, MMSMAN, MMSSCSI, MMSFLAG, MMSACSLs, and MMSSHARE. Include NA and PVR trace classes when tracing MMS (suggested).	Used to diagnose problems with tape libraries, library volume inventories, or other general library issues.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
NA	<p>This trace class displays information about path information for the server or storage agent. This relates to the DEFINE PATH, UPDATE PATH, DELETE PATH, and QUERY PATH commands. This trace class is also useful for identifying issues related to operations involving NDMP file servers, for example, DEFINE DATAMOVER, UPDATE DATAMOVER, BACKUP NODE, and RESTORE NODE commands.</p> <p>This aggregate trace class uses NALOCK, NAPATH, NAMOVER, NADISK, and NACONFIG. It might be best to include MMS and PVR trace classes when tracing NA.</p>	<p>Use this trace class to diagnose problems with paths to devices.</p>
PRODCONS	<p>If there are problems with work getting dispatched to batches, PRODCONS displays information about the problem and whether it is in the PC object or in replication.</p>	<p>Use PRODCONS to trace the internal workings of the Producer/Consumers objects used in the server.</p>
PROXYNODE	<p>This trace class displays information about proxynode sessions and the commands related to proxynode associations (GRANT, REVOKE, QUERY PROXYNODE).</p>	<p>Use this trace class to diagnose problems with proxynode sessions and related commands. It might be best to include SESSION trace when analyzing proxynode session problems.</p>
PVR	<p>This trace class displays information about sequential media devices and the server or storage agent use of these devices. PVR is an aggregate trace class that enables PVRVOL, PVRCLASS, and PVRMP.</p> <p>The PVR trace class contains everything in the PVRIO aggregate trace class and the PVRNOIO trace class.</p>	<p>Use this trace class to diagnose problems with tape drives, failures reading or writing tape volumes, or other tape-volume-related issues.</p>
PVRIO	<p>This trace class displays tracing of read, write, or POS operations for sequential media devices and the server or storage agent use of these devices.</p>	<p>Use this trace class to diagnose problems with tape drive failures reading or writing tape volumes.</p>

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
PVRNOIO	This trace class displays PVRVOL, PVRCLASS, and PVRMP information.	Use this trace class to diagnose problems with tape drive mounts or other tape-volume-related issues.
REPL	REPL is an aggregate trace class that enables REPLBATCH, REPLCMD, REPLFS, REPLINV, REPLPROC, REPLSTATS, and REPLSESS.	Use this trace class to diagnose problems with replication.
REPLBATCH	This trace class displays tracing related to batch processing, where individual files are sent from the source server to the target server.	Use this trace class to diagnose replication problems with batch processing.
REPLCMD	This trace class displays tracing related to command parsing and the resolution of file space replication rules.	Use this trace class to diagnose replication problems with command parsing and the resolution of file space replication rules.
REPLFS	This trace class displays tracing related to the iteration of the file spaces to decide what is to be replicated, updated, or deleted.	Use this trace class to diagnose replication problems with iterating file spaces to decide what is to be replicated, updated, or deleted.
REPLINV	This trace class displays tracing related to the inventory updates (IM tables) as part of replication.	Use this trace class to diagnose replication problems with inventory updates.
REPLPROC	This trace class displays tracing of the overall replication process. This is the main thread and dispatcher.	Use this trace class to diagnose replication problems with the replication process.
REPLSESS	This trace class displays tracing related to the establishment of sessions for replication, including the session management on both source and target servers.	Use this trace class to diagnose replication problems with the establishment of sessions.
REPLSTATS	This trace class displays tracing related to updating the statistics as replication runs. Also includes insertion or update of history records in the replication history table.	Use this trace class to diagnose replication problems with statistical updates.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
RETPROT	Trace class for the archive retention protection functions.	Use this trace class to debug problems using the RETINIT and RETMIN parameters in the archive copy group. You can also use this trace class for problems caused by using the VB_SignalObject verb (only supported via the client API) to signal an object's event or to hold or release an object, or problems during expiration or deletion of retention protected objects.
ROWMGR	Traces activities for row-based operations. Row-based operations are the following operations: <ul style="list-style-type: none"> • Abbrev • Delete • Fetch • FetchNext • FetchPrev • Insert • SearchBounds • Update 	Use this trace class to trace the activities for row-based operations.
SCHED	Trace class for the central scheduler functions. This trace class applies to classic and enhanced schedules equally.	Use this trace class to debug problems related to schedule commands like DEFINE/UPDATE/QUERY SCHEDULE or DEFINE ASSOCIATION . Also use this trace class to debug problems related to the central scheduler background processes, such as the schedule manager and schedule prompter.
SESSION	This trace class displays information about sessions connected to the server, including all verbs sent and received by the server.	This trace class is generally recommended for protocol violations, transaction processing errors, or in cases where the client is stopped and not responding.
SESSREMOTE	Traces communication between the Tivoli Storage Manager server and the Tivoli Storage Manager client during NDMP backup and restore operations.	This trace class is used to identify NDMP-related backup or restore operations that are initiated using the Tivoli Storage Manager web or command-line client.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
SHRED	This trace class displays information related to data-shredding operations on the server.	This trace class is used to diagnose problems with data shredding. Data shredding is only applicable if one or more storage pools on the server has a non-zero value for the SHRED attribute. Activity that is related to data shredding occurs primarily during the EXPIRE INVENTORY, DELETE FILESPACE, DELETE VOLUME, MOVE DATA, MIGRATE, and SHRED DATA commands. Other trace classes that report activity related to data shredding are BFDESTROY, DFDESTROY, DSALLOC, DSDEALLOC, and CRCDATA .
SPI/SPID	Traces the server NDMP protocol interface.	The SPI and SPID trace classes are used to identify issues related to NDMP backup or restore operations of NAS file servers. These trace classes are specific to the functions that implement the NDMP protocol and communicate with a NAS file server. The SPID trace class provides more detailed tracing, including tracing all NDMP file history records sent by the NAS file server.
SSLDATA	Detailed Secure Sockets Layer (SSL) trace used to display byte-level information about data that is sent or received between the backup-archive client and the server.	Use the SSLDATA trace class to debug the session data corruption issues that might be caused by SSL that is running through the SSLTCP or SSLTCPADMIN server options. Because this is a byte-level trace, it can collect a large amount of data.
SSLINFO	General SSL trace used to display setup and characteristics of SSL sessions between the backup-archive client and the server.	Use the SSLINFO trace class to debug session connection and handshake errors that might be caused by the SSL that is running through the SSLTCP or SSLTCPADMIN server options. This can be used in tandem with the TCPINFO and SESSION trace classes.
TBREORG	Collects information regarding table and index reorganization activities initiated by the server.	Use the TBREORG trace class to debug server-initiated reorganization activity.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
TBLMGR	Traces activities for table-based operations.	Use the TBLMGR trace class to view table-based operations such as table registration, table open, and table close.
TCP	Collects information regarding TCP/IP used between the client and either server or storage agent. TCP is an aggregate trace class. It enables TCPINFO and TCPERROR.	Use this trace class to debug session connection errors or data corruption issues that might be caused by the network.
TCPDATA	Detailed TCP/IP trace used to display byte-level information about data that is sent or received.	Use this trace class to debug session data corruption issues that might be caused by the network.
TCPINFO	General TCP/IP trace used to display setup and characteristics of TCP/IP on the server or storage agent.	Use this trace class to debug session data corruption issues that might be caused by the network.
TEC	Provides information regarding events sent to a TEC server. Corresponds to the 'tivoli' event receiver.	To debug connection issues encountered with TEC event logging.
TOC	General trace class for the Table Of Contents (TOC) component, used during file-level NDMP operations. TOC is an aggregate trace class that enables TOCBUILD, TOCLOAD, TOCREAD, and TOCUTIL.	Use this trace class to debug problems during file-level NDMP operations, such as an NDMP backup with the TOC=YES parameter, or an NDMP restore with the FILELIST parameter.
TOCBUILD	Table Of Contents (TOC) build functions.	Use this trace class to debug problems during an NDMP backup with the TOC=YES parameter.
TOCLOAD	Table Of Contents (TOC) load functions.	Use this trace class to debug problems while displaying files and directories on the client graphical user interface (GUI).
TOCREAD	Table Of Contents (TOC) read functions.	Use this trace class to debug problems during a QUERY TOC command or while trying to load a TOC for displaying files and directories on the client GUI.
TOCUTIL	Table Of Contents (TOC) utility functions.	Use this trace class to debug problems related to TOC component initialization or TOC retention.

Table 15. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
UNICODE	This trace class displays information about code page conversions and Unicode filespace operations.	Use this trace class to debug problems related to code page conversion problems or unicode filespace problems.
XI	This trace class displays general information for the IMPORT and EXPORT commands.	Use this trace class to debug problems related to IMPORT and EXPORT commands.

Related tasks:

“Enabling a trace for the server or storage agent” on page 129

Show commands for the server or storage agent

SHOW commands are unsupported diagnostic commands that are used to display information about in-memory control structures and other runtime attributes. The **SHOW** commands are used by development and service only as diagnostic tools. Several **SHOW** commands exist for the backup-archive client.

Depending upon the information that a **SHOW** command displays, there might be instances where the information is changing or cases where it might cause the application (client, server, or storage agent) to stop. The **SHOW** commands should only be used at the recommendation of IBM Software Support. The **SHOW** commands that are included here are a portion of the available **SHOW** commands.

Table 16. Server or storage agent **SHOW** commands

SHOW Command	Description	Recommendation
AGGREGATE	Displays information about an aggregate object in the server storage hierarchy. The syntax is SHOW AGGRegate aggrID-high aggrID-low . <i>aggrID-high</i> and <i>aggrID-low</i> are the high-order and low-order 32-bit words of the 64-bit aggregate ID of the aggregate that is being queried.	Issue this command to determine the existence and logical files stored in an aggregate object in the server's storage hierarchy. The offset, length, and active state of backup files is displayed for files within the aggregate. You might issue this command if you are having trouble restoring or retrieving files, expiring or moving data, backing up primary storage pools, copying active data to active data pools, or auditing volumes.

Table 16. Server or storage agent *SHOW* commands (continued)

SHOW Command	Description	Recommendation
ASQUEUED	Displays the mount point queue. The syntax is SHOW ASQueued .	In order to use a drive, a client session or server process must first obtain a mount point. The mount point management on the server allows for queuing waiters for mount points if more mount points are needed than are available. This command is useful for determining the state of a mount point request, especially if a session or process appears to be stopped and waiting for a mount point.
ASVOL	Displays assigned volumes. The syntax is SHOW ASVo1 .	As sequential media volumes are assigned for use by a session or a process, they are tracked in an in-memory list. You can view this list to determine the state of in-use volumes, as well as stoppages or deadlock situations where a session or process appears to be stuck waiting for a volume or holding a volume and waiting for something else.
BFOBJECT	Displays the following information in the server storage hierarchy data: <ul style="list-style-type: none"> • The active/inactive state of logical files within an aggregate • The offset/length of logical files within an aggregate • The active state or owner bitfile ID of logical files within an aggregate • The link bitfile ID if the deduplicated extent is linked to another extent The syntax is SHOW BFObject .	This command helps you determine the existence and attributes of a bitfile object in the server's storage hierarchy. You might issue this command if you are having trouble restoring, retrieving, expiring, or auditing the object.
BUFSTATS	Displays usage statistics for the database buffer pool. The BUFSTATS command shows the cache hit percentage of the buffer pool, which is suggested to be above 98%. The syntax is SHOW BUFStats .	Issue this command to determine if the configured database buffer pool size is large enough.

Table 16. Server or storage agent SHOW commands (continued)

SHOW Command	Description	Recommendation
BUFVARS	Displays database buffer pool global attributes. The syntax is SHOW BUFVars . Important: Ensure that the IMEXP trace class is active.	Issue this command to determine if the configured database buffer pool size is large enough. This might also be useful for diagnosing cases where the server is stopped, or when the server runs out of recovery log space. The database buffer pool performance and characteristics can influence the server running out of recovery log space because the ability to write (flush) the changed pages to the database volumes can impact the ability of the recovery log to manage its space.
CMD DEDUPDELETEINFO	Displays the status of background deletion threads for dereferenced deduplicated objects.	Issue this command to check the status of the background deletion process for deduplicated objects. When a file is deleted or moved out of a deduplicated storage pool, the extents are queued to a background processor for attempted removal from the storage pool. This command is useful for checking the backlog of queued extents and the status of each deletion thread.
CONFIGURATION	The CONFIGURATION command is a summary SHOW command that actually issues many different show commands and queries. The syntax is SHOW CONFIGuration .	Issue this command to provide general configuration and other information about the server to IBM service.
DB2CONNECTIONS	The DB2CONNECTIONS command shows the defined DB2 connections from the various connection pools. This command does not require any additional parameters. The syntax is SHOW DB2COnnections .	Issue this command to show how many DB2 connections are defined, in-use, and free in total and within a given pool.

Table 16. Server or storage agent *SHOW* commands (continued)

SHOW Command	Description	Recommendation
DB2TABLES	The DB2TABLES command shows the registered tables and their column attributes. This command does not require any additional parameters. The syntax is SHOW DB2TABLEs .	Issue this command to show the registered tables and their column attributes.
DBTXNTABLE	Displays information about transactions that are performing database operations. The syntax is SHOW DBTXNTable .	Issue this command to display the following information: <ul style="list-style-type: none"> • Database tables that are open (in-use) • Recovery-log usage information such as the number of log records written and the recovery log space used • The first, last, and next recovery log records that were written • Whether or not a transaction is valid or being rolled back
DBVARS	Displays database global attributes. The syntax is SHOW DBVars .	Issue this command to view the current state and attributes of the server database.
DEDUPOBJECT	Displays data deduplication information for files. When you issue this command, you must specify the objectID parameter. Issue the SHOW VERSION command to determine the value of this parameter. The syntax is SHOW DEDUPObject .	Issue this command to display data deduplication information, such as: <ul style="list-style-type: none"> • The bit file ID for each extent • The owning bit file ID • The offset and length of the owning bit file • The digest type and value of the data deduplication object

Table 16. Server or storage agent SHOW commands (continued)

SHOW Command	Description	Recommendation
DEVCLASS	Displays information about device classes. The syntax for this command is SHOW DEVClass .	Issue this command to display the states of allocated drives, device class attributes, and other information. This command is often used to diagnose problems with devices or locks up waiting for a drive, library, or volume. The command SHOW LIBRARY also gives good complementary information about drives and libraries.
GROUPLEADERS	Displays all backup group leaders for an object in the server inventory. The syntax is SHOW GROUPLeaders objID-high objID-low . <i>objID-high</i> and <i>objID-low</i> are the high-order and low-order 32-bit words of the 64-bit object ID of the object being queried. The high-order word is optional; if not specified, a value of zero is assumed. The object must be a backup object.	Issue this command to determine the backup group relationships of an object in the server's inventory. You might issue this command if you are having trouble restoring, retrieving, expiring, or auditing the object.
GROUPMEMBERS	Displays all backup group members for an object in the server inventory. The syntax is SHOW GROUPMembers objID-high objID-low . <i>objID-high</i> and <i>objID-low</i> are the high-order and low-order 32-bit words of the 64-bit object ID of the object being queried. The high-order word is optional; if not specified, a value of zero is assumed. The object must be a backup object.	Issue this command to determine the backup group relationships of an object in the server's inventory. You might issue this command if you are having trouble restoring, retrieving, expiring, or auditing the object.

Table 16. Server or storage agent SHOW commands (continued)

SHOW Command	Description	Recommendation
INVOBJECT	Displays information about an inventory object in the server. The syntax is SHOW INVOBJECT <i>objID-high objID-low. objID-high</i> and <i>objID-low</i> are the high-order and low-order 32-bit words of the 64-bit object ID of the object being queried. The high-order word is optional; if not specified, a value of zero is assumed. The object can be a backup object, an archive object, a space-managed object, and so on.	Issue this command to determine the existence and attributes of an object in the server inventory. You might issue this command if you are having trouble restoring, retrieving, expiring, or auditing the object. The INVOBJECT command reports the following items: <ul style="list-style-type: none"> • New information for archive retention protected objects. • Whether the archive object is in deletion hold. • Whether the object uses event-based retention.
LIBINVENTORY	Displays the current state of the library inventory for the library specified. The syntax is SHOW LIBINVENTORY <i>libraryName</i> where <i>libraryName</i> is optional, and if left out, the command returns the inventory information for all libraries.	Issue this command if there is a problem with the library inventory information. The command displays current in-memory properties of the library inventory.
LIBRARY	Use the LIBRARY command to display the current state of the specified library and all of its drives. The syntax is SHOW LIBRARY <i>libraryName</i> where <i>libraryName</i> is optional. If left out, the command returns information for all the libraries.	This command is useful to gather a quick view of all in-memory information about a library and its drives. This output can be gathered for any problem related to libraries or drives (e.g. mounting problems).

Table 16. Server or storage agent *SHOW* commands (continued)

SHOW Command	Description	Recommendation
LOCK	Displays lock holders and waiters. The syntax is SHOW LOCK .	The server and storage agent use locks as a mechanism to serialize access and updates to information and other constructs. This information is used to diagnose stoppages or other resource contention issues.
LOGPINNED	Evaluates and determines whether or not the server recovery log is pinned. A pinned recovery log might cause the recovery log to run out of space and possibly cause the server to stop. The syntax is SHOW LOGPInned . To recover from a pinned recovery log, issue ' SHOW LOGPInned Cancel ' to cause the server to cancel or terminate the session, transaction, or process. Under some conditions, the pinning session or transaction might not terminate after issuing the CANCEL parameter.	This SHOW command interrogates a number of server control structures and correlates the data to determine if a session, transaction, or process is pinning the recovery log. If it determines that something is pinning the recovery log, this information is reported.
LOGVARS	Displays recovery log global attributes. The syntax is SHOW LOGVars .	Issue this command to determine the state of the recovery log.

Table 16. Server or storage agent SHOW commands (continued)

SHOW Command	Description	Recommendation
MEMTREND	<p>The MEMTREND command reports the memory used by the server, in megabytes, recorded at hourly intervals for the last 50 hours (this is a constant in the server code and is not user-configurable). The command also displays a histogram to help visualize the usage trend. The syntax is SHOW MEMTREnd.</p>	<p>Issue this command to determine if the server has a memory leak. If the memory usage is constantly increasing, this might indicate a leak. Note that for the measurements to be valid, the measurement period (the last 50 hours) should represent normal, steadystate server activity. The reported usage represents the amount of memory that internal server routines request from the pseudo-kernel memory routines. It does NOT represent the total amount of memory that the server is using. Nevertheless, it is still useful in determining the server's memory usage trend.</p>
MP	<p>Displays mount points. The syntax is SHOW MP.</p>	<p>Issue this command to determine which volume is in-use by a given mount point and other attributes for the assigned mount points. SHOW LIBRARY and SHOW DEVCLASS have useful complementary information with this command to display the current state of drives and current devclass mount point counts.</p>

Table 16. Server or storage agent *SHOW* commands (continued)

SHOW Command	Description	Recommendation
NASDEV	Displays the SCSI devices attached to a network-attached storage (NAS) file server associated with a NAS datamover definition. The syntax is SHOW NASDev.	Create a Network Data Management Protocol (NDMP) connection to the specified NAS file server and display the attached SCSI devices on the file server. This command only requires a NAS node and datamover definition.
NASFS	Displays the file systems on a NAS file server associated with a NAS datamover definition. The syntax is SHOW NASFs.	Create an NDMP connection to the specified NAS file server and display the file systems defined on the file server. Any file systems displayed by this command might be backed up by IBM Tivoli Storage Manager. This command requires only a NAS node and datamover definition.
NASINFORMATION	Displays configuration information about the NAS file server associated with a NAS datamover definition. The syntax is SHOW NASInformation.	Create an NDMP connection to the specified NAS file server and display general configuration information retrieved from the file server. This command is useful for identifying basic communication problems with NAS file servers such as authentication errors. This command only requires a NAS node and datamover definition.
NASWORKLOAD	Displays the workload of NAS filers that are used for all Tivoli Storage Manager operations. The syntax is SHOW NASWorkload.	Issue this command to determine the workload of backend data movement, as well as backup and restore operations.
REPLICATION	Displays all known replication servers and their globally unique identifier (GUID) as well as all running replication processes. The processes might include the individual statistics of each file space and the status of each replication session.	Issue this command if replication does not appear to be progressing or if replication is not working correctly.

Table 16. Server or storage agent *SHOW* commands (continued)

SHOW Command	Description	Recommendation
RESQUEUE	Displays the resource queue. The syntax is SHOW RESQueue .	Use the resource queue to monitor common resources on the server. If a resource appears to be stopped or holding a resource for an unreasonable amount of time, the resource monitoring algorithms for the server takes action and cancels or terminates the resource user. Typically, this is used to display information about transactions, locks, and other resources used by a storage agent on the database server that it is configured to use.
SESSIONS	Displays information about sessions connected to the server or storage agent. The syntax is SHOW SESSIONS .	Issue this command to diagnose stoppages or other general session problems while a session is still connected to the server. This is also useful in cases where a session is canceled or terminated and still appears in QUERY SESSION .

Table 16. Server or storage agent *SHOW* commands (continued)

SHOW Command	Description	Recommendation
SLOTS	Displays the current state of the specified library's slot information (for example: which volumes are in the library and in which slots). The syntax is SHOW SLOTS <i>libraryName</i> .	<p>The information displayed is what is saved directly from the library hardware to in-memory values and can be used to determine if this information is out-of-sync, incorrect, or to determine if the values returned from the library hardware itself are incorrect.</p> <p>Alternatively, issue this command to determine the drive element numbers for a SCSI library if QUERY SAN is unavailable for a particular library (for example: 3570 library).</p>
SSPOOL	Displays information storage pools. The syntax is SHOW SSPool1 .	Issue this command to display the states and attributes of defined storage pools.
THREADS	<p>Displays information about all threads known to the server. The syntax is SHOW THReads.</p> <p>Important: On some platforms (as an example: HP), the information reported is obtained without serialization. On a busy system, the information might be inconsistent, multiple threads might report holding the same mutex, or a thread might report that it is waiting on a mutex held by another thread that does not claim to hold it.</p>	<p>The server displays information about each thread, typically including the Tivoli Storage Manager thread ID, the system thread id, the thread name, mutexes it holds (if any), and mutex or condition it is awaiting (if any). This command is platform-specific, so each platform might have slightly different information. You might want to issue this command if the server or a particular server process appears to be stopped, in order to see if there are threads waiting for resources held by another thread.</p>

Table 16. Server or storage agent SHOW commands (continued)

SHOW Command	Description	Recommendation
TOCSETS	<p>Displays all Table Of Contents (TOC) sets known to the server. The syntax is SHOW TOCsets DELETE=setNum TOUCH=setNum. The DELETE parameter causes the specified TOC set number to be deleted. The TOUCH parameter updates the last used date of the specified TOC set number. A TOC set is retained for the TOC retention period following the last used date (see SET TOCRETENTION command).</p>	<p>A TOC set is used during file-level NDMP operations. During an NDMP backup with the TOC=YES parameter, a TOC is built in the server database. During a restore, one or more TOCs might be loaded into the server database in order to provide file and directory names to the client GUI. This command displays the status of the TOC set (e.g. building or loading) and how much temporary database space is in use for each TOC set. You might issue this command if you are having trouble doing an NDMP backup with the TOC=YES parameter, or having trouble restoring files from an NDMP backup, or if TOC sets are being retained in the server database too long or not long enough.</p>
TOCVARS	<p>Displays information about the TOC component of the server. The syntax is SHOW TOCvars.</p>	<p>Issue this command to determine the status of the TOC component. You might issue this command if you are having trouble doing an NDMP backup with the TOC=YES parameter, or having trouble restoring files from an NDMP backup.</p>

Table 16. Server or storage agent *SHOW* commands (continued)

SHOW Command	Description	Recommendation
TXNTABLE	Displays information about transactions that are on the in-use list on the server. The syntax is SHOW TXNTable .	The transactions that are mined by this command are used by server processes, sessions, or other operations to read information from the database, make updates to the database (such as insert, update, or delete information), or to manage locks. This information is useful for diagnosing stoppages or other transaction-related failures while the transaction is still open on the server.
VALIDATE LANFREE	Validates whether the definitions are in place on the server in order for a given client to perform LAN-free data movement operations. In cases where these definitions are not present or are incorrect, it might be difficult to determine if the LAN-free environment is configured correctly. The syntax is VALIDATE LANFREE <i>nodeName storageAgent</i> . Note: The VALIDATE LANFREE command replaced the SHOW LANFREE command.	This command evaluates all possible destination storage pools for this client node and reports whether or not the storage pool is capable of LAN-free data movement operations.
VERSIONS	Issue the SHOW VERSIONS command to retrieve an objectID . The objectID is necessary to issue the SHOW DEDUOBJECT command. The syntax is SHOW Versions .	Issue this command to display object IDs.
VOLINUSE	Displays whether the volume specified is currently in the server's in-use list. The VOLINUSE command displays additional information that might be helpful, including whether the volume is currently pending removal from the in-use list. The syntax is SHOW VOLINUSE <i>volumeName</i> . If the volume must be removed from the in-use list, you can specify the following additional parameter to remove the volume from the list: SHOW VOLINUSE <i>volumeName</i> REMOVE=YES .	Issue this command to determine whether a volume is on the in-use list and, if necessary, to remove it from that list. Operations that are associated with this volume might fail if the volume is removed from the in-use list.

Enabling a trace for the Tivoli Storage Manager device driver

Tracing is available for the IBM Tivoli Storage Manager device driver. The Tivoli Storage Manager device driver can be traced from the server console, an administrative client, or from a shell running on the system where the device driver is installed.

The tracing instructions apply to the Tivoli Storage Manager device driver on all platforms where the device driver is supported. For devices that use device drivers other than the Tivoli Storage Manager device driver, the ability to trace and instructions on how to trace those device drivers is provided by the device vendor.

Related reference:

“Tracing from the server console”

“Tracing data from a command shell for AIX, Oracle Solaris, and Windows” on page 159

Tracing from the server console

To trace the driver from the server, you must first issue the proper commands.

Issue the **TRACE ENABLE** and **TRACE BEGIN** commands to trace the driver from the server.

The Tivoli Storage Manager device driver actually consists of three drivers: one for library-autochanger devices, one for tape devices, and one for optical drives. You might choose which one you want to trace. The following syntax is for the command:

```
DDTRACE START [ LIBRARYDD | TAPEDD | OPTICALDD ]
[flags=EE |, FULL |, SYSLOG | BASE ]
DDTRACE GET [ LIBRARYDD | TAPEDD | OPTICALDD ]
DDTRACE END [ LIBRARYDD | TAPEDD | OPTICALDD ]
```

The following options are available:

START Turns on tracing and writes the trace to a memory buffer based on the default or specified **FLAGS** option.

GET Writes the memory buffer to the same file that was specified with the server **TRACE BEGIN** command.

END Stops writing trace to the memory buffer but does not wipe out the contents of the buffer, so you might run **END** before running **GET**.

LIBRARYDD

Traces the device driver that controls library-autochangers.

TAPEDD Traces the device driver that controls tape drives.

OPTICALDD

Traces the device driver that controls optical drives.

For the options listed above, you might specify any one device driver or the library device driver, and one of the other two. These are space delimited. For example:

DDTRACE START TAPEDD - Starts tracing the device driver that controls tape drives.

DDTRACE START OPTICALDD Starts tracing the device driver that controls optical drives.

DDTRACE START LIBRARYDD Starts tracing the library-autochanger.
DDTRACE START LIBRARYDD TAPEDD Traces both the library and the tape drives.
DDTRACE START LIBRARYDD OPTICALDD Traces both the library and the optical drives.

Whichever of these you use, specify the same ones for all commands in the start-get-end series.

The **FLAGS** parameter is optional and usually not required. The following values are for the **FLAGS** parameter:

EE Traces all device driver routine entries and exits.

FULL Turns on more debug tracing and provides more detail. Because the memory buffer size is fixed, however, fewer events are traced. Does not trace routine entry and exit points.

SYSLOG

On some platforms, SYSLOG directs the trace statements to be written to the system log in addition to the memory buffer. This offering is most useful in debugging kernel stoppages or in when the trace wraps in the memory buffer.

BASE BASE is the default and cannot be specified with any other flags. It is only used to turn off the EE, FULL, and SYSLOG flags without turning off trace.

Tracing data from a command shell for AIX, Oracle Solaris, and Windows

AIX

Solaris

Windows

The stand-alone utility, `ddtrace`, exactly mimics the **DDTRACE** server commands.

The stand-alone `ddtrace` utility is installed in the `devices` directory, which is the same directory as the `mttest`, `lbtest`, and `optest` utilities. Its syntax and options are identical to the **DDTRACE** server command. For example:

```
$ ddtrace start librarydd tapedd flags=EE - Start tracing both the library and tape drivers, and get additional entry/exit trace.
```

```
$ ddtrace get librarydd tapedd - Get the trace from memory and write it to the file ddtrace.out.
```

```
$ ddtrace end librarydd tapedd - Stop tracing to memory.
```

The main use of this stand-alone utility is primarily for cases when the driver needs to be traced during the IBM Tivoli Storage Manager server initialization. The `ddtrace` utility writes the memory buffer to the “`ddtrace.out`” file in the current directory. If the file exists, it appends to the file and does not overwrite it.

Tracing to detect a code page conversion failure

The IBM Tivoli Storage Manager server uses operating system functions to convert between Unicode and the server code page. If the system is not set up correctly, the conversion fails.

AIX **Linux** **Solaris** **Windows** The conversion failure is most likely to occur when you are using the Administration Center to access the server.

Perform the following steps to attain more information on the failure:

1. Begin tracing the UNICODE trace class.
2. Repeat the action that caused the error message to occur.
3. Check the server README file for any platform-specific requirements for language installation.
4. Ensure that the locales indicated by the problem code pages are installed and any requirements listed in the README file are installed.

Tracing data for the client

You can enable tracing on the client or client application programming interface (API) by altering the client options file.

Perform the following steps to enable tracing on the client or client API:

1. Determine the trace classes to enable from the following table:

Trace Class Name	Description	When to use	Additional Notes
SERVICE	Display general processing information for the client.	Useful in many cases. Generally recommended for protocol violations, transaction processing errors, or in cases where the client is stopped and not responding.	
VERBINFO	Collect information regarding the client-server protocol used by IBM Tivoli Storage Manager.	To debug protocol violations, transaction processing errors, or in cases where the client is stopped and not responding.	
VERBDETAIL	Detailed information regarding the client-server protocol used by Tivoli Storage Manager. This displays internal memory buffers containing the verbs sent and received by the client.	To debug session data corruption issues that might be caused by the network.	This generates a large amount of output.

2. Enable the trace by adding the following text to the client options file:
`traceflag <trace class name>.`

Attention: *<trace class name>* might be a comma-delimited list of trace classes. For example, this text could be entered as `traceflag service,verbinfo,verbdetail`.

3. Configure trace to begin and issue the trace messages to a file by adding the following text to the client options file: `tracefile <file name>`.
4. Perform the operation that is causing the problem.

Tip: Tracing might also be configured and started by invoking the client from a command prompt and specifying the flags above. For example, `dsm -traceflags=service -tracefile=file.out`.

Client and Journal Daemon traceflags

To run journal-based backup, you must use the Journal Daemon process. This process is used to track file system changes and maintain change journal databases.

The Journal Daemon uses the same tracing mechanism as the client, but the trace settings are specified in the journal configuration file (`tsmjbbd.ini`) as follows:

```
[JournalSettings]
TraceFlags=all_jbb
;
; the following two settings allow tracefile segmentation
;
TraceMax=100
TraceSegMax=1
tracefile=tracefiles\trace.out
```

Journal Daemon specific trace settings:

- BTREEDB - low-level BTREE database base class
- CACHEDB - disk cache backup and Windows 2003 exclude cache processing
- DBPERF - low-level database operation performance
- DBSTATS - performance tracking of database query, insert/update, delete, and tree walk operations
- FILEOPS - internal database activity
- JBBCOMM - listening thread
- JBBDAEMON - process manager
- JBBFILEMON - file system monitor
- JBBDBACCESS - database controller thread
- JBBDBINFO - low-level database access
- JBBNPCOMM - named pipe communications
- JBBSERVICE - Windows platform-specific service tracing
- JBBVERBINFO - detailed verb information
- ALL_JBB - aggregate traceflag that includes all of the above settings

Trace Settings for the backup-archive client specified in `dsm.opt`:

- JOURNAL - journal based backup tracing

Client trace classes

The client provides individual and aggregate trace classes. Aggregate trace classes are a shortcut for enabling many related trace classes by simply specifying the aggregate trace class name. For the documented trace classes, there might be references to trace classes that are enabled as part of an aggregate trace class, but are not explicitly discussed on their own.

The documented trace classes in table Table 17 are those that are most typically requested or used for diagnosing problems. The trace class name should be used with the TRACEFLAG options in the dsm.opt file.

Table 17. Trace classes

Trace Class	Description	Recommendation
ALL_BACK	Displays general backup processing information for the client. Aggregate of TXN, INCR, POLICY, and PFM trace classes and implicitly included in the SERVICE trace class.	Use this trace class for problems related to selective or incremental backups.
ALL_FILE	Displays general backup processing information for the client. Aggregate of DIROPS, FILEOPS, and FIOATTRIBS trace classes and implicitly included in the SERVICE trace class.	Use this trace class for problems related to reading and writing data and obtaining file attribute information.
ALL_IMAGE	Displays image processing information for the client. Aggregate of several image-related trace classes and implicitly included in the SERVICE trace class.	Use this trace class for problems related to all aspects of volume image backup and restore operations.
ALL_JBB	Displays journal-based backup processing information for the client. Aggregate of several journal-based backup-related trace classes and implicitly included in the SERVICE trace class.	Use this trace class for problems related to all aspects of journal-based backups.
ALL_NAS	Displays NDMP processing information for the client. Aggregate of several NDMP-related trace classes and implicitly included in the SERVICE trace class.	Use this trace class for problems related to all aspects of NDMP backup and restore operations.

Table 17. Trace classes (continued)

Trace Class	Description	Recommendation
ALL_SESS	Displays all session and verb information sent between the client and the server. Aggregate of SESSION, VERBINFO, SESSVERB, VERBADMIN, and VERBDETAIL trace classes. All of the trace classes in this aggregate are implicitly included in the SERVICE trace class, except VERBDETAIL	Use this trace class for problems related to the client and server session, such as communication timeouts, protocol violations, and instances where the client appears to be stopped and waiting for the server, or vice versa.
ALL_SNAPSHOT	Displays information relating to volume snapshot operations. Aggregate of several volume snapshot-related trace classes and implicitly included in the SERVICE trace class.	Use this trace class to determine problems related to volume snapshots that are used in online image backup and open file-support operations.
ALL_WAS	Displays Web Application Server (WAS) processing information for the client. Aggregate of several WAS-related trace classes and implicitly included in the SERVICE trace class.	Use this trace class for problems related to all aspects of WAS backup and restore operations.
AUDIT	Displays auditing information for backup and restore processing. Part of the SERVICE trace aggregate.	Use this trace class to keep record of files processed, committed, and restored in a file.
CLIENTTYPE	Displays client type on each trace output line.	Use this trace class for tracing situations when more than one client component is involved, such as the client acceptor and the file system agent.
COMPRESS	Displays compression information. Part of the SERVICE trace aggregate.	Use this trace class to determine how much data is compressed on a per-file basis.
DELTA	Displays adaptive subfile backup processing information. Part of the SERVICE trace aggregate.	Use this trace class to determine errors in adaptive subfile backup and restore operations.
DIROPS	Displays directory read and write operations. Part of the SERVICE and ALL_FILE trace aggregates.	Use this trace class when problems occur in a read or write directory.

Table 17. Trace classes (continued)

Trace Class	Description	Recommendation
DOMAIN	Displays incremental domain processing information. Part of the SERVICE trace aggregates.	Use this trace class for determining how DOMAIN statements are resolved during backup processing, such as problems in resolving the ALL-LOCAL domain.
ENCRYPT	Displays data encryption information. Part of the SERVICE trace aggregate.	Use this trace class to determine if a file is included for encryption processing.
ERROR	Displays operating system specific error information. Part of the SERVICE trace aggregate.	Use this trace class to determine error codes generated by the operating system.
FILEOPS	Displays file read and write operations. Part of the SERVICE and ALL_FILE trace aggregates.	Use this trace class when problems occur in a file open, read, write, or close operation.
FIOATTRIBS	Displays comparisons of file attributes between the local client version and the active version on the Tivoli Storage Manager server. Part of the SERVICE, ALL_BACK, and ALL_FILE trace aggregates.	Use this trace class in determining why a file was backed up during an incremental backup.
INCR	Displays incremental list processing comparisons between the client and server. Part of the SERVICE and ALL_BACK trace aggregates.	Use this trace class for determining if files are candidates for incremental backup, especially in conjunction with the FIOATTRIBS trace class.
INCLEXCL	Displays include-exclude status for the object being processed. This flag is also used for the Preview function.	Use this trace class to determine which object (usually file or directory) is included or excluded during backup-archive/preview.
MEMORY	Displays memory allocation and free requests. This trace class writes a large amount of information into the trace file and is not included in any aggregate classes.	Use this trace class to determine memory leaks, memory spikes, and other memory-related problems.
OPTIONS	Displays current processing options. Part of the SERVICE trace aggregate.	Use this trace class to determine which options are in effect for the current session, and for problems in accepting processing options from server client-options sets.

Table 17. Trace classes (continued)

Trace Class	Description	Recommendation
PASSWORD	Displays password file-access information (does not show passwords). Part of the SERVICE trace aggregate.	Use this trace class to determine problems with reading the Tivoli Storage Manager server passwords from local storage, for example, PASSWORDACCESS=GENERATE errors.
PID	Displays process ID on each trace statement. Part of the SERVICE trace aggregate.	Use this trace class to diagnose problems that might involve multiple processes.
POLICY	Displays policy information available to the backup-archive client. Part of the SERVICE and ALL_BACK trace aggregates.	Use this trace class to see which policies are available during a backup or archive operation.
SCHEDULER	Displays general processing information for the scheduler. An aggregate that includes most of the client trace classes listed in this table. Aggregate of all trace classes except MEMORY, THREAD_STATUS, and *DETAIL classes.	Useful in many cases. This trace class is generally recommended for diagnosing scheduler problems when the nature of the problem is unknown. If the SCHEDULER trace flag is used, it generally is not necessary to specify any other trace flags because it already includes most of the basic trace classes.
SERVICE	Displays general processing information for the client. An aggregate that includes most of the client trace classes listed in this table. Aggregate of all trace classes except MEMORY and *DETAIL classes. The SERVICE trace flag can generate a substantial amount of information. Consider using the TRACEMAX option in conjunction with the SERVICE trace flag.	Useful in many cases. This trace class is generally recommended when the nature of the problem is unknown. If the SERVICE trace flag is used, it will generally not be necessary to specify any other trace flags because it already includes most of the basic trace classes.
SESSION	Displays minimal session information between the client and the server. Part of the SERVICE and ALL_SESS trace aggregates.	Use this trace class to give session context to general processing errors, or in conjunction with one of the VERB* trace classes, to determine session problems such as session timeouts and protocol violations.

Table 17. Trace classes (continued)

Trace Class	Description	Recommendation
SESSVERB	Displays additional session information between the client and the server. Part of the SERVICE and ALL_SESS trace aggregates.	Use this trace class to give session context to general processing errors, or in conjunction with one of the VERB* trace classes, to determine session problems such as session timeouts and protocol violations.
STATS	Displays final processing statistics in the trace file. Part of the SERVICE trace aggregate.	Use this trace class for collecting final processing statistics into a file.
THREAD_STATUS	Displays thread status. Part of the SERVICE trace aggregate.	Use this trace class when diagnosing problems related to threading.
TXN	Displays transaction processing information. Part of the SERVICE and ALL_BACK trace aggregates.	Use this trace class when diagnosing problems related to transaction processing problems on the server, and for such problems as transaction stops and retries.
VERBDETAIL	Displays detailed verb information pertinent to client-server sessions. Part of the ALL_SESS trace aggregates.	Use this trace class to determine the contents of verbs sent between the client and server.
VERBINFO	Displays verb information pertinent to client-server sessions. Part of the SERVICE and ALL_SESS trace aggregates.	Use this trace class in conjunction with the SESSION traceflag to give session context to general processing errors or to determine session problems like session timeouts and protocol violations.
WIN2K	Displays Windows system object or system state processing. Part of the SERVICE trace aggregates. Only valid on the Windows backup-archive client.	Use this trace class to determine errors with backup or restore of the system state information.

Enabling a backup-archive client trace

There are two methods of tracing that are available for the backup-archive client.

The first method is to configure trace parameters prior to starting the backup-archive client. The second is to enable tracing while the client is running. Choose which method of tracing to enable.

Enabling a client trace using the command line

You can trace the available backup-archive client by enabling client trace on the command line.

Perform the following steps to enable client tracing on the command line:

1. Determine the trace classes to enable.
2. Choose which trace classes to enable by adding the following text to the `dsm.opt` client options file: `traceflags <trace class name>`
3. Use a minus sign (-) in front of a trace class to subtract the class from tracing. Make sure that the subtracted trace classes are placed at the end of the trace class list. For example, if you want to collect a SERVICE trace without the SESSION or SESSVERB classes, then specify the following text:

Correct: `traceflags service,-session,-sessverb`

Incorrect: `traceflags -session,-sessverb,service`

Attention: `<trace class name>` might be a comma-delimited list of trace classes. For example, this text could be entered as `traceflags service,verbdetail`

4. Choose the location of the trace messages output by adding the following text to the client options file: `tracefile <file name>`.

The *tracefile* name should be fully-qualified, for example:

 `tracefile c:\service\trace.out`

    `tracefile /home/spike/trace.out`

`tracefile trace.txt` (Macintosh)

5. Set a maximum size for the trace file between 1 and 4,294,967,295 MB by specifying the following variable in the client options file: `tracemax <size in mb>`

If a maximum value is specified, the client starts writing information from the beginning of the trace file (that is, wrapping) when the trace reaches its maximum size. This information can be useful if you are trying to capture an event that happens at the end of a long-running process. For example, to specify a maximum trace file size of 10 MB: `tracemax 10` After a tracefile reaches the limit specified with `tracemax`, “Continued at beginning of file” is written to the end of the trace file and tracing continues from the top of the file. The end of the tracefile is indicated with “END OF DATA.” You can locate the end of the trace by searching for this string. If you specify a TRACEMAX size of 1001 or higher and TRACESEGSIZE is not specified, then the trace file is automatically split into multiple segments of 1000 MB per segment (see TRACESEGSIZE discussion).

You can choose to let the client split the trace into smaller segments (between 1 and 1,000 MB per segment) by specifying the following variable in the client options file: `tracesegsize <trace segment size in MB>`

Splitting the trace into small segments allows you to more easily manage very large amounts of trace data, avoiding the problems associated with compressing very large files and eliminating the need to use a separate “file splitter” utility. For example, issue the following command to specify a trace segment size of 200 MB: `tracesegize 200`

A trace file segment name is specified with the `tracefile` option, plus an extension indicating the segment number. For example, if you specify `tracefile tsmtrace.out` and `tracesegsize 200`, then the trace will be segmented into multiple separate files of no more than 200 MB each, with file names `tsmtrace.out.1`, `tsmtrace.out.2`, and so on. Note that when specifying the segment size, do not use any comma separators:

Correct: tracemax 1000

Incorrect: tracemax 1,000

If you use the TRACESEGSIZE option, the trace file segments are named using the name specified in the option file with an additional extension using the segment number. For example, trace.out.1

6. Perform the operation that exhibits the problem.

Tracing might also be configured and started by invoking the client from a command prompt and specifying the previously-defined flags. For example:

```
dsmc -traceflags=service,verbdetail -tracefile=tsmtrace.out  
-tracemax=2500 -traceseysize=200
```

Related reference:

“Client trace classes” on page 162

Enabling a trace while the client is running

You can trace the available backup-archive client while the client is running.

- The backup-archive client must be installed to use dynamic tracing.
- The DSMTRACELISTEN YES option must be in effect when the client is started.
 - **AIX** **HP-UX** **Linux** **Solaris** This option is specified in the system options file (dsm.sys) in the stanza that the client uses. Users must be logged in as root in order to use dsmtrace.
 - **Windows** This option is specified in the client options file (usually dsm.opt). Users must be logged in as a member of the Administrators group.

When the client starts, it launches a separate “trace listener” thread. This thread “listens” on a named pipe, waiting to be contacted by the dsmtrace utility. In order to make the named pipe name unique, the client process ID (PID) is a part of the pipe name. When you use dsmtrace to configure tracing, it contacts the client through the named pipe on which the client is listening and passes to it the preferred trace configuration operation. The client then passes the results of the operation back to dsmtrace through another similarly-named output pipe. dsmtrace displays the results to the console. The client starts the trace listener thread only when client option DSMTRACELISTEN YES is in effect. If DSMTRACELISTEN NO is in effect, then the listener thread is not started and dynamic tracing is not available to that client. DSMTRACELISTEN NO is currently the default value.

The steps for gathering an IBM Tivoli Storage Manager client trace are as follows:

1. Stop the Tivoli Storage Manager backup-archive client.
2. Configure the client options file with the preferred trace options.
3. Restart the backup-archive client and reproduce the problem.
4. Stop the backup-archive client.
5. Remove the trace options from the backup-archive client options file.
6. Send the resulting trace file to IBM technical support for analysis.

You can also use the dsmtrace utility to start, stop, and configure client tracing dynamically without having to stop the client or modify the options file. Dynamic tracing is especially useful when you need to trace only the beginning of a long-running backup-archive client operation, or when you need to start tracing after the backup-archive client is running for some time.

The dsmtrace utility includes the following features:

- Identify running processes and their process PIDs
- Enable client tracing

- Disable client tracing
- Query client trace status

The following table summarizes the availability of this feature:

Table 18. Availability of the *dsmtrace* utility

Client Component	AIX, HP-UX, Linux, or Oracle Solaris Program Name	Windows Program Name
Backup-Archive Client (command line)	dsmc	dsmc.exe
Backup-Archive Client (GUI)	N/A	dsm.exe
Client acceptor	dsmcad	dsmcad.exe
Remote Client Agent	dsmagent	dsmagent.exe
Scheduler Service	N/A	dsmcsvc.exe
Journal Service	N/A	tsmjbbd.exe
Data Protection for Domino (command line)	domdsmc	domdsmc.exe
Data Protection for Domino (GUI)	N/A	domdsm.exe
Data Protection for Microsoft Exchange (command line)	N/A	tdpexcc.exe
Data Protection for Microsoft Exchange (GUI)	N/A	tdpexc.exe
Data Protection for Microsoft SQL Server (command line)	N/A	tdpsqlc.exe
Data Protection for Microsoft SQL Server (GUI)	N/A	tdpsql.exe

Note:

- The center column in Table 18 includes Macintosh OS X.
- Tracing for the Data Protection components is for the Tivoli Storage Manager application programming interface (API) only.
- The Tivoli Storage Manager API tracing is available with any multithreaded application that uses the Tivoli Storage Manager API. The executable file name is the name of the application program that loads the API.

The following example shows you how to enable client trace while the client is running:

1. Identify the process PID of the backup-archive client that you want to trace (make sure that DSMTRACELISTEN YES is in effect). Issue the following command to show all running instances of the client: `dsmtrace query pids`

Example output:

```
D:\tsm>dsmtrace query pids
```

```
IBM Tivoli Storage Manager
dsmtrace utility
dsmtrace Version 5, Release 3, Level 0.0
dsmtrace date/time: 10/24/2004 21:07:36
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.
```

PROCESS ID	PROCESS OWNER	DESCRIPTION	EXECUTABLE NAME
4020	andy	Backup-Archive Client (CLI)	dsmc.exe

D:\tsm>

Important: Linux The threading model for some versions of Linux is to run each thread as a separate process. This means that when you query process information, you might see several processes for each instance of the client. The process that you need to identify is the dsmc parent process. For example:

```
fvrtlinuxppc:/opt/tivoli/tsm/client/ba/bin # dsmtrace q p
```

```
IBM Tivoli Storage Manager
dsmtrace utility
dsmtrace Version 5, Release 3, Level 0.0
dsmtrace date/time: 10/24/04 08:07:37
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.
```

PROCESS ID	PROCESS OWNER	DESCRIPTION	EXECUTABLE NAME
28970	root	Backup-Archive Client (CLI)	dsmc
28969	root	Backup-Archive Client (CLI)	dsmc
28968	root	Backup-Archive Client (CLI)	dsmc
28967	root	Backup-Archive Client (CLI)	dsmc

```
fvrtlinuxppc:/opt/tivoli/tsm/client/ba/bin #
```

In such a situation, issue the **PS** command to identify the parent dsmc process:

```
linuxppc:~ # ps -ef | grep dsmc
```

```
root 28967 1151 0 Oct22 pts/16 00:00:00 dsmc
root 28968 28967 0 Oct22 pts/16 00:00:00 dsmc
root 28969 28968 0 Oct22 pts/16 00:00:00 dsmc
root 28970 28968 0 Oct22 pts/16 00:00:00 dsmc
root 24092 24076 0 08:15 pts/93 00:00:00 grep dsmc
```

```
linuxppc:~ #
```

Notice that the parent for processes 28969 and 28970 is 28968. The parent for 28968 is 28967. The parent for 28967 is 1151, but the 1151 process does not appear in this display output. Process 1151 is the process that launched dsmc. So, the correct parent process ID is 28967.

- Issue the following command to enable tracing on the client:

```
dsmtrace enable 4020 -traceflags=service -tracefile=d:\trace.txt
```

Example output:

```
C:\program files\tivoli\tsm\baclient>dsmtrace enable 4020 -traceflags=service
-tracefile=d:\trace.txt
```

```
IBM Tivoli Storage Manager
dsmtrace utility
dsmtrace Version 5, Release 3, Level 0.0
dsmtrace date/time: 10/24/2004 21:45:54
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.
```

ANS2805I Tracing has been enabled.

```
C:\program files\tivoli\tsm\baclient>
C:\program files\tivoli\tsm\baclient>
```

Important: When tracing an API application, the `-pipenameprefix` option must be included.

- AIX
HP-UX
Linux
Solaris
 Use prefix `/tmp/TsmTraceTargetAPI`

- **Windows** Use prefix `\\.pipe\TsmTraceTargetAPI`
3. After sufficient trace data is collected, disable the tracing by issuing the following command:


```
dsmtrace disable 4020
```

 Example output:


```
C:\program files\tivoli\tsm\baclient>dsmtrace disable 4020

IBM Tivoli Storage Manager
dsmtrace utility
  dsmtrace Version 5, Release 3, Level 0.0
  dsmtrace date/time: 10/24/2004 21:47:43
  (c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.

ANS2802I Tracing has been disabled.
```

Other examples of enabling client trace while the client is running are defined in the following list:

dsmtrace query pids

This command displays all running processes whose names are listed in the table in the Background section.

dsmtrace query pids -filter=*

This command displays all running processes.

dsmtrace query pids -filter=dsm*

This command displays all running processes whose name begins with "dsm"

dsmtrace query pids -filter=dsm?

This command displays all running processes whose name begins with "dsm" plus one other character.

dsmtrace enable 2132 -traceflags=service -tracefile=c:\trace.txt

This command turns on SERVICE tracing for process 2132. Trace output is written to file c:\trace.txt.

dsmtrace enable 2132 -traceflags=-extrc

This command turns off extrc tracing for process 2132 (presumably tracing is already running for this process).

dsmtrace enable 4978 -traceflags=fileops -tracefile=/tmp/dsmtrace.out -tracemax=1000 -tracesegsize=200

This command turns on FILEOPS tracing for process 4978. The trace is written to files /tmp/dsmtrace.out.1, /tmp/dsmtrace.out.2, and so on, with each file being no larger than 200 MB. After 1000 MB are written, tracing wraps back to /tmp/dsmtrace.out.1.

dsmtrace query trace 4978 -on

This command displays basic trace information and lists trace flags that are turned on for process 4978.

dsmtrace disable 4978

This command disables tracing for process 4978.

dsmtrace disable 364 -pipenameprefix=/tmp/TsmTraceTargetAPI

This command disables tracing for API application process 364.

Known trace problems and limitations

The known problems and limitations of trace processes are gathered to help you resolve problems that you might have encountered when running a trace process.

- If tracing is not currently active for a process and `dsmtrace` is used only with the `-TRACEFLAGS` option, e.g. `dsmtrace enable 2346 -traceflags=service`, then you still see the following message:

```
ANS2805I Tracing has been enabled.
```

In this case, the trace flags were enabled (turned on), but tracing is not actually active until a trace file is specified by using the `-TRACEFILE` option.

- Do not use the `dsmtrace enable` command to start tracing the application programming interface (API) for Data Protection applications if the Data Protection application is run in a manner that does not cause it to connect to the IBM Tivoli Storage Manager server. For example, The Data Protection for Lotus Domino command line interface has several such commands:

- `domdsmc help`
- `domdsmc set`
- `domdsmc query domino`
- `domdsmc query pendingdbs`
- `domdsmc query preferences`

If you use `dsmtrace` to enable tracing for such commands, the result can be a stoppage of the `dsmtrace` process and (AIX, HP-UX, Linux, and Oracle Solaris only) a residual named pipe in the `/tmp` directory.

- **Windows** You must be logged in with a local administrative account to use `dsmtrace`.
- You must be logged in as root to use `dsmtrace`. If a client process stops or is stopped, it might leave a named pipe (UNIX FIFO) in the `/tmp` directory. These FIFOs have names beginning with `TsmTrace` and they include a process ID (PID) number. If a client process stops or is stopped, and then a new client process is started whose PID happens to match that of the old residual FIFO, then the trace listener thread will probably not start. Any old FIFOs with process numbers that do not match those of running the Tivoli Storage Manager processes can be safely deleted. Do NOT delete the FIFO of a running process.
- The threading model for some versions of Linux is to run each thread as a separate process, meaning that when you query process information, you might see several processes for each instance of the client. The process you need to identify is the `dsmc` parent process.
- When multiple instances of the same program are running, you must identify the PID of the instance that you want to trace. In such a situation, using other clues such as process information from the operating system might be available to narrow down the preferred PID. For example, if you want to trace `dsmc` that is being run by user 'andy' and there are two instances of `dsmc` (one owned by user 'andy' and the other owned by user 'kevin'), you can use the process owner to identify which process to trace.
- If an options file contains a false option and the client does not start, you might see some named pipe errors in the `dsmerror.log` file. These error messages might be safely ignored. A fix for this problem is planned for a future release.

Trace options

Trace has several options that you can employ.

DSMTRACEListen

DSMTRACEListen No | Yes

- No** The client does not start the trace listener thread and dynamic tracing is not available. The default is No.
- Yes** The client starts the trace listener thread and dynamic tracing is available.

Windows The DSMTRACEListen option is specified in the client options file (usually dsm.opt).

HP-UX The DSMTRACEListen option is specified in the system options file (dsm.sys) in the stanza that the client uses and cannot be specified from the command line.

dsmtrace

dsmtrace enable <pid> <options>

Use this command to start or modify tracing for a process.

pid The process ID (PID) for the client. Use dsmtrace query pids or your operating system facilities to identify the correct PID.

options The client trace options.

dsmtrace disable <pid>[<options>]

Use this command to stop tracing for a process. The trace file closes and the trace flags, maximum trace size, maximum trace segment size, and trace file name are all cleared.

<pid> The PID for the client. Use **dsmtrace query pids** or your operating system facilities to identify the correct PID.

<options>
The client trace options.

dsmtrace help

This command displays basic syntax for dsmtrace.

dsmtrace query pids [-Filter=<spec>]

<spec> The client process name filter specification, which can include the wildcard characters "?" (match exactly one character) or "*" (match zero or more characters).

If no filter is specified, then the default behavior is to display process information for any running instances of the program names listed in the table in the Background section above.

Important: **AIX** **HP-UX** **Linux** **Solaris** When using the FILTER, put the * symbol before and after the search text. This adjustment is necessary because the executable file name often includes the path in front of it, and in some cases, the executable file name might have additional characters at the end of it. For example:

- /opt/tivoli/tsm/client/ba/bin/dsmc

- domdsmc_DominoUserID

Thus, instead of `-filter=dsmc` or `-filter=domdsmc`, use `-filter=*dsmc*` or `-filter=*domdsmc*`.

dsmtrace query trace *<pid>* [*<options>*] [*<displayType>*] [-ALI | -ON | -Off | -BASic]

<pid> The process ID (PID) for the client. Use `dsmtrace query pids` or your operating system facilities to identify the correct PID.

<options>
The client trace options.

<displayType>
The display type can be one of the following entries:

ALI Displays all trace flags and, for each flag, indicates whether it is turned on or off. The information shown with the `-BASIC` display type is also included.

ON Displays the names of the trace flags that are turned on. The information shown with the `-BASIC` display type is also included.

Off Displays the names of the trace flags that are turned off. The information shown with the `-BASIC` display type is also included.

BASic Displays the name of the trace file and the maximum trace and trace segment sizes. This display type also indicates whether tracing is enabled or disabled.

-PIPENameprefix

-PIPENameprefix=*<pipeNamePrefix>*

The `-PIPENameprefix` option must be used when tracing application programming interface (API) applications:

- **AIX** **HP-UX** **Linux** **Solaris** Use prefix `/tmp/TsmTraceTargetAPI`
- **Windows** Use prefix `\\.\pipe\TsmTraceTargetAPI`

-TRACEFile

-TRACEFile=*<traceFileName>*

The `-TRACEFile` option must specify a valid file name to which the trace is written. If tracing is already running, then this option has no effect.

-TRACEFlags

-TRACEFlags=*<traceFlags>*

Specify one or more trace flags. Typically, the trace flag `SERVICE` is used. Separate multiple trace flags with a comma. Trace flags can also be turned off by prefixing the flag name with a minus sign. When combining trace flags that you want to turn on, with trace flags that you want to turn off, put the flags that you want to turn off at the end of the list. For example, if you want to turn on `SERVICE` tracing except for `VERBDETAIL`, specify `-TRACEFLAGS=SERVICE,-VERBDETAIL`. If tracing is already running, then this option can be used to turn on additional trace flags or turn off trace flags.

-TRACEMax

-TRACEMax=<maximumTraceSize>

This option limits the maximum trace file length to the specified value (by default the trace file grows indefinitely). When the maximum length is reached, then the trace wraps back to the beginning of the file. Specify a value in MB between 1 and 4095. If tracing is already running, this option has no affect.

-TRACESegsize

-TRACESegsize=<maximumTraceSegmentSize>

This option is used when you anticipate a large trace file and you want the trace file to be written in smaller, more easily-managed segments. Each segment is no larger than the specified size. When this option is used, a segment number is appended to the trace file name for each segment. Specify a value in MB between 1 and 1000. If tracing is already running, this option has no affect.

Note:

- To turn tracing on for a process, you must use the -TRACEFLAGS and -TRACEFILE options (and -PIPENAMEPREFIX when tracing an API application).
- To modify trace flags for an existing process, use -TRACEFLAGS (and -PIPENAMEPREFIX when tracing an API application).
- If you need to modify the trace file name, maximum trace size, or maximum trace segment size, then you need to first disable tracing altogether (see the **dsmtrace disable** command).

Determining if data is encrypted or compressed during backup-archive through trace

You must perform several steps to determine whether the data during backup-archive is compressed or encrypted, or both.

1. Add the trace options listed to the client options file prior to backing up or archiving objects:

- TRACEFILE <trace file name>
- TRACEFLAGS api api_detail

2. Examine the trace file after the operation and locate a statement that looks similar to the following statement:

```
dsmSendObj ENTRY:... objNameP: <the file name>
```

This output is followed by the following trace message that indicates whether the object is compressed, encrypted, or both compressed and encrypted:

```
tsmEndSendObjEx: Total bytes send * *, encryptType is *** encryptAlg is ***  
compress is *, totalCompress is * * totalLFBytesSent * *
```

```
+-----+  
| encryptType/compress | 0 | 1 |  
+-----+  
| NO | not compressed, not encrypted | compressed, not encrypted |  
| CLIENTENCRKEY | not compressed, encrypted | compressed, encrypted |  
| USER | not compressed, encrypted | compressed, encrypted |  
+-----+
```

Alternatively, your application itself can determine encryption type/strength and compression of your data by using the **dsmEndSendObjEx** function call and the **dsmEndSendObjExOut_t** data structure.

```

/*-----+
| Type definition for dsmEndSendObjExOut_t
+-----*/
typedef struct dsmEndSendObjExOut_t
{
dsUint16_t      stVersion;          /* structure version */
dsStruct64_t    totalBytesSent;     /* total bytes read from app */
dsmBool_t       objCompressed;     /* was object compressed */
dsStruct64_t    totalCompressSize; /* total size after compress */
dsStruct64_t    totalLFBytesSent;  /* total bytes sent LAN Free */
dsUInt8_t       encryptionType;    /* type of encryption used */
}dsmEndSendObjExOut_t;

```

objCompressed - A flag that displays if the object was compressed.
encryptionType - A flag that displays the encryption type.

For example:

```

...
rc = dsmEndSendObjEx(&endSendObjExIn, &endSendObjExOut);
if (rc)
{
printf("*** dsmEndSendObjEx failed: ");
rcApiOut(dsmHandle, rc);
}
else
{
printf("Compression:      %s\n",
endSendObjExOut.objCompressed == bTrue ? "YES" : "NO");

printf("Encryption:      %s\n",
endSendObjExOut.encryptionType & DSM_ENCRYPT_CLIENTENCRKEY ?
"CLIENTENCRKEY" :
endSendObjExOut.encryptionType & DSM_ENCRYPT_USER ? "USER" : "NO");
printf("Encryption Strength: %s\n\n",
endSendObjExOut.encryptionType & DSM_ENCRYPT_AES_128BIT ? "AES_128BIT" :
endSendObjExOut.encryptionType & DSM_ENCRYPT_DES_56BIT ? "DES_56BIT" :
"NONE");
}
...

```

See the *API Function Calls* in *Using the Application Programming Interface* for more information.

Tracing data for the API

You can enable tracing for the application programming interface (API).

To enable tracing for the IBM Tivoli Storage Manager API, add the following lines to the `dsm.opt` file or another file designated as the client options file:

```

TRACEFILE trace file name
TRACEFLAGS trace flags

```

trace file name

The name of the file where you want to write the trace data.

trace flags

The list of trace flags to enable. Separate each trace flag by a space. The following trace flags are specific to the Tivoli Storage Manager API:

api Information about the API function calls

api_detail

Detailed information about the API function calls

You can also specify other Tivoli Storage Manager backup-archive client and Tivoli Storage Manager API trace flags. Refer to the backup-archive client documentation for a list of any available trace classes. For example:

- TRACEFILE /log/trace.out
- TRACEFLAGS api api_detail verbinfo verbdetail time stamp

Important: If you do not have write permission for the file pointed by the TRACEFILE option, dsmSetup or dsmInitEx/dsmInit fail with return code DSM_RC_CANNOT_OPEN_TRACEFILE (426).

To enable tracing for the multithreaded API after an application is started, use the dsmttrace utility. The dsmttrace utility lets you turn on tracing while the problem is occurring, without having trace constantly enabled. Refer to the *dsmttrace* section.

Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on an AIX or Linux system

AIX

Linux

By using Tivoli Monitoring for Tivoli Storage Manager, you can create and configure agent instances that monitor Tivoli Storage Manager servers. To activate tracing of monitoring agents for Tivoli Storage Manager servers running on AIX or Linux systems, stop all agent instances, modify the configuration files, and restart the agent instances.

Before activating tracing, you can also open the Tivoli Enterprise Portal Agent log workspace and view the agent activities. The Agent Log workspace contains information about any Tivoli Storage Manager server that has an agent instance configured to monitor it. The Agent Log attribute group lets you see output from the trace file without having to enable the trace file.

Complete the following steps to activate tracing:

1. Stop monitoring agent instances by completing one of the following steps:
 - Stop the monitoring agents using the CandleManage graphical user interface by issuing the following commands:
 - a. Go to the directory where the CandleManage program is located by issuing the following command:

```
cd /opt/tivoli/tsm/reporting/itm/bin
```
 - b. Run the CandleManage program by issuing this command:

```
./CandleManage &
```
 - c. In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is stopped. If it is not stopped, select the applicable agent instance, right-click it and select **Stop**.
 - Stop the monitoring agents from the command line, by issuing the following commands:
 - a.

```
/opt/tivoli/tsm/reporting/itm/bin  
./itmcmd agent -o instance_name stop sk
```
2. To ensure that all agents are stopped, complete the following steps:
 - a. Wait until the CandleManage graphical interface reports that the agent is stopped.
 - b. Verify whether the following process is running by issuing the following command:

```
ps -ef | grep -i SK
```

- c. If the process is running, stop the process by issuing the following command:

```
kill -9 process_ID
```

3. Locate the directory where the configuration files are stored by issuing the following command:

```
install_dir/itm/config
```

For example, the configuration files are typically in the default directory:
`/opt/tivoli/tsm/reporting/itm/config`

4. To activate monitoring agent tracing, ensure that the following value is set in the `sk_agentInstance.config` file:

```
KSK_TRACE='1'
```

You must also ensure that the following value is set in the `sk.ini` configuration file:

```
KSK_TRACE=1
```

5. If you are asked by an IBM Support representative to activate tracing for the API, ensure that the following value is set in the `sk_agentInstance.config` file:

```
KSK_APITRACE='1'
```

You must also ensure that the following value is set in the `sk.ini` configuration file:

```
KSK_APITRACE=1
```

6. Start the Tivoli Monitoring for Tivoli Storage Manager agent instances by completing one of the following steps:

- From the command line, issue the following commands:

```
cd /opt/tivoli/tsm/reporting/itm/tables  
../bin/itmcmd agent -o instance_name start sk
```

- From the CandleManage graphical user interface, select each monitoring agent, right-click it and select **Start**.

To review trace results, locate the log files in the `/install_dir/itm/logs/` directory.

The log file that contains the trace information has the following format:

`aaaapppptttt.log`, and the API trace has the following format:

`aaaappppttttDsmQuery.out`, where:

aaaa is the agent instance name

pppp is the server port number

tttt is the timestamp

For example:

`instancename15001111103143325000.log`, and `hostname1500DsmQuery.out`

Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on a Windows operating system

Windows

By using Tivoli Monitoring for Tivoli Storage Manager, you can create and configure agent instances that monitor Tivoli Storage Manager servers. To activate tracing of monitoring agents for Tivoli Storage Manager servers running on Windows operating systems, stop all agent instances, modify the configuration file, and restart the agent instances.

Before activating tracing, you can also open the Tivoli Enterprise Portal Agent log workspace and the Agent Log attribute group, and view the agent activities. The Agent Log workspace contains information about any Tivoli Storage Manager server that has an agent instance configured to monitor it.

Complete the following steps to activate tracing:

1. Stop the monitoring agent instances by completing the following steps:
 - a. On the Tivoli Monitoring server, click **Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services**.
 - b. Select each monitoring agent instance, right-click, and select **Stop**.
2. Locate the directory where the configuration file is stored:

```
install_dir\itm\tmaitm6
```

For example:

```
C:\IBM\itm\tmaitm6
```

3. To activate agent tracing, ensure that the following value is set in the `kskenv_agentInstance` file:

```
KSK_TRACE=1
```
4. The application programming interface (API) can also be traced, but is not required unless requested by a Tivoli support representative. To activate tracing for the API, ensure that the following value is set in the `kskenv_agentInstance` file:

```
KSK_APITRACE=1
```
5. Start the Tivoli Monitoring for Tivoli Storage Manager agent instances by completing the following steps:
 - a. On the Tivoli Monitoring server, click **Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services**.
 - b. Select each monitoring agent, right-click, and select **Start**.

The trace results are located in the same directory as the configuration file:

```
install_dir\itm\tmaitm6\logs
```

The API trace results are located in the following directory:

```
install_dir\itm\tmaitm6
```

For example:

```
C:\IBM\itm\tmaitm6\logs
```

```
C:\IBM\itm\tmaitm6
```

The log file that contains the trace information has the format `aaaappptttt.log`, and the API trace has the format `aaaapppttttDsmQuery.out`, where:

aaaa is the agent instance name
pppp is the server port number
tttt is the timestamp

For example:

instancename15001111103143325000.log, and hostname1500DsmQuery.out

Chapter 8. Resolving Tivoli Monitoring for Tivoli Storage Manager problems

You can use these troubleshooting topics to diagnose and resolve Tivoli Monitoring for Tivoli Storage Manager problems.

Resolving installation issues with Tivoli Monitoring for Tivoli Storage Manager

AIX

Linux

Windows

Tivoli Monitoring for Tivoli Storage Manager installs several components including IBM DB2, IBM Tivoli Monitoring, and the monitoring agent. Any errors during installation can be investigated.

Reviewing the logs.zip file to resolve installation failures

AIX

Linux

Windows

If the Tivoli Monitoring for Tivoli Storage Manager installation fails, you are prompted to review the files in the logs.zip compressed file. Understanding how to read the logs.zip file can help you to determine the cause of the failure and provide ways to resolve it.

Symptom

The installation of Tivoli Monitoring for Tivoli Storage Manager failed.

Causes

You can often resolve installation failures by reviewing the logs.zip compressed file.

Resolving the problem

When you install Tivoli Monitoring for Tivoli Storage Manager, several components are installed and several configuration scripts are run. If an error occurs when the components are installing, the installation stops at the current component and any remaining components are not installed. If an error occurs when the installation software runs any of the post configuration scripts, the remaining scripts are not run. After the failure occurs, review the files in the logs.zip file to find the cause of the failure. The logs.zip file contains several files for each component that was installed, and the files contain the results from each script that was run.

The components are installed in the following order: IBM DB2, IBM Tivoli Monitoring, and the monitoring agent.

The configuration scripts run in the following order and perform the following tasks: create the *WAREHOUS* database, configure the Tivoli Enterprise Portal, configure the Warehouse Proxy agent, configure the Summarization and Pruning

agent, configure history configuration, configure the data source, only if the Administration Center is installed, and create the KSK tables in the *WAREHOUS* database.

To begin reviewing files in the *logs.zip* file, complete the following steps:

1. Uncompress the *logs.zip* file. The *logs.zip* file is in the installation directory:

AIX **Linux** The default location for *logs.zip* is
 /opt/tivoli/tsm/reporting.

Windows The default location for *logs.zip* is C:\IBM.

2. Review the *MachinePlan_localhost.log* file to determine where the installation failed.

coi/plan/install/MachinePlan_localhost/logs/[INSTALL_xxxx_xx.xx]/
 MachinePlan_localhost.log where the *xxx_xxx.xx* is the timestamp of the failure.

3. Further identify the error by using the following steps:
 - If a failure occurred when DB2 was being installed, see DB2 installation logs for a list of files to review.
 - If a failure occurred when IBM Tivoli Monitoring was being installed, see IBM Tivoli Monitoring installation logs for a list of files to review.
 - If a failure occurred when the monitoring agent was being installed, see Monitoring agent installation logs for a list of files to review.
 - If a failure occurred when the scripts were being run, see Script failures for a list of files to review.
4. After the error is identified, and you know how to resolve it, uninstall Tivoli Monitoring for Tivoli Storage Manager before you reinstall it. For more information about uninstalling, see *Uninstalling Tivoli Monitoring for Tivoli Storage Manager* in the *Installation Guide*.

The tables list the files that you can review to determine the cause of failure. Search for words like *fail*, *err*, *severe*, *abort*, or *exception*.

Tip: Most return codes that you see are operating system return codes. Looking up the return codes specific to your operating system can reveal the cause of the failure.

Table 19. DB2 installation logs

Logs	Action
de/installUser/de_trace.log where the <i>installUser</i> is the system user ID that was used during the installation	Search this file, from the bottom up, for the letters <i>ACU</i> . Error messages in this file begin with <i>ACU</i> and end with <i>E</i> for error.
coi/plan/install/MachinePlan_localhost/00001_DB2/DB2.log	Review this log for any errors.
coi/plan/install/MachinePlan_localhost/00001_DB2/db2setup.log	Review this log for any errors.
coi/plan/logs	Review any DB2 related log files for any errors.

Table 20. IBM Tivoli Monitoring installation logs

Logs	Action
itm/itm/InstallITM/Abort IBM Tivoli Monitoring xxxxxxxx xxxx.log where the xxx_xx.xx is the timestamp of the failure	Review this log for any errors. This file contains information as to why the failure occurred.
installdir/InstallITM/TEPS_ODBC.log	The abort log might direct you to look in this file for the error. The file is found in the installation directory.
installdir/InstallITM/Warehouse_Configuration.log	The abort log might direct you to look in this file for the error. The file is found in the installation directory.

Table 21. Monitoring agent installation logs

Header	Action
coi/plan/install/MachinePlan_localhost/00003_TSM_Agent	Review any log files in this directory for any errors.
ia/log.txt or ia/log.txt.0	Review this log for any errors.
itm/itm/InstallITM/Abort IBM Tivoli Monitoring xxxxxxxx xxxx.log where the xxx_xx.xx is the timestamp of the failure	Review this log for any errors. This file might contain information as to why the failure occurred.

Script failures

After the components are installed, a series of scripts are run to configure the product. If a script fails, you can review the coi/plan/logs directory for the .log and .err files that correspond to the script that failed. For example, if history configuration failed, review the configure_history.log and configure_history.err files.

History configuration failed

AIX Linux Windows

After the Tivoli Monitoring for Tivoli Storage Manager components are installed, the installation wizard runs scripts to configure the IBM Tivoli Monitoring component. In some cases, the script to configure history configuration fails.

Symptom

The installation fails and the following message is displayed:
Failed to configure the history configuration

Causes

The history configuration script can fail because of a possible timing or firewall issue.

Resolving the problem

1. Confirm that the history configuration did not get configured, by completing the following steps:
 - a. Start the Tivoli Enterprise Portal, and log on by using the sysadmin ID and password.
 - b. Click **Edit > History Configuration**.

c. Expand **Tivoli Storage Manager** and highlight it.

In the expanded node, if there are no attributes that begin with KSK, for example, KSK_DATABASE, the history configuration script failed. In the right panel, under **Select Attributes Group(s)** pane, several of the rows in the **Group**

column must have a  icon, depicting a running man, otherwise the history configuration script failed. The icon is not displayed in every row by default.

2. Run the `setup_history` script by completing the following steps:

a. Open a command prompt window and change to the directory that contains the history configuration script. For example:

```
installMedia/COI/PackageSteps/Config_ITM
```

Tip: The script is in the installation media (DVD or downloaded package).

b. Run the `setup_history` script, issue the following command:

```
AIX Linux
setup_history.sh install_dir/itm sysadmin sysadmin_password
hostname_HUB > setupHistoryOutput.log
```

```
Windows
setup_history.bat install_dir\itm sysadmin sysadmin_password
hostname_HUB > setupHistoryOutput.log
```

where:

install_dir is the installation location of the software

sysadmin is a fixed value

sysadmin_password is the password of the *sysadmin* ID

hostname is the host name of the operating system

_HUB is case-sensitive

c. Review the `setupHistoryOutput.log` file to determine whether the script ran successfully. If it failed, resolve the error before you move on.

d. Start the Tivoli Enterprise Portal, and log on with the *sysadmin* ID and password.

e. Click **Edit > History Configuration**.

f. Look for several attribute groups that begin with KSK under the Tivoli Storage Manager node in the left pane.

g. In the right pane, under the Group column, look for several groups with the  icon, that depicts a running man, to the left of the name.

h. If you see the icons, it indicates that the attributes are summarizing and storing data. If there are no icons, review the `setupHistoryOutput.log` output file for any errors.

i. If Tivoli Common Reporting is installed, you must run the `datasrc` script to set up the data source. The data source connection allows Tivoli Common Reporting to connect to the *WAREHOUS* database to generate reports. For more information, see *Configuring a data source using the configuration wizard* in the *Administrator's Guide*.

Tivoli Enterprise Portal configuration fails on SUSE Linux Enterprise 11 during installation

Linux

When you install Tivoli Monitoring for Tivoli Storage Manager on SUSE Linux Enterprise 11, the installation might hang or fail when the Tivoli Enterprise Portal server is installed or configured.

Symptom

The Tivoli Monitoring for Tivoli Storage Manager installation fails when the installation sets up and configures the Tivoli Enterprise Portal server.

The installation hangs during the configuration of the Tivoli Enterprise Portal server.

The `_uninst/plan/logs/configure_cq.log` file contains the following errors:

```
DB20000I The TERMINATE command completed successfully.
SQL1013N The database alias name or database name "TEPS"
could not be found. SQLSTATE=42705
stty: standard input: the argument is invalid
Agent configuration has been started...
...InstallPresentation.sh is execute.
```

The following messages are seen in the `_uninst\plan\install\MachinePlan_localhost\00002_ITM\itm62_xx_error.txt` file:

```
insserv: Starting ITMAgents1 depends on SuSEfirewall2_setup and
          therefore on system facility `$all' which can not be true!
insserv: Max recursions depth 99 reached
insserv: loop involving service xdm at depth 2
insserv: There is a loop at service ITMAgents1 if started
insserv: loop involving service boot.udev at depth 1
insserv: loop involving service boot.swap at depth 8
insserv: loop involving service boot.sysctl at depth 8
insserv: loop involving service boot.localnet at depth 11
insserv: exiting now without changing boot order!
/sbin/insserv failed, exit code 1
```

Causes

A limitation exists when you install and configure Tivoli Enterprise Portal server on Linux platforms. This limitation is documented in the IBM Tivoli Monitoring manual.

For more information about this limitation, see the hardware and software requirements in the *IBM Tivoli Monitoring guide*.

Resolving the problem

TCP/IP network services such as NIS, DNS, and the `/etc/hosts` file must be configured to return the fully qualified host name, for example, `server1.ibm.com`, where `server1` is the host name of the system. You must define the fully qualified host name after the dotted-decimal host address value and before the short host name in the `/etc/hosts` file.

1. Verify that the `/etc/hosts` file has only one line that contains the short name of the host system. A system with multiple host names can result in Tivoli

Enterprise Portal server errors when either starting, reconfiguring, or running the buildPresentation.sh, InstallPresentation.sh, or migrate-export scripts.

2. Review the /etc/hosts file of the host where Tivoli Monitoring for Tivoli Storage Manager is installed. If the file contains the same host short name on different lines, an installation error can occur.

On SUSE Linux Enterprise 11, by default, the hosts file contains the following lines:

```
127.0.0.1 localhost localhost.mydomain
::1 localhost ipv6-localhost ipv6-loopback
```

3. Edit the hosts file to remove the extraneous short name entry, in this case, localhost. The updated file contains the following lines:

```
127.0.0.1 localhost localhost.mydomain
::1 ipv6-localhost ipv6-loopback
```

4. Uninstall Tivoli Monitoring for Tivoli Storage Manager and restart the installation.

Resolving errors with the Tivoli Enterprise Portal component of IBM Tivoli Monitoring

AIX

Linux

Windows

Tivoli Enterprise Portal is a graphical user interface that is part of the Tivoli Monitoring for Tivoli Storage Manager installation. You can use this interface to monitor Tivoli Storage Manager.

No data, or old data, is displayed in the Tivoli Enterprise Portal workspaces

AIX

Linux

Windows

A problem can occur with the monitoring agent if no data is displayed, or if no current data is displayed in the Tivoli Enterprise Portal workspaces.

Symptom

No data is displayed in some or all of the Tivoli Enterprise Portal workspaces, or data that is displayed is old and no longer current.

Causes

If all workspaces do not show data, it indicates that the monitoring agent is not collecting data. This monitoring agent might not be correctly configured or is not started.

If one or more workspaces do not show data, there might an issue with the agent query that runs on the Tivoli Storage Manager server.

If current data is not being displayed, an error occurred with the monitoring agent.

Resolving the problem

1. Create and configure an agent instance, if you did not do so already. For detailed procedures about creating and configuring an agent instance, see the *Tivoli Storage Manager Installation Guide*.
2. Ensure that the monitoring agent is started by completing the following steps:

AIX

Linux

- a. Issue the following command to navigate to the directory where the CandleManage program is located:

```
cd /opt/tivoli/tsm/reporting/itm/bin
```

- b. Run the CandleManage program by issuing this command:
./CandleManage &

If the logon failed, reconfigure the monitoring agent to update the configuration settings, for example, the password, server address, and port number.

- c. In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is started. If it is not started, select the applicable agent instance, right-click, and select **Start**.

Tip: To start the monitoring agent on the command line, see the information about configuring the agent instance with the command-line interface in the *Tivoli Storage Manager Installation Guide*.

Windows

- a. Click **Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services**.
 - b. In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is started. If it is not started, select the applicable agent instance, right-click, and select **Start**.
3. Ensure that the monitoring agent is communicating with the Tivoli Storage Manager server by completing the following steps:
 - a. Open the Tivoli Enterprise Portal, and check the Agent Log workspace for messages that indicate whether the logon to the Tivoli Storage Manager server was successful. You might see one of the following phrases within the message: TSM Login Succeeded or TSM Login Failed. If you see the TSM Login Failed message, reconfigure the monitoring agent to update the password, server address, and port number settings.
 - b. Optionally, you can review the activity log on the Tivoli Storage Manager server to determine whether queries from the monitoring agent are being logged.
 4. Determine whether an error occurred with a monitoring agent query that can be causing no data to be displayed in the workspace. Follow the steps in Empty workspaces when agent queries run too long.
 5. Review the agent log files for error messages. You can review these log files in the following directory:

AIX

Linux

```
install_dir/ITM/logs
```

Windows

```
install_dir\ITM\TMAITM6\logs
```

6. If the error is not resolved, you can enable tracing on the monitoring agent and examine the output for additional information. For tracing instructions, see Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on an AIX or Linux system, or Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on a Windows operating system.

Empty workspaces when agent queries run too long

AIX

Linux

Windows

The monitoring agent query runs on the Tivoli Storage Manager server. If a query takes a long time to run and does not return any data, the result is that one or more workspaces are empty.

Symptom

One or more Tivoli Enterprise Portal workspaces do not display data. These workspaces are empty.

Causes

Queries are sent to the Tivoli Storage Manager server and run one at a time. After the query returns data, the next one runs. An error can occur when a query takes too long to run. The query times out and no data is returned. The remaining queries are not run. The result is that one or more workspaces are empty.

Resolving the problem

You can identify a query that is taking too long to run by using the following techniques:

Agent log

Review the Agent Log workspace for the monitoring agent instance. Whenever a query runs, you can see a TSM Login Succeeded message that is associated with an attribute group. If the message is not being logged to the Agent Log workspace, you can use that information to help identify the query that is causing the problem. For example, if the Database workspace is empty, you might not see the TSM Login Succeeded message that is associated with the Database attribute group in the Agent Log workspace.

Tivoli Storage Manager activity log

Review the Tivoli Storage Manager server activity log to determine the order that the queries run. By default, the queries are run after every hour. In the activity log, find the last query that is run in the hour. Manually run that query from the Tivoli Storage Manager server administrative command line to determine whether it takes more than a few minutes to run. If a query takes a long time to run, that query is likely to be the one causing the error.

Agent trace

You can enable an agent trace to identify the query that is taking too long to run. For tracing instructions, see [Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on an AIX or Linux system](#), or [Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on a Windows operating system](#).

After you allow the trace to run for a few hours, review the trace file. By default, the queries are run every hour. In a span of an hour, review the attribute groups for the corresponding workspaces that are empty. The table includes a list of all attribute group names. In the trace, look for the message TSM Login Succeeded. The attribute name is displayed before a message similar to `NodeActivity: TSM Login Succeeded`.

Any of the attribute groups that are missing is associated with a query that is causing the issue on the Tivoli Storage Manager server. Use the table to turn on or off the attribute groups, one at a time, until you find the query with the issue.

After you identify the attribute group that is causing the error, you can stop the query from running. Use the table to find the corresponding environment variable that you must modify to stop the query from running. The following file contains the environment variables that you can update:

```

AIX Linux install_dir/itm/config/KSKENV_agentName
Windows install_dir\itm\tmaitm6\KSKENV_agentName

```

By default, the values of the environment variable are set to 1, which is On. You can update the variable to 0 to turn it Off. Save the file and restart the monitoring agent and determine whether there are still several empty workspaces.

Note: The workspace for the attribute group that you turned off no longer collects current data.

Table 22. Attribute group name and the corresponding environment variable. The attribute group name corresponds to the workspace name you see in the Tivoli Enterprise Portal.

Attribute group name	Environment Variable
Activity summary	KSK_SUMM_ON
Activity log	KSK_ACTLOG_ON
Client node storage	KSK_CNSTG_ON
Client missed files	KSK_CMF_ON
Client node status	KSK_CNS_ON
Database	KSK_DB_ON
Drives	KSK_DRV_ON
Libraries	KSK_LIB_ON
Node activity	KSK_NODEA_ON
Occupancy	KSK_OCC_ON
PVU details	KSK_PVU_ON
Replication details	KSK_REPLDETAIL_ON
Replication status	KSK_REPLSTAT_ON
Schedule	KSK_SCHED_ON
Storage pool	KSK_STGPOOL_ON
Server	KSK_SERVER_ON
Storage device	KSK_STGDEV_ON
Tape usage	KSK_TAPEUSG_ON
Tape volume	KSK_TAPEVOL_ON

Note: The Sessions attribute group is not listed in this table because data is collected directly from the Tivoli Storage Manager server. The queries for Sessions are immediately sent to the Tivoli Storage Manager server and the workspace is populated by data that is returned from those queries. The data is not stored. If there are any issues with the query, do not click the workspace.

Cannot log on to the Tivoli Enterprise Portal

AIX

Linux

Windows

You are unable to log on to the Tivoli Enterprise Portal.

Symptom

You cannot log on to the Tivoli Enterprise Portal and an error message is displayed in the logon window.

Causes

Various errors can occur that can prevent you from logging on to the Tivoli Enterprise Portal.

Resolving the problem

The following table provides solutions for problems you might encounter when you log on to the Tivoli Enterprise Portal.

Table 23. Reasons that you cannot log on to the Tivoli Enterprise Portal.

Problem	Action and Solution
KFWITM393E User ID or password is invalid. - OR - KFWITM215E: Unable to process logon request	<ul style="list-style-type: none">• Start Manage Tivoli Monitoring Services and verify that the Tivoli Enterprise Portal server is started.• Ensure that the user ID and password are correct when you log on to the Tivoli Enterprise Portal. Use the sysadmin user ID and password.• Determine whether the sysadmin account password is expired. If so, update the password and corresponding services by following the steps in Resetting passwords for Tivoli Monitoring for Tivoli Storage Manager.• Verify that the sysadmin account is not locked. The account is an operating system account.
KFWITM010I: Tivoli Enterprise Portal Server not ready. - OR - KFWITM402E: Communication with the Tivoli Enterprise Server could not be established.	<ul style="list-style-type: none">• Verify that the Tivoli Enterprise Portal server is started.• Restart the Tivoli Enterprise Portal server.• The Portal server initialization can take up to 20 minutes. Wait 20 minutes and log on again. <p>You can review the trace log for more information. Start Manage Tivoli Monitoring Services, right-click the Tivoli Enterprise Portal server, and click Advanced > View trace log.</p>
KFWITM392E: Internal error occurred during logon.	Start the DB2 instance. If it is already started, restart it. From a DB2 command window, issue the db2stop and db2start commands to stop and start the DB2 instance.

Table 23. Reasons that you cannot log on to the Tivoli Enterprise Portal. (continued)

Problem	Action and Solution
<p>The portal server cannot initialize because of a DB2 shutdown.</p> <p>- OR -</p> <p>KFWITM009I: The Tivoli Enterprise Portal Server is still being initialized and is not ready for communications.</p>	<ul style="list-style-type: none"> Start the DB2 instance and wait a few minutes for it to finish initializing. If you receive message KFWITM009I, you can look at the most recent trace log to verify that the portal server is initialized. Search for the text string Waiting for requests. Startup completed. To view the trace log, start Manage Tivoli Monitoring Services, right-click the Tivoli Enterprise Portal server, and click Advanced > View trace log.
<p>KFWITM008W The Tivoli Enterprise Portal Server has lost contact with the Tivoli Enterprise Monitoring Server.</p>	<p>Start Manage Tivoli Monitoring Services and verify that the Tivoli Enterprise Monitoring server is started.</p>

Missing monitoring agent from the Tivoli Enterprise Portal

AIX Linux Windows

After you install Tivoli Monitoring for Tivoli Storage Manager, no monitoring agent is listed in the Navigator pane of the Tivoli Enterprise Portal.

Symptom

The monitoring agent is not displayed in the Tivoli Enterprise Portal.

Causes

The monitoring agent is not created and configured, or the agent is not started after its configuration.

Resolving the problem

- Verify that a monitoring agent is created. For detailed procedures about creating and configuring an agent instance, see the *Tivoli Storage Manager Installation Guide*.
- Ensure that the monitoring agent is started by completing the following steps:

AIX Linux

- Go to the directory where the CandleManage program is located by issuing the following command:


```
cd /opt/tivoli/tsm/reporting/itm/bin
```
- Run the CandleManage program by issuing this command:


```
./CandleManage &
```
- In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is started. If it is not started, select the applicable agent instance, right-click, and select **Start**.

Tip: To start the monitoring agent by using the command line, see the information about creating and configuring the agent instance using the command-line interface in the *Tivoli Storage Manager Installation Guide*.

Windows

- a. Click **Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services**.
 - b. In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is started. If it is not started, select the applicable agent instance, right-click, and select **Start**.
3. Start the Tivoli Enterprise Portal and determine whether the monitoring agent is displayed in the **Navigator** pane.

Start and end times display 00:00:00 in Node Activity workspace

AIX

Linux

Windows

When you view the Node Activity workspace in the Tivoli Enterprise Portal, the Start Date and End Date columns always show 00:00:00 for the time.

Symptom

The Node Activity workspace displays 00:00:00 as a time for every row in the Start Date and End Date columns.

Causes

The column that displays 00:00:00 in all rows are correct values by the design.

Resolving the problem

The time is displaying 00:00:00 because the data can be gathered from the Tivoli Storage Manager server activity log table by multiple sessions and streams, and therefore is only totaled for a 24-hour period.

Removing agent instances does not stop them

AIX

Linux

Windows

Removing the monitoring agent instance from Manage Tivoli Enterprise Monitoring Services window does not always stop the process from running.

Symptom

When you remove the monitoring agent instance, the agent instance service continues to run.

Causes

This problem is caused by failing to stop the monitoring agent instance before removing it.

Resolving the problem

Restart the system to stop the service.

Requirement: Stop the monitoring agent before you remove it.

AIX

Linux

1. Verify whether a process is running by issuing the following command:
`ps -ef | grep -i SK`
2. If the process is running, stop the process by issuing the following command:
`kill -9 process_ID`

Resolving historical data reporting problems in the Warehouse Proxy workspace

Windows

If you experience errors in the Warehouse Proxy workspace in the Tivoli Enterprise Portal browser when using a non-English system, you might have to re-configure the ITMUser ID. You must add the ITMUser ID to the Administrator group or a group with the same authority.

If the ITMUser ID is not part of the Administrators group or part of a group with authority like the Administrators group, errors are generated on non-English systems. The following are examples of possible errors:

- Initialization with Datasource "ITM Warehouse" failed. A retry will be attempted in 10 minutes.
- The bufferpool ITMBUF8K could not be created.

Perform the following steps to resolve the reporting problem:

1. Right-click **My Computer**.
2. Select **Manage**.
3. Select **Local Users and Groups**.
4. Select **Groups**.
5. Select **Administrators** or a group with the same authority.
6. Select **Add**.
7. Type **ITMUser** in the space provided.

Unable to start Tivoli Enterprise Portal using a non-English administrator name

You cannot access Tivoli Enterprise Portal using a non-English locale with a system administrator name that contains non-English characters.

Symptom

When you log on to the Tivoli Enterprise Portal, the window displays an error message indicating an APPCRASH. For example, you might see this error message if you are running with a Russian locale, and are trying to access the Tivoli Enterprise Portal with a Russian administrator ID. When you logon to the Tivoli Enterprise Portal, instead of the log on window, the following information is displayed:

```
Problem signature:  
Problem Event Name: APPCRASH  
Application Name: java.exe  
Application Version: 5.0.0.0  
Application Timestamp: 4930ade2  
The name of the module with the error: webrenderer.dll  
Version of the module with the error: 0.0.0.0  
The timestamp of the module with the error: 4a00f901
```

Exception Code: c0000005
Exception Offset: 00002f44
OS Version: 6.0.6002.2.2.0.274.10
Language Code: 1049
Additional Information 1: 94c5
Additional Information 2: 69ca74015f8538bb7b0988b56b6976e5
Additional Information 3: 37de
Additional Information 4: 516f45eb25de8caa79c140ad69cfd8d

Read our privacy statement:
<http://go.microsoft.com/fwlink/?linkid=50163&clcid=0x0419>

Causes

You might have a problem due to logging in to the Tivoli Enterprise Portal with an administrator system account ID that contains non-English characters. The Tivoli Enterprise Portal tries to create or access files in the /Users/administrator folder, but fails to access these directories, because the administrator name is non-English.

Resolving the problem

Verify whether any non-English characters were used when logging in. If there are non-English characters, create an administrator user account with English characters and assign the appropriate administrator authority. Log off and log on again using the new English administrator account to access the Tivoli Enterprise Portal.

Resolving errors with the Tivoli Common Reporting reports

AIX Linux Windows

Tivoli Common Reporting is an optional component of the Administration Center. Use Tivoli Common Reporting to obtain historical data about your Tivoli Storage Manager servers. If problems occur in Tivoli Common Reporting, you can diagnose and resolve the problems in a number of ways.

Resolving no data in Tivoli Common Reporting reports

AIX Linux Windows

The flowchart provides steps to help guide you through resolving problems with the Tivoli Common Reporting reports.

Symptom

After you run a Cognos® or BIRT report, the report does not display data. The report either states that there is no data to be displayed, or an error message is displayed.

For BIRT reports, if there is no error message and no data is displayed, the following message is displayed:

Empty data set.

For Cognos reports, if there is no error message and no data is displayed, the following message is displayed:

No data is available based on your selections.

Causes

Several errors might cause data to be missing from the Tivoli Common Reporting reports. The flowchart covers the possible causes and resolutions.

Resolving the problem

Use the flowchart as a step-by-step procedure to diagnose problems where no data is displayed in the Tivoli Common Reporting reports. The diamonds are decision questions to be answered. The boxes are the tasks to complete. Answering *Yes* or *No* advances you to the next question or final task. Use the two tables for more information about the question and task. Decisions are numbered 1 - 8, and tasks are lettered A - H.

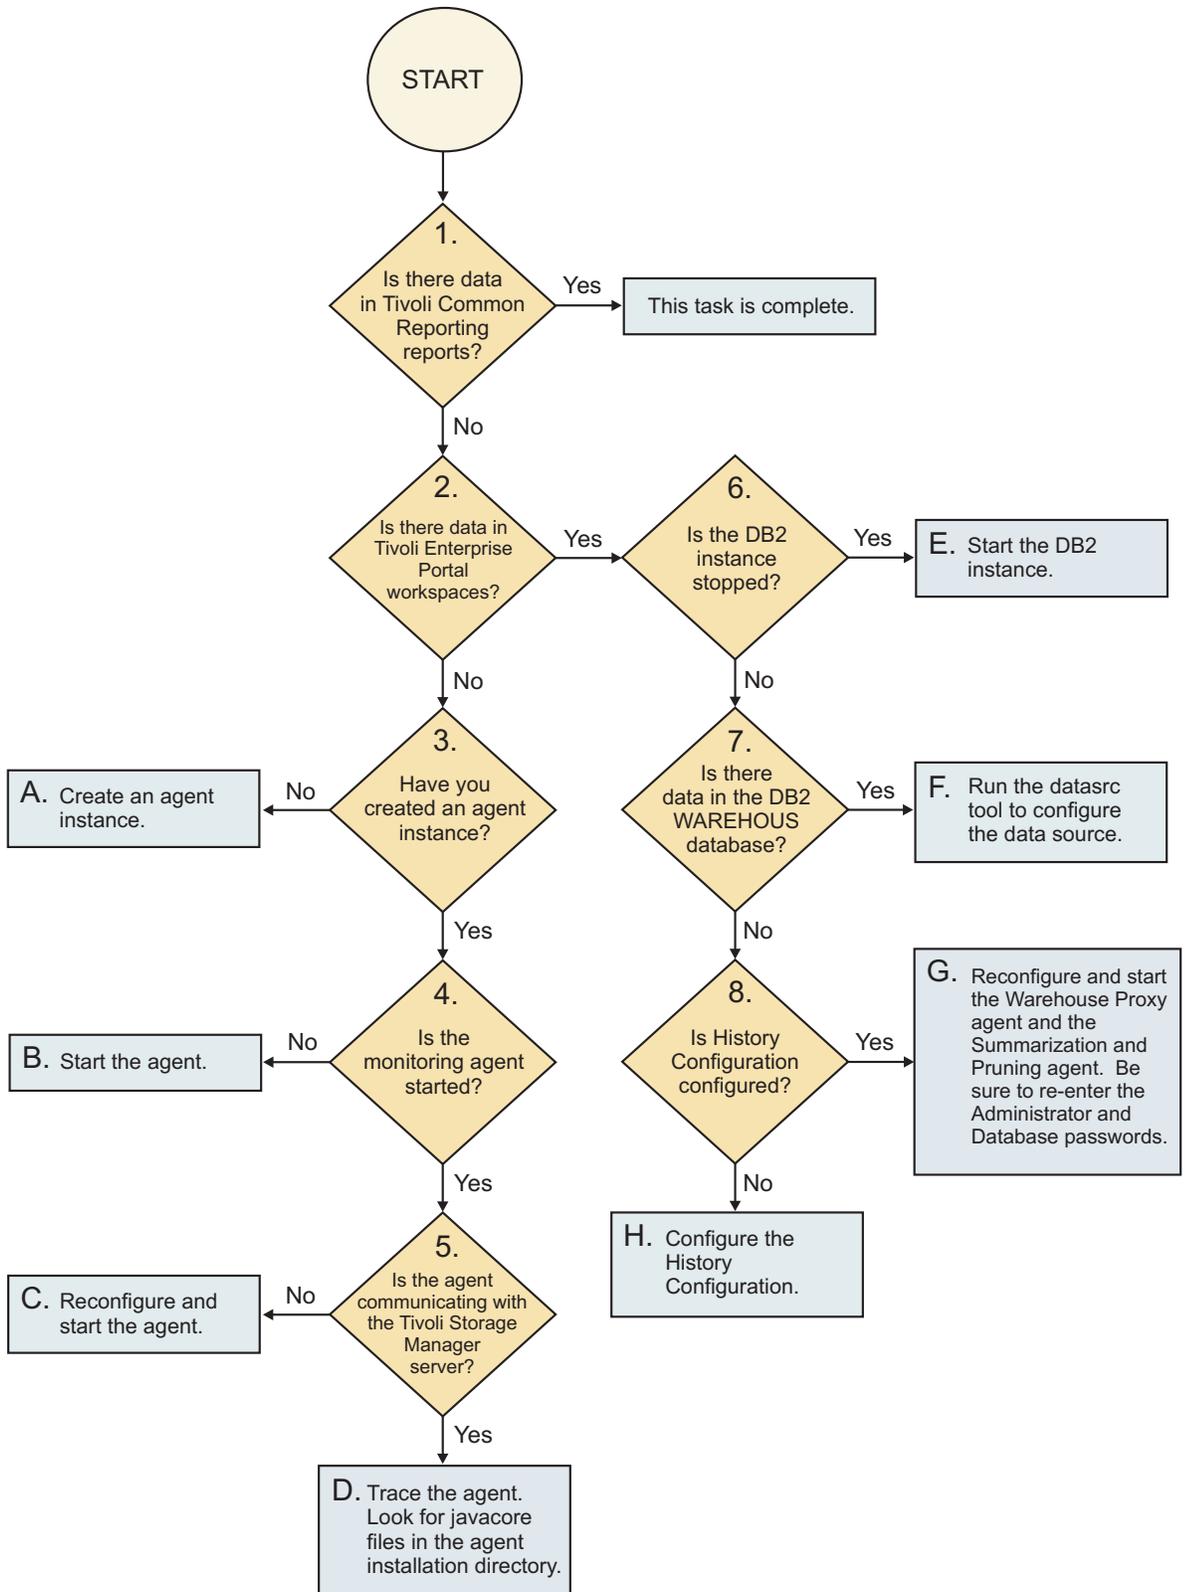


Figure 1. Flowchart to resolve problems with the Tivoli Common Reporting reports.

Table 24. Explanation of the flowchart questions and the next step to advance to, based on your response.

Question	Question Explanation	Decision
<p>Decision 1: Is there data in Tivoli Common Reporting reports?</p>	<p>Log on to the Administration Center and view any of the Tivoli Common Reporting reports.</p> <p>If no data is displayed, BIRT reports might display Empty data set or some SQL error.</p> <p>The Cognos reports might display <i>No data is available based on your selections</i>, or An attempt to connect to the data source failed, or some other error.</p>	<p>If yes, the task is complete.</p> <p>If no, go to decision 2.</p>
<p>Decision 2: Is there data in Tivoli Enterprise Portal workspaces?</p>	<p>Start the Tivoli Enterprise Portal, and log on by using the syadmin ID and password. View the workspaces to see whether they contain data.</p>	<p>If yes, go to decision 6.</p> <p>If no, go to decision 3.</p>
<p>Decision 3: Have you created an agent instance?</p>	<p>Start Manage Tivoli Monitoring Services:</p> <p>AIX Linux</p> <p>/opt/tivoli/tsm/reporting/itm/bin> CandleManage &</p> <p>Windows Click Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services.</p> <p>In the Manage Tivoli Enterprise Monitoring Services window, verify that you created and configured a monitoring agent instance for the Tivoli Storage Manager server.</p>	<p>If yes, go to decision 4.</p> <p>If no, go to task A.</p>
<p>Decision 4: Is the agent started?</p>	<p>Start Manage Tivoli Enterprise Monitoring Services:</p> <p>AIX Linux</p> <p>/opt/tivoli/tsm/reporting/itm/bin> CandleManage &</p> <p>Windows Click Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services.</p> <p>In the Manage Tivoli Enterprise Monitoring Services window, find the monitoring agent. Look in the Status column to see whether the monitoring agent is started.</p>	<p>If yes, go to decision 5.</p> <p>If no, go to task B.</p>
<p>Decision 5: Is the agent communicating with the Tivoli Storage Manager server?</p>	<p>Start the Tivoli Enterprise Portal, and log on by using the sysadmin ID and password.</p> <p>Review the Agent Log workspace for the monitoring agent. Search for messages like TSM Login Succeeded, or TSM Login Failed. The TSM Login Failed message means that the agent is not connected to the Tivoli Storage Manager server.</p>	<p>If yes, go to task D.</p> <p>If no, go to task C.</p>

Table 24. Explanation of the flowchart questions and the next step to advance to, based on your response. (continued)

Question	Question Explanation	Decision
Decision 6: Is the DB2 instance stopped?	Verify whether the DB2 instance is started or stopped.	If yes, go to task E. If no, go to decision 7.
Decision 7: Is there data in DB2?	To verify whether there is data in the tables of the <i>WAREHOUS</i> database, follow the steps in “Verifying the <i>WAREHOUS</i> data for each report” on page 219. Find the table that corresponds to the report that you are viewing. Issue the DB2 <code>select</code> statement from the DB2 command window to determine whether there is any data in the table and if the data is current.	If yes, go to task F. If no, go to decision 8.
Decision 8: Is history configuration configured?	Start the Tivoli Enterprise Portal, and log on by using the <code>syadmin</code> ID and password. Click Edit > History Configuration . Click and highlight the Tivoli Storage Manager node. In the right pane, in the Groups column, if there are no rows with the  icon that depicts a running man, the history configuration is not configured.	If yes, go to task G. If no, go to task H.

Table 25. Explanation of the tasks to resolve issues with the Tivoli Common Reporting reports.

Task	Task Explanation
Task A: Create an agent instance.	To create an agent, see <i>Creating and configuring the agent instance using the command-line interface</i> in the <i>Installation Guide</i> .
Task B: Start the agent.	Start Manage Tivoli Monitoring Services: <div style="display: flex; justify-content: space-around; margin-bottom: 10px;"> AIX Linux </div> <pre style="margin: 0;">/opt/tivoli/tsm/reporting/itm/bin> CandleManage &</pre> <div style="margin-bottom: 10px;"> Windows Click Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services. </div> In the Manage Tivoli Enterprise Monitoring Services window, double-click the monitoring agent to start it.
Task C: Reconfigure and start the agent.	To configure the agent, see <i>Creating and configuring the agent instance</i> in the <i>Installation Guide</i> .

Table 25. Explanation of the tasks to resolve issues with the Tivoli Common Reporting reports. (continued)

Task	Task Explanation
Task D: Trace the agent. Look for javacore files in the agent installation directory.	<p>Enable tracing by following the steps in Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on an AIX or Linux system or Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on a Windows operating system.</p> <p>Determine whether there are long running queries that can cause no data to be displayed in one or more workspaces by following the steps in Empty workspaces when agent queries run too long.</p> <p>If javacore files are on your operating system, it is an indication that there is a major error when you attempt to start the agent. You can look for files with the word javacore in the file name. Search for these files in the <code>itm</code> directory and its subdirectories.</p>
Task E: Start the DB2 instance.	Start the DB2 database by using the <code>db2start</code> command from the DB2 command window.
Task F: Run the <code>datasrc</code> tool to configure the data source.	To reconfigure the data source, see <i>Configuring a data source using the configuration wizard</i> in the <i>Installation Guide</i> .
Task G: Reconfigure and start the Warehouse Proxy agent and the Summarization and Pruning agent. Ensure that you reenter the Admin and Database passwords.	<p>To reconfigure the Warehouse Proxy agent, see <i>Manually configuring the Warehouse Proxy agent</i> in the <i>Installation Guide</i>.</p> <p>To reconfigure the Summarization and pruning agent, see <i>Manually configuring the Warehouse Summarization and Pruning agent</i> in the <i>Installation Guide</i>.</p>
Task H: Configure the history configuration.	To configure history configuration, follow the steps in History configuration failed.

No current data displayed in the Tivoli Common Reporting reports

AIX Linux Windows

Tivoli Common Reporting reports are populated by data from the DB2 `WAREHOUS` database. If current data is not displayed in the reports, data collection is stopped.

Symptom

Current data is not displayed in the reports.

Causes

The monitoring agent is not collecting data, or an error occurred in DB2 that prevents the data from being exported to the `WAREHOUS` database.

Resolving the problem

1. Ensure that the monitoring agent is started by completing the following steps:

AIX Linux

- a. Go to the directory where the CandleManage program is located by issuing the following command:


```
cd /opt/tivoli/tsm/reporting/itm/bin
```
- b. Run the CandleManage program by issuing this command:


```
./CandleManage &
```
- c. In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is started. If it is not started, select the applicable agent instance, right-click, and select **Start**.

Tip: To start the monitoring agent using the command line, see the information about creating and configuring the agent instance using the command-line interface in the *Tivoli Storage Manager Installation Guide*.

Windows

- a. Click **Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services**.
 - b. In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is started. If it is not started, select the applicable agent instance, right-click, and select **Start**.
2. Ensure that the monitoring agent is communicating with the Tivoli Storage Manager server.
 - a. Start the Tivoli Enterprise Portal, and review the Agent Log workspace for messages that indicate whether the logon to the Tivoli Storage Manager server was successful. You might see one of the following phrases within the message: TSM Login Succeeded or TSM Login Failed. If you see the TSM Login Failed message, reconfigure the monitoring agent to update the password, server address, and port number settings.
 - b. Optionally, you can review the activity log on the Tivoli Storage Manager server to determine whether the queries from the monitoring agent are being logged.
 3. Review several of the Tivoli Enterprise Portal workspaces to determine whether current data is displayed. Several of the workspaces display timestamps that you can use to determine whether data is current. If data is not current or there is no data at all, follow the steps in “No data, or old data, is displayed in the Tivoli Enterprise Portal workspaces” on page 186.
 4. Ensure that the DB2 instance is started.
 5. Reconfigure the Warehouse Proxy agent and reenter the Admin and Database passwords. Wait a few hours for data to be exported to the *WAREHOUS* database. For more information, see *Manually configuring the Warehouse Proxy agent* in the *Installation Guide*.

Missing server name in the Tivoli Common Reporting report prompt page

AIX

Linux

Windows

When you run a report, in certain conditions, you can select a Tivoli Storage Manager server from a list or menu. If the list or menu is empty or an expected server is not displayed, the problem can be caused by a configuration error.

Symptom

In a Tivoli Common Reporting report, the option to select a Tivoli Storage Manager server name either does not list any servers or the server name that you want to select is not displayed.

Causes

The server name might not be displayed in the server list or on the menu for several reasons. Common reasons include:

- A monitoring agent is not created.
- The monitoring agent is not collecting data.
- History configuration is not configured.
- The DB2 database is not running.
- There is an authentication failure with the Warehouse Proxy agent.

Resolving the problem

1. Verify that a monitoring agent is created and configured. For detailed procedures about creating and configuring an agent instance, see the *Tivoli Storage Manager Installation Guide*.
2. Ensure that the monitoring agent is started by completing the following steps:

AIX

Linux

- a. Go to the directory where the CandleManage program is located by issuing the following command:

```
cd /opt/tivoli/tsm/reporting/itm/bin
```
- b. Run the CandleManage program by issuing this command:

```
./CandleManage &
```
- c. In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is started. If it is not started, select the applicable agent instance, right-click, and select **Start**.

Tip: To start the monitoring agent by using the command line, see the information about creating and configuring the agent instance using the command-line interface in the *Tivoli Storage Manager Installation Guide*.

Windows

- a. Click **Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services**.
 - b. In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is started. If it is not started, select the applicable agent instance, right-click, and select **Start**.
3. Ensure that the monitoring agent is communicating with the Tivoli Storage Manager server by completing the following steps:
 - a. Start the Tivoli Enterprise Portal, and review the Agent Log workspace for messages that indicate whether the logon to the Tivoli Storage Manager server was successful. You might see one of the following phrases within the message: TSM Login Succeeded or TSM Login Failed. If you see the TSM Login Failed message, reconfigure the monitoring agent to update the password, server address, and port number settings.
 - b. Review a few workspaces to verify that data is displayed.

- c. Optionally, you can review the activity log on the Tivoli Storage Manager server to determine whether queries from the monitoring agent are being logged.
 4. Verify that history configuration is configured by completing the following steps:
 - a. Start the Tivoli Enterprise Portal, and log on with the *sysadmin* ID and password.
 - b. Click **Edit > History Configuration**.
 - c. Look for several attribute groups that begin with KSK under the Tivoli Storage Manager node in the left pane.
 - d. In the right pane, under the Group column, look for several groups with the  icon, that depicts a running man, to the left of the name.
 - e. If you see the icons, the history configuration script ran successfully and the attributes are summarizing and storing data. If you do not see the icons, follow the steps in History configuration failed.
 5. Ensure that the DB2 instance is started.
 6. Reconfigure the Warehouse Proxy agent and reenter the Admin and Database passwords. Wait a few hours for data to be exported to the *WAREHOUS* database. For more information, see *Manually configuring the Warehouse Proxy agent* in the *Installation Guide*.

No data is available message displayed in a Cognos report

AIX Linux Windows

Tivoli Common Reporting reports are populated by data from the DB2 *WAREHOUS* database. If no data is returned by an SQL query, no data is populated in the reports.

Symptom

If no data is populated in the reports, the following message is displayed:
No data is available based on your selections.

Causes

The SQL query that is sent to the *WAREHOUS* database is created based on the parameters that are selected for the report. If no data is returned from the query, a message is displayed that states that there is no data available. The message is displayed if there is no summarized data or if data collection is stopped.

Resolving the problem

1. Ensure that the monitoring agent is started by completing the following steps:

AIX Linux

- a. Go to the directory where the CandleManage program is located by issuing the following command:


```
cd /opt/tivoli/tsm/reporting/itm/bin
```
 - b. Run the CandleManage program by issuing this command:


```
./CandleManage &
```

- c. In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is started. If it is not started, select the applicable agent instance, right-click, and select **Start**.

Tip: To start the monitoring agent by using the command line, see the information about creating and configuring the agent instance using the command-line interface in the *Tivoli Storage Manager Installation Guide*.

Windows

- a. Click **Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services**.
 - b. In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is started. If it is not started, select the applicable agent instance, right-click, and select **Start**.
2. Ensure that the monitoring agent is communicating with the Tivoli Storage Manager server by completing the following steps:
 - a. Start the Tivoli Enterprise Portal, and check the Agent Log workspace for messages that indicate whether the login to the Tivoli Storage Manager server was successful. You might see one of the following phrases within the message: TSM Login Succeeded or TSM Login Failed. If you see the TSM Login Failed message, reconfigure the monitoring agent to update the password, server address, and port number settings.
 - b. Review a few workspaces to verify that data is displayed.
 - c. Optionally, you can review the activity log on the Tivoli Storage Manager server to determine whether queries from the monitoring agent are being logged.
 3. Review several of the Tivoli Enterprise Portal workspaces to determine whether current data is displayed. Several of the workspaces display timestamps that you can use to determine whether data is current. If data is not current or there is no data at all, follow the steps in No data, or old data, is displayed in the Tivoli Enterprise Portal workspaces.
 4. Ensure that the DB2 instance is started.
 5. Reconfigure the Warehouse Proxy agent and reenter the Admin and Database passwords. Wait a few hours for data to be exported to the *WAREHOUS* database. For more information, see *Manually configuring the Warehouse Proxy agent* in the *Installation Guide*.
 6. If data is pruned too frequently, data might be deleted from the *WAREHOUS* database. Review the frequency of data that is being pruned by completing the following steps:
 - a. Start the Tivoli Enterprise Portal, and log on with the sysadmin ID and password.
 - b. Click **Edit > History Configuration**.
 - c. Click the **Tivoli Storage Manager** node and highlight the attribute groups.
 - d. Review the pruning settings in the **Configuration Controls** section of the pane. You can increase the values to retain data for a longer period.

Empty data set message in Tivoli Common Reporting

AIX

Linux

Windows

Tivoli Common Reporting reports are populated by data from the DB2 *WAREHOUS* database. The Empty data set message occurs when no data is returned from an SQL query.

Symptom

The issue applies only to BIRT reports. BIRT reports display the following message:
Empty Data Set
No data returned by query. Try another set of parameters.

Causes

The SQL query is created based on the parameters you select before you run the report. If the query returns no data, no data is displayed in the report. The report displays the Empty data set message, and you are instructed to select another set of parameters.

Resolving the problem

In some cases, the Empty data set message might be expected. For example, if you specify start and end dates that are before the Tivoli Monitoring for Tivoli Storage Manager was installed, this message can occur. If the Tivoli Monitoring for Tivoli Storage Manager monitoring agent should be collecting data, follow these steps to diagnose and resolve the problem:

1. Verify that a monitoring agent is created and configured. For detailed procedures about creating and configuring an agent instance, see the *Tivoli Storage Manager Installation Guide*.
2. Ensure that the monitoring agent is started by completing the following steps:

AIX

Linux

- a. Go to the directory where the CandleManage program is located by issuing the following command:

```
cd /opt/tivoli/tsm/reporting/itm/bin
```
- b. Run the CandleManage program by issuing this command:

```
./CandleManage &
```
- c. In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is started. If it is not started, select the applicable agent instance, right-click, and select **Start**.

Tip: To start the monitoring agent by using the command line, see the information about creating and configuring the agent instance using the command-line interface in the *Tivoli Storage Manager Installation Guide*.

Windows

- a. Click **Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services**.
 - b. In the Manage Tivoli Enterprise Monitoring Services window, verify that the monitoring agent is started. If it is not started, select the applicable agent instance, right-click, and select **Start**.
3. Ensure that the monitoring agent is communicating with the Tivoli Storage Manager server by completing the following steps:

- a. Start the Tivoli Enterprise Portal, and check the Agent Log workspace for messages that indicate whether the logon to the Tivoli Storage Manager server was successful. You might see one of the following phrases within the message: TSM Login Succeeded or TSM Login Failed.
- b. Review a few workspaces to verify that data is displayed.
- c. Optionally, you can review the activity log on the Tivoli Storage Manager server to determine whether queries from the monitoring agent are being logged.

If you see the TSM Login Failed message, reconfigure the monitoring agent to update the password, server address, and port number settings.

4. Review several of the Tivoli Enterprise Portal workspaces to determine whether current data is displayed. Several of the workspaces display timestamps that you can use to determine whether data is current. If data is not current or there is no data at all, follow the steps in No data, or old data, is displayed in the Tivoli Enterprise Portal workspaces.
5. Ensure that the DB2 instance is started.
6. Reconfigure the Warehouse Proxy agent and reenter the Admin and Database passwords. Wait a few hours for data to be exported to the *WAREHOUS* database. For more information, see *Manually configuring the Warehouse Proxy agent* in the *Installation Guide*.
7. If data is pruned too frequently, data might be deleted from the *WAREHOUS* database. Review the frequency of data that is being pruned by completing the following steps:
 - a. Start the Tivoli Enterprise Portal, and log on with the sysadmin ID and password.
 - b. Click **Edit > History Configuration**.
 - c. Click the **Tivoli Storage Manager** node and highlight the attribute groups.
 - d. Review the pruning settings in the **Configuration Controls** section of the pane. You can increase the values to retain data for a longer period. For more information, see *Configuring summarization and pruning values*

Error message RQP-DEF-0177 is displayed in a Cognos report

AIX Linux Windows

When you click a Cognos report, an SQL query is run against the DB2 *WAREHOUS* database. If the database table does not exist, an error is displayed and you are unable to run the report.

Symptom

When you attempt to run a Cognos report, the following error is displayed:

```
RQP-DEF-0177 An error occurred while performing operation
'sqlPrepareWithOptions' status='-201'.
```

Causes

The RQP-DEF-0177 message is displayed if the database table that is referenced in the query does not exist. The tables of the *WAREHOUS* database are created during the installation of Tivoli Monitoring for Tivoli Storage Manager. If the Tivoli Monitoring for Tivoli Storage Manager is not installed, or the script to create the tables fail, the RQP-DEF-0177 message is displayed.

Resolving the problem

1. Ensure that Tivoli Monitoring for Tivoli Storage Manager is installed, and that it is on the same system as the Administration Center.
2. Verify that the Warehouse Proxy agent can connect to the *WAREHOUS* database by completing the following steps:
 - a. Start the Tivoli Enterprise Portal and log on by using the sysadmin ID and password.
 - b. In the **Navigator** pane, find the **Warehouse Proxy** node and click to highlight it.
 - c. In the **Database Information** pane to the right, review the cell under **DB Connectivity**. The background color of the cell must be green have a value of **Yes**. If the background color is red with a value **No**, reconfigure the Warehouse Proxy agent and enter a valid password in the **Admin Password** field.

For more information, see *Manually configuring the Warehouse Proxy agent in the Installation Guide*.

3. Ensure that the DB2 instance is started.
4. Verify that the *WAREHOUS* database tables exist by completing the following steps:
 - a. Open a DB2 command-line interface and list the *WAREHOUS* tables by issuing the following commands:

```
db2 connect to WAREHOUS
db2 list tables for schema itmuser > warehouTables.txt
```

The *warehouTables.txt* file is created and contains the output from the `db2 list tables` command.

- b. Open the *warehouTables.txt* file and verify that there are several tables that begin with KSK, for example KSK_DATABASE. If there are no tables that begin with KSK, the script to create the tables failed. If you completed Step 2c to reenter a valid password, wait one to two hours for the necessary tables to be created.

Note: If you recently upgraded the Administration Center, wait approximately one or two hours for the new tables to be created. The tables are created only when there is data to be stored in the *WAREHOUS* database.

5. Verify that the setup of Tivoli Monitoring for Tivoli Storage Manager is correctly configured to allow data to be stored in the *WAREHOUS* database. Follow the steps in “No data is available message displayed in a Cognos report” on page 202.

Unable to find data source 'TDW' message is displayed in a Cognos report

AIX Linux Windows

A Cognos report might not run if a problem occurred with the Tivoli Data Warehouse data source.

Symptom

The following error is displayed when you select a Cognos report:

```
QE-DEF-0288 Unable to find the data source 'TDW' in the content store
```

Causes

The QE-DEF-0288 message is caused because of the Tivoli Data Warehouse data source was incorrectly set, a Tivoli Data Warehouse data source definition is missing, or because IBM Tivoli Monitoring for Tivoli Storage Manager is not installed.

Resolving the problem

Verify that the Tivoli Monitoring for Tivoli Storage Manager software is installed on the same system as the Administration Center. If it is not installed, install it and run the report again. If the error persists, continue with the following steps:

1. Log on to the Administration Center.
2. Click **Reporting > Common Reporting**.
3. In the right pane titled Work with Reports, click **Launch** and select **Administration**.

Tip: If this option is not available, log on to the Administration Center with a user ID with more privileges.

4. Click the **Configuration** tab and click **Data Source Connections**. The navigation trail displayed is now **Directory > Cognos**.
5. If you see the Tivoli Data Warehouse data source in the **Name** column, continue with the instructions for Editing the TDW data source. Otherwise, continue with the instructions for Creating the TDW data source.

Editing the Tivoli Data Warehouse data source	Creating the Tivoli Data Warehouse data source
<ol style="list-style-type: none"> 1. Click the TDW data source in the Names column. 2. In the Actions column, click More. 3. Click Set properties. 4. Click the Connection tab. 5. Scroll down to the Connection String field. Click the pencil icon to the right of the field. 6. Verify that the DB2 database name is WAREHOUS. 7. Scroll down to the Signon section and select The signons of this connection, and ensure that Password is selected. 8. In the Testing section, click Test the connection. 9. In the Test the connection - Tivoli Data Warehouse window, click Test. 10. In the View the results window, in the Status column, the message must read Succeeded. 11. Click Close to return to the previous screen and click Close again. 12. If the test does not work, your signon is invalid. 13. At the upper right of window, click Close (X) next to Help. 14. In the Set properties - Tivoli Data Warehouse window, click Cancel. 15. In the Perform an action - Tivoli Data Warehouse window, click Cancel. 16. The breadcrumb trail at the top must read: Directory > Cognos > TDW > TDW. Click TDW in the Name column. 17. You must see <i>itmuser</i> in the Name column. In the Action column, click More. 18. Click Set properties. Click the Signon tab, and click Edit the signon. 19. In the Enter the signon - itmuser window, ensure that the user ID is itmuser and enter the correct password in the two fields. Click OK. 20. In the Set properties - itmuser window, click OK. 21. The breadcrumb trail must read: Directory > Cognos > TWD > TDW. Click Cognos and then repeat these instructions to edit the Tivoli Data Warehouse data source to test the connection. 	<ol style="list-style-type: none"> 1. Click the New Data source icon at the upper right. The icon looks like a blue disk drive with a plus sign on it. 2. In the Name field, enter: TDW and click Next. 3. Select DB2 from the Type list and click Next. 4. In the DB2 Database name field, enter WAREHOUS. 5. Select Password. 6. Complete the user ID and Password fields to connect to the data source. The user ID is itmuser. 7. In the Testing section, click Test the connection. 8. Click Test. You must see the Succeeded message or the user ID or password that you entered in are incorrect. Click Close twice, and then click Next. 9. Click Finish.

Missing BIRT or Cognos reports

AIX

Linux

Windows

If BIRT or Cognos reports are missing from Tivoli Common Reporting, you can import them.

Symptom

You are unable to find the Tivoli Common Reporting reports in the Administration Center.

Causes

If BIRT and Cognos packaged or custom reports are removed from Tivoli Common Reporting, you can import them.

Resolving the problem

The BIRT reports are packaged in the `BirtReports.zip`. The Cognos reports are packaged in the `TSM_Cognos.zip`.

Import the packaged BIRT or Cognos reports by completing the following steps:

1. Log on to the system where the Administration Center is installed, and open a command window.
2. Find the `BirtReports.zip` file or the `TSM_Cognos.zip` file from in the Administration Center installation media (DVD or downloaded package).
3. From a command window, go to the following directory:

AIX

Linux

```
/opt/IBM/tivoli/tipv2Components/TCRComponent/bin/
```

Windows

```
C:\IBM\tivoli\tipv2Components\TCRComponent\bin\
```

4. Import BIRT and Cognos reports by issuing the following commands:

- a. Import the BIRT reports by issuing the following command:

AIX

Linux

```
./trcmd.sh -import -bulk path/BirtReports.zip -username tipadmin  
-password tipadmin password
```

where *path* refers to the path to the compressed files from Step 2.

Windows

```
trcmd.bat -import -bulk path\BirtReports.zip -username tipadmin  
-password tipadmin password
```

where *path* refers to the path to the compressed files from Step 2.

- b. Import the Cognos reports by issuing the following command:

AIX

Linux

```
./trcmd.sh -import -bulk path/TSM_Cognos.zip -username tipadmin  
-password tipadmin password
```

where *path* refers to the path to the compressed files from Step 2.

Windows

```
trcmd.bat -import -bulk path\TSM_Cognos.zip -username tipadmin  
-password tipadmin password
```

where *path* refers to the path to the compressed files from Step 2.

If the commands were successful, the following message is displayed:

```
CTGTRQ092I Import operation successfully performed.
```

5. Reconfigure the data source. For more information, see *Configuring a data source using the configuration wizard* in the *Installation Guide*.

No group membership exists to access Tivoli Common Reporting reports

AIX Linux Windows

If you cannot access Tivoli Common Reporting from the Administration Center in Tivoli Integrated Portal, the user ID that you used might not exist or have the correct roles associated with it.

Symptom

You are unable to access Tivoli Common Reporting from the Administration Center.

Causes

No group exists that enables a newly defined user ID to access the Tivoli Common Reporting from Administration Center, or the user ID was not granted the roles necessary to access the reports.

Resolving the problem

You can create a user ID, or add a role to an existing user ID, by completing the following steps:

1. Create a user ID by completing these steps:
 - a. Log on to the Tivoli Integrated Portal with the `tipadmin` credentials.
 - b. Click **Users and Groups** > **Manage Users**, and click **Create**.
 - c. Specify a new user ID, for example: `tipadmin2`.
 - d. Click **Group Membership** > **Search**.
 - e. Select **TSM_Administrator** and click **Add** > **Close**.
2. Add a role to an existing user ID by completing these steps:
 - a. Log on to the Tivoli Integrated Portal with the `tipadmin` credentials.
 - b. Click **Users and Groups** > **User Roles**.
 - c. Click **Group Membership**, and then click **search**.
 - d. Select **TSM_Administrator**, click **Add**, and then click **Close**.

Tip: Repeat step 2d to add the user ID to the `tcPortalOperator` and `chartAdministrator` groups.

Start and end times display 12:00:00 AM in a Cognos report

AIX

Linux

Windows

When you view a custom Cognos report in Tivoli Common Reporting, the Start Date and End Date columns always show 12:00:00 AM for the time.

Symptom

In a custom Cognos report, the Start Date and End Date times display 12:00:00 AM in every row.

Causes

The column that displays 12:00:00 AM in all rows are correct values by the design.

Resolving the problem

The time is showing up as 12:00:00 AM because data can be gathered from the Tivoli Storage Manager server activity log table by multiple sessions and streams, and therefore is only totaled for a 24-hour period.

Reports do not display in Japanese, Korean, or Traditional Chinese languages

Windows

If your reports are not being generated in Japanese, Korean, or Traditional Chinese, you might need to rebuild or reconfigure the UTF8TEST table in the *WAREHOUS* database.

IBM Tivoli Storage Manager reports are displayed in several languages, but problems occur with the output of Japanese, Korean, and Traditional Chinese reports. One reason might be that the historical collection configuration was not set up correctly in the Tivoli Enterprise Portal.

Ensure that the following items are completed:

1. Connect to DB2 by issuing the following command:
`connect to WAREHOUS`
2. Drop the UTF8TEST table in the *WAREHOUS* database by issuing the following command:
`drop UTF8TEST`
3. Stop and restart the Warehouse Proxy agent by completing the following steps:
 - a. Open the Manage Tivoli Enterprise Monitoring Services application.
 - b. Select **Warehouse Proxy**, right-click, and select **Stop**.
 - c. Select **Warehouse Proxy**, right-click, and select **Start**.

Tip: Stopping and restarting the Warehouse Proxy agent rebuilds the UTF8TEST table within the DB2 *WAREHOUS* database.

If you are still experiencing problems with reports in Japanese, Korean, or Traditional Chinese, ensure that the following items are set up:

- Verify that the DB2CODEPAGE=1208 environment variable exists. To add this variable, right-click on **My Computer**, select **Properties > Advanced > Environment Variables > New** for the system variables.
- Open the DB2 Control Center and select the *WAREHOUS* database.
- Select and open the UTF8TEST table and change the “Results” column on the third row to zero.
- Stop and restart the Warehouse Proxy Agent within the Tivoli Enterprise Monitoring Server.

Resolving a java.net.ConnectException: CTGRI0011E error

Windows

You can resolve java.net connection errors by installing and enabling the remote registry service.

By using the remote registry service, a user can log on to a system with different credentials, and complete various tasks. If the service is disabled, you cannot log on to a system with the db2admin user ID and password to install software.

The remote registry access service is enabled by default on Windows systems, but if the service is disabled, you might see error message CTGRI0011E, when the RXA connection is attempted.

Messages with the prefix CTGRI are related to the Service Request Manager installation program. The messages are in the CTGInstallTrace.log. The program cannot access remote systems that were designated as part of the product deployment.

The following message indicates that either the remote registry service is disabled or that the logon credentials are not valid.

```

java.net.ConnectException: CTGRI0011E An error occurred when accessing the remote registry
or service control manager.
at com.ibm.tivoli.remoteaccess.BaseProtocol.beginSession(BaseProtocol.java:638)
at RXAPrototypeMain.makeConnection(RXAPrototypeMain.java:176)
at RXAPrototypeMain.main(RXAPrototypeMain.java:54)
Caused by: java.net.ConnectException: CTGRI0011E An error occurred when accessing
the remote registry or service control manager.
at com.ibm.tivoli.remoteaccess.WindowsProtocol$RegistryKeyContext.<init>
(WindowsProtocol.java:11525)
at com.ibm.tivoli.remoteaccess.WindowsProtocol.getRegistryValue(
WindowsProtocol.java:8607)
at com.ibm.tivoli.remoteaccess.WindowsProtocol.getRegistryKeyString_internal
(WindowsProtocol.java:8810)
at com.ibm.tivoli.remoteaccess.WindowsProtocol.getRegistryKeyString
(WindowsProtocol.java:5943)
at com.ibm.tivoli.remoteaccess.WindowsProtocol.getTempDir
(WindowsProtocol.java:3483)
at com.ibm.tivoli.remoteaccess.WindowsProtocol.getCurrentDirectory
(WindowsProtocol.java:3340)
at com.ibm.tivoli.remoteaccess.WindowsProtocol.beginSessionImpl
(WindowsProtocol.java:1186)
at com.ibm.tivoli.remoteaccess.BaseProtocol$BeginSession.invoke
(BaseProtocol.java:1756)
at com.ibm.tivoli.remoteaccess.util.ThreadedOperation$OperationThread.run
(ThreadedOperation.java:154)
Caused by: com.starla.smb.SMBException: Access denied
at com.starla.smb.client.admin.IPCPipeFile.checkStatus
(IPCPipeFile.java:683)
at com.starla.smb.client.admin.WinregPipeFile.openKey
(WinregPipeFile.java:166)
at com.ibm.tivoli.remoteaccess.WindowsProtocol$RegistryKeyContext.<init>
(WindowsProtocol.java:11502)
... 8 more

```

Enable the remote registry service, and retry to failed operation.

Resolving web browser issues with the Administration Center or Tivoli Enterprise Portal

If you experience web browser issues, your security settings might be too restrictive, or you might be using an unsupported version. To avoid potential compatibility issues you can also use Java Web Start.

If the web browser window is blank when using Internet Explorer 7, the default security settings might be too restrictive.

Complete these steps to change your security settings:

1. In an Internet Explorer version 7 window, click **Tools > Internet Options**, and on the Security tab, click **Custom level**.
2. Locate the Miscellaneous category, disable the **Access data sources across domains** option, and click **OK** to apply the changes.

If your browser still does not display correctly, consider using a different browser.

Notes: The following browsers are supported

- Microsoft Internet Explorer 7 and 8
- Firefox 3.5

Restriction: Firefox 3.6 and later is not supported.

IBM Java version 5 is required to start Tivoli Enterprise Portal.

Tip: You can also view Tivoli Enterprise Portal remotely using the Java Web start applet by completing the following steps:

- Open a web browser and connect to `http://hostname:1920`.
- Click **IBM Tivoli Enterprise Portal Web Start Client** to start the Java Web Start client.
- Create a desktop shortcut to start Tivoli Enterprise Portal in the future.

Java Web start applet does not start when Tivoli Enterprise Portal server is running on a Windows system

Windows

The `http://teps_hostname:1920///cnp/kdh/lib/tep.jnlp` does not display, and the Java Web start applet does not start.

Symptom

The `http://teps_hostname:1920///cnp/kdh/lib/tep.jnlp` does not display because it is empty.

Resolving the problem

1. Open a command prompt window.
2. Create a temporary directory, such as `C:\temp`, by issuing the following command:

```
mkdir C:\temp
```
3. Change directory to the temporary directory by issuing the following command:

```
cd C:\temp
```
4. Issue the **jar** command that is provided with the IBM Tivoli Monitoring JRE to inflate the `jnlp` files in the temporary directory. If you installed IBM Tivoli Monitoring in the default directory, that directory is:

```
C:\IBM\ITM\CNPSJ\java\bin\jar xvf C:\IBM\ITM\InstallITM\plugin\plugins\plugin_cq-install_all-1.jar cq/install/shell/component.jnlp cq/install/shell/tep.jnlp
```
5. Move the inflated files to the IBM Tivoli Monitoring config directory by issuing the following commands:

```
move C:\temp\cq\install\shell\component.jnlp C:\IBM\ITM\config  
move C:\temp\cq\install\shell\tep.jnlp C:\IBM\ITM\config
```
6. Reconfigure Tivoli Enterprise Portal server by completing the following steps:
 - a. Start the Manage Tivoli Enterprise Monitoring Services application.
 - b. Select the Tivoli Enterprise Portal server and right-click **Reconfigure**.
 - c. Accept all of the defaults to reconfigure and restart the service.

File system is full or almost full

AIX

Linux

Windows

Files on the system can grow rapidly and fill up the file system. If the file system is full, the application might stop functioning and can prevent other applications that write to the file system from functioning.

Symptom

The file system is full, or almost full. You might notice that the *WAREHOUS* database is large, or some files in the installation directory are several gigabytes in size and growing rapidly.

Causes

The error can occur for several reasons:

- The Warehouse Proxy agent cannot export data to DB2:
 - The DB2 instance is not running.
 - Incorrect passwords are stored in the Warehouse Proxy agent configuration.
- The *WAREHOUS* database is not being pruned:
 - Pruning settings are not configured.
 - Incorrect passwords are stored in the Summarization and Pruning agent configuration.

Resolving the problem

If you notice that your file system is filling up, you can check the size of the binary files and *WAREHOUS* database. Use the steps in the table to identify and resolve the problem.

Table 26. Items to check to identify the reason that the file system is filling up

Check the size of the binary files	Check the size of the <i>WAREHOUS</i> database
<p>Data that comes from the Tivoli Storage Manager server is exported to temporary binary files before it is exported to the <i>WAREHOUS</i> database. If there is an issue, some binary files can grow rapidly and fill up the file system.</p> <ol style="list-style-type: none"> Go to the agent directory where the binary files are located: <div style="margin-left: 20px;"> <div style="display: flex; justify-content: space-around; margin-bottom: 5px;"> AIX Linux </div> <pre style="margin-left: 20px;">install_dir\itm\tmaitm6\logs\ History\KSK\agentName</pre> <div style="margin-left: 20px;"> Windows <pre style="margin-left: 20px;">install_dir\itm\tmaitm6\ logs\History\KSK\agentName</pre> </div> </div> Check the size of the binary files to see whether any are over 5 gigabytes. The binary files do not have an extension, for example, KSKDATABAS. If there are large files, a problem exists when data is trying to be exported to the <i>WAREHOUS</i> database. Ensure that the DB2 instance is started. Reconfigure the Warehouse Proxy agent to reenter the Admin and Database passwords. For more information, see <i>Manually configuring the Warehouse Proxy agent</i> in the <i>Installation Guide</i>. 	<p>If pruning on the database is not enabled, or the pruning values are low, the database can unexpectedly grow large, and fill up the system.</p> <ol style="list-style-type: none"> Check the size of the database directory. By default, the location of the database is: <div style="margin-left: 20px;"> <div style="display: flex; justify-content: space-around; margin-bottom: 5px;"> AIX Linux </div> <pre style="margin-left: 20px;">/home/db2inst1</pre> <div style="margin-left: 20px;"> Windows <pre style="margin-left: 20px;">C:\DB2</pre> </div> </div> Verify that pruning is enabled by looking at the history configuration settings from the Tivoli Enterprise Portal. Pruning the database ensures that data is removed periodically. For more information, see <i>Configuring summarization and pruning values</i> in the <i>Installation Guide</i>. If the directories are large enough that you suspect this is the reason that your file system is filling up, you can increase the pruning values. Review the history configuration settings from the Tivoli Enterprise Portal to increase the values to further reduce the size of the database. For more information, see <i>Configuring summarization and pruning values</i> in the <i>Installation Guide</i>. Reconfigure the Summarization and Pruning agent. Click Test database connection. For more information, see <i>Manually configuring the Warehouse Summarization and Pruning agent</i> in the <i>Installation Guide</i>.

Resetting passwords for Tivoli Monitoring for Tivoli Storage Manager

AIX
Linux
Windows

Some of the Tivoli Monitoring for Tivoli Storage Manager user account passwords might expire and cause the application to stop functioning. You can reset any of the passwords to continue monitoring your Tivoli Storage Manager servers.

Symptom

The following symptoms might occur on the operating system on which a password expired:

- An authentication error can occur when you try to start the Tivoli Enterprise Portal and log on.
- The Warehouse Proxy agent and Summarization Pruning agent can no longer access the *WAREHOUS* database. Errors for the two agents might be displayed in the Tivoli Enterprise Portal.
- Tivoli Common Reporting reports can fail to display recent historical data.

- Tivoli Common Reporting reports can fail to run because it cannot connect to the *WAREHOUS* database.

Causes

When you install Tivoli Monitoring for Tivoli Storage Manager, several operating system account passwords are created. Passwords are configured during the installation. If you forget a password, or if a password expires, you must reset the password, and then reconfigure one or more IBM Tivoli Monitoring services to update the password field.

Resolving the problem

You can reset any of the Tivoli Monitoring for Tivoli Storage Manager passwords, and then update the necessary services for the new password to be accepted. The accounts that are created are described in the table.

Complete the following steps to reset a Tivoli Monitoring for Tivoli Storage Manager account password and update the necessary services.

AIX

Linux

1. From a command prompt window, use the **passwd** operating system command to reset any of the Tivoli Monitoring for Tivoli Storage Manager passwords.
passwd *userID*
2. Use the table to reconfigure the services that are based on the user account password that was reset.

Windows

1. From the operating system, right-click **Computer > Manage**.
2. Expand **Local Users and Groups**.
3. Click **Users**.
4. Right-click the user account for which the password was reset, and select **Set Password**. Enter in the new password.
5. Use the information in the following table to reconfigure the services that are based on the user account password that was reset.

Review the table to determine which services need to be reconfigured.

Table 27. List of user IDs and corresponding services that must be reconfigured if passwords are reset

User ID	Action	Description of the user ID
sysadmin	After you reset this password, no other services must be updated.	Default Tivoli Enterprise Portal user ID

Table 27. List of user IDs and corresponding services that must be reconfigured if passwords are reset (continued)

User ID	Action	Description of the user ID
<p>AIX Linux db2inst1</p> <p>Windows db2admin</p>	<p>After you reset this password, open the Manage Tivoli Enterprise Monitoring Services window. Right-click the following services and click Reconfigure.</p> <ul style="list-style-type: none"> • Warehouse Proxy agent • Tivoli Enterprise Portal server <p>Windows From the Windows Start menu, click Run and type services.msc. Find all of the DB2 services that have Log on As as db2admin. For each service, right-click the service, and then click Properties. Click the Log on tab and enter in the new password.</p>	Database administrator user ID
<p>AIX Linux itmuser</p> <p>Windows ITMUser</p>	<p>After you reset this password, open the Manage Tivoli Enterprise Monitoring Services window. Right-click the following services and click Reconfigure.</p> <ul style="list-style-type: none"> • Warehouse Proxy agent • Summarization and Pruning agent • Tivoli Enterprise Portal server <p>Run the datasrc script to reconfigure the data source. The data source defines the connection to the WAREHOUS database from Tivoli Common Reporting.</p>	Tivoli Data Warehouse user ID
TEPS	If this password expires, update the password. No services need to be updated.	Tivoli Enterprise Portal server database user ID

For details about how to configure the Tivoli Enterprise Portal server, see the *Installation Guide*.

For details about how to reconfigure the Warehouse Proxy agent, see the *Installation Guide*.

For details about how to reconfigure the Summarization and Pruning agent, see the *Installation Guide*.

For details about how to configure the data source, see the *Installation Guide*.

Verifying the WAREHOUS data for each report

AIX

Linux

Windows

If the data that you expect is not displayed in the reports, you must verify that data is being stored in the tables in the *WAREHOUS* database tables.

Each table generates several different summarized views. The data within each table contains the following summarized views: *_H* (hourly), *_D* (daily), *_W* (weekly), *_M* (monthly), *_Q* (Quarterly), and *_Y* (Yearly). Some reports are generated from the summarized view while other reports use the details data.

The summarized view, in the form of a table, is named *KSK_table_name_HV*, where *table_name* is the name of the table that is used by the report. The report names are specified in the following tables, and the value for HV is the summarized view such as *_H*, *_D*, *_W*, *_M*, *_Q*, or *_Y*.

Table 28. Client BIRT report names, and the associated WAREHOUS table used to generate the report.

Client BIRT report name	WAREHOUS table that is used by the query	Requires summarized views to generate reports
Client activity details	KSK_NODE_ACTIVITY	No
Client activity history	KSK_NODE_ACTIVITY	Yes
Client backup currency	KSK_CLIENT_NODE_STATUS	No
Client backup missed files	KSK_CLIENT_MISSED_FILES	No
Client storage pool media details	KSK_CLIENT_NODE_STORAGE	No
Client storage summary	KSK_CLIENT_NODE_STORAGE	Yes
Client storage summary details	KSK_CLIENT_NODE_STORAGE	Yes
Client top activity	NODE_ACTIVITY	No
Node replication details	KSK_REPLDETAILS and KSK_REPLSTATUS	No
Node replication growth	KSK_REPLDETAILS	Yes
Node replication summary	KSK_REPLDETAILS	No
Schedule status	KSK_SCHEDULE	Yes

Table 29. Server BIRT report names, and the associated WAREHOUS table used to generate the report.

Server BIRT report name	WAREHOUS table that is used by the query	Requires summarized views to generate reports
Server activity details	KSK_NODE_ACTIVITY	No
Server activity log details	KSK_ACTIVITY_LOG	No
Server database details	KSK_DATABASE	Yes
Server resource usage	KSK_SERVER	Yes
Server throughput	KSK_SERVER	Yes

Table 29. Server BIRT report names, and the associated WAREHOUS table used to generate the report. (continued)

Server BIRT report name	WAREHOUS table that is used by the query	Requires summarized views to generate reports
Server tape volume capacity analysis	KSK_TAPE_VOLUME	Yes

Table 30. Cognos report names, and the associated WAREHOUS table used to generate the report.

Cognos Report name	WAREHOUS table that is used by the query	Requires summarized views to generate reports
Client activity status	KSK_ACTIVITY_SUMMARY	No
Client activity success rate	KSK_ACTIVITY_SUMMARY	Yes
Client backup currency	KSK_CLIENT_NODE_STATUS	No
Client backup status	KSK_CLIENT_NODE_STATUS	No
Client schedule status	KSK_SCHEDULE	No
Client schedule success rate	KSK_SCHEDULE	No
Client storage pool usage summary	KSK_OCCUPANCY	No
Current client occupancy summary	KSK_OCCUPANCY	No
Current storage pool summary	KSK_STORAGE_POOL	No
Highest storage space usage	KSK_CLIENT_NODE_STORAGE	No
Server database growth trends	KSK_DATABASE	Yes
Server schedule status	KSK_SCHEDULE	No
Server storage growth trends	KSK_SERVER	Yes
Yesterday's missed and failed client schedules	KSK_SCHEDULE	No

You can verify whether data is being stored in the DB2 tables by running DB2 SQL queries against the table view.

Table 31. Methods for querying the data in a WAREHOUS database table

Method 1	Method 2
<p>1. AIX Linux Using the root user ID, query the data in the KSK_DATABASE table by using the DB2 interface. Issue the following commands:</p> <pre>db2 connect to WAREHOUS select * from ITMUser."KSK_DATABASE"</pre> <p>2. Windows Query the data in the KSK_DATABASE table by using the DB2 interface. Completing the following steps:</p> <ol style="list-style-type: none"> Open a command prompt window and issue the following command: db2cmd Issue the following commands in the DB2 command prompt window: db2 connect to WAREHOUS select * from ITMUser."KSK_DATABASE" 	<p>1. AIX Linux Complete the following steps to query data in the KSK_DATABASE table by using a DB2 script:</p> <ol style="list-style-type: none"> Create a file called database.sql. Add the following information to the database.sql file: connect to WAREHOUS; select * from ITMUser."KSK_DATABASE"; Issue the following command: db2 -tvf database.sql > output.txt <p>2. Windows Query the data in the KSK_DATABASE table by using a DB2 script. Completing the following steps:</p> <ol style="list-style-type: none"> Create a file called database.sql. Add the following statements to the database.sql file: connect to WAREHOUS; select * from ITMUser."KSK_DATABASE"; Issue the following command from a DB2 command window: db2cmd Issue the following command: db2 -tvf database.sql > output.txt

Tip: The KSK_DATABASE table is one example of the tables in the WAREHOUS database.

Chapter 9. Resolving data storage problems

If you are experiencing a problem in storing or retrieving data, several methods are available to help you resolve the problem.

Resolving unreadable data problems

You might receive unreadable data during import or node replication processes related to a lack of code page conversion during these processes.

If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and uses extended ASCII characters can be affected.

To resolve the issue, update the fields with the appropriate UPDATE commands after the import or node replication operation.

Checking the server activity log to resolve data storage issues

Check the server activity log for other messages occurring 30 minutes before and 30 minutes after the time of the error.

Issue the **QUERY ACTLOG** command to check the activity log. Often, other messages that are issued can offer additional information about the cause of the problem and how to resolve it.

Checking HELP for messages issued for a data storage problem

Check HELP for any messages issued by IBM Tivoli Storage Manager.

The Tivoli Storage Manager messages provide additional information in the **Explanation**, **System Action**, or **User Response** sections of the message. Often, this supplemental information about the message might provide the necessary steps necessary to resolve the problem.

Recreating the data storage problem

If a problem can be easily or consistently recreated, it might be possible to isolate the cause of the problem to a specific sequence of events.

Data read or write problems might be sequence-related, in terms of the operations being performed, or might be an underlying device error or failure.

Typical problems related to the sequence of events occur for sequential volumes. One example would be that a volume is in use for a client backup and that volume is preempted by a data restore from another client node. This situation might surface as an error to the client backup session that was preempted. However, that client backup session might succeed if it was retried or if it was not preempted in the first place.

Resolving data storage errors related to reading or writing to a device

If there is an error due to reading or writing data from a device, many systems and devices record information in a system error log file. For example, the `errpt` file for AIX and the Event Log file for Windows.

If a device or volume used by Tivoli Storage Manager is reporting an error to the system error log file, it is likely a device issue. The error messages recorded in the system error log file might provide enough information to resolve the problem.

Changing the storage hierarchy to resolve data storage problems

The storage hierarchy includes the defined storage pools and the relationships between the storage pools on the server.

The storage pool definitions are also used by the storage agent. If attributes of a storage pool were changed, the change might affect data store and retrieve operations. Review any changes to the storage hierarchy and storage pool definitions. Issue the **QUERY ACTLOG** command to see the history of commands or changes that might affect storage pools. Also, use the following **QUERY** commands to determine if any changes were made:

- **QUERY STGPOOL F=D**

Review the storage pool settings. If a storage pool is **UNAVAILABLE**, then data in that storage pool cannot be accessed. If a storage pool is **READONLY**, then data cannot be written to that pool. If either situation is the case, review why these values were set and consider issuing the **UPDATE STGPOOL** command to set the pool to **READWRITE**. Another consideration is to review the number of scratch volumes that are available for a sequential media storage pool.

- **QUERY DEVCLASS F=D**

The storage pools can be influenced by changes to device classes. Review the device class settings for the storage pools, including checking the library, drive, and path definitions. Issue the **QUERY LIBRARY**, **QUERY DRIVE**, and **QUERY PATH** commands for sequential media storage pools.

Changing the server policies to resolve data storage problems

The server policy attributes that directly relate to data storage are the copy group destinations for backup and archive copy groups. Similarly, the management class, **MIGDESTINATION**, also impacts where data is stored.

Review any changes to the server storage policies. Issue the **QUERY ACTLOG** command to view the history of commands or changes that might affect storage policies. Also, use the following **QUERY** commands to determine if any changes were made:

- **QUERY COPYGROUP F=D**

Review the **DESTINATION** settings for the **TYPE=BACKUP** and **TYPE=ARCHIVE** copy groups. Also review the "Migration Destination" for management classes used by HSM clients. If storage pool destinations were changed and resulting data read or write operations are now failing, either evaluate the changes made and correct the problem or revert to the previous settings.

- **QUERY NODE F=D**

Assigning a node to a different domain might impact data read-and-write operations for that client. Specifically, the node might now be going to storage

pool destinations that are not appropriate, based on the requirements of this node. For example, it might be assigned to a domain that does not have any TYPE=ARCHIVE copy group destinations. If this node tries to archive data, it fails.

Resolving a data storage backup or copy problem that occurs only with a specific node

If you cannot backup or copy data to a specific node, you might not have an active data pool listed in your active destinations. These are specified in the node policy domain.

Issue the `QUERY NODE nodeName F=D` command to verify that the node that is storing the data is authorized. The `QUERY NODE` command finds the policy domain name to which the node is assigned. Issue the `QUERY DOMAIN domain_name` where *domain_name* is the output gathered from the previous `QUERY NODE` command. Look in the `ACTIVEDESTINATION` parameter for the list of active data ports. If the active data pool into which you want to store data is not on the list, issue the `UPDATE DOMAIN` command to add the active data pool to the list.

Resolving a data storage problem that occurs only for a specific volume

If problems occur only for a specific storage volume, there might be an error with the volume itself, whether the volume is sequential media or DISK.

If your operation is a data write operation, issue the `UPDATE VOLUME volumeName ACCESS=READONLY` command to set this volume to READONLY, then retry the operation. If the operation succeeds, try setting the original volume back to READWRITE by issuing the `UPDATE VOLUME volumeName ACCESS=READWRITE` command. Retry the operation. If the operation fails only when using this volume, consider issuing the `AUDIT VOLUME` command to evaluate this volume and issue the `MOVE DATA` command to move the data from this volume to other volumes in the storage pool. After the data is moved off of this volume, delete the volume by issuing the `DELETE VOLUME` command.

Hints and tips for storage

The hints and tips that are gathered here are from actual problem experiences. You might find that one of the solutions is right for addressing your IBM Tivoli Storage Manager problem.

Device driver hints and tips

Device driver problems might be attributed to the operating system, the application using the device, the device firmware, or the device hardware itself.

Whenever a device problem is encountered, ask “Has anything been changed?”

If the adapter firmware changed, this change might cause a device to exhibit intermittent or persistent failures. Try reverting back to an earlier version of the firmware to see if the problem continues.

If cabling between the computer and the device was changed, this change often accounts for intermittent or persistent failures. Check any cabling changes to verify that they are correct.

A device might exhibit intermittent or persistent failures if the device firmware was changed. Try reverting back to an earlier version of the firmware to see if the problem continues.

For SCSI connections, a bent pin in the SCSI cable where it connects to the computer (or device) can cause errors for that device or any device on the same SCSI bus. A cable with a bent pin must be repaired or replaced. Similarly, SCSI buses must be terminated. If a SCSI bus is improperly terminated, devices on the bus might exhibit intermittent problems, or data that is transferred on the bus might be or appear to be corrupted. Check the SCSI bus terminators to ensure that they are correct.

Remember: If the “hints and tips” information does not adequately address your device driver issue or this is the initial setup of your system's device drivers, please refer to the *IBM Tivoli Storage Manager Administrator's Guide* and *Tivoli Storage Manager Installation Guide*. Check also, that your hardware devices are supported by Tivoli Storage Manager. See <http://www.ibm.com/support/entry/portal/>.

Adjusting to operating system changes

Operating system maintenance can change kernel levels, device drivers, or other system attributes that can affect a device.

Similarly, upgrading the version or release of the operating system can cause device compatibility issues. If possible, revert the operating system back to the state prior to the device failure. If reverting is not possible, check for device driver updates that might be needed based on this fix level, release, or version of the operating system.

Adjusting to changes in the HBA or SCSI adapter connecting to the device

A device driver communicates to a given device through an adapter.

If it is a fibre channel-attached device, the device driver uses a host bus adapter (HBA) to communicate. If the device is SCSI attached, the device driver uses a SCSI adapter to communicate. In either case, if the adapter firmware was updated or the adapter itself was replaced, the device driver might have trouble using the device.

Work with the vendor of the adapter to verify that it is installed and configured appropriately. The following list shows the other possible steps:

- If the adapter was changed, try reverting back to the previous adapter to see if the issue is resolved.
- If other hardware in the computer was changed or the computer was opened, reopen the computer and check to make sure that the adapter is properly seated in the bus. By opening and changing other hardware in the computer, the cards and other connections in the computer might have loosened, which might cause intermittent problems or total failure of devices or other system resources.

Resolving a loose cable connection

Problems with the device might occur if a connection is loose from the computer to the cable, or from the cable to the device.

Check the connections and verify that the cable connections are correct and secure.

For SCSI devices, check that the SCSI terminators are correct and that there are no bent pins in the terminator itself. An improperly terminated SCSI bus might result in difficult problems with one or more devices on that bus.

Resolving error messages in the system error log

A device might try to report an error to a system error log.

The following are examples of various system error logs:

- errpt for AIX
- Event Log for Windows

The system error logs can be useful because the messages and information recorded might help to report the problem or the messages might include recommendations on how to resolve the problem.

Check the appropriate error log and take any actions based on the messages issued to the error log.

Supporting 64- or 32-bit Linux kernel modules for 64- or 32-bit applications

Linux

The Linux kernel modules dictate the bit mode of the Linux SCSI generic device driver, all different Host Bus Adapter (HBA) drivers, and other settings.

All of these kernel modules only support applications that have the same bit mode with running kernel modules. In other words, 64-bit kernel modules only support 64-bit applications on 64-bit Linux systems.

If a 32-bit application runs on a 64-bit Linux system and invokes a 64-bit kernel module, the 32-bit application causes a kernel segmentation fault. The same will happen if a 64-bit application invokes a 32-bit kernel module on a 32-bit Linux system.

To avoid a segmentation fault, ensure that the bit mode of the Linux kernel module and its applications are the same. That means 32-bit applications can only invoke 32-bit kernel modules on 32-bit Linux systems. 64-bit applications can only invoke 64-bit kernel modules on 64-bit Linux systems.

Running a Tivoli Storage Manager Linux server on x86_64 architecture

Linux

The 32-bit and 64-bit Linux operating systems can run on the AMD64 and EM64T systems, which are 64-bit systems.

A 64-bit IBM Tivoli Storage Manager Linux server and storage agent can only run on a AMD64/EM64T system with a 64-bit Linux operating system. Likewise, a 32-bit Tivoli Storage Manager Linux server and storage agent can only run on an AMD64/EM64T system with 32-bit Linux operating system.

A 64-bit Tivoli Storage Manager server issuing the **QUERY SAN** command requires a 64-bit host bus adapter (HBA) application programming interface (API) on an AMD64/EM64T system. If an AMD64 system is equipped with a Qlogic HBA, it could create a problem since, by default, Qlogic only provides a 32-bit HBA API on the AMD64 system. You must install the 64-bit HBA API on the system before issuing the 64-bit **QUERY SAN** command.

Adjusting to HBA driver changes on the Linux 2.6.x kernels

The most distinct change for HBA drivers on the Linux 2.6.x kernels is that all drivers have “ko” as a new suffix.

The following list shows the driver names and locations in 2.6.x kernels:

Adaptec

The driver (aic7xxx.ko) is located in the `/lib/modules/kernel-level/drivers/scsi/aic7xxx/` directory.

Emulex

The driver (lpfcdd.ko) is located in the `/lib/modules/kernel-level/drivers/scsi/lpfc/` directory.

Qlogic

Its driver names are `qla2xxx.ko`, `qla2100.ko`, `qla2200.ko`, `qla2300.ko`, `qla2322.ko`, and so on. There is a certain order to load the HBA drivers. The `qla2xxx.ko` is a base driver and should be loaded first. After loading the `qla2xxx.ko` driver, the system should then load the `qla2300.ko` driver if it is equipped with a Qla2300 card. All of drivers are located in the `/lib/modules/kernel-level/drivers/scsi/qla2xxx/` directory.

Enabling multiple LUN support on Linux kernels

Linux

To configure SCSI devices with multiple LUNs on a Linux system, the Linux kernel must be set to enable multiple LUN support.

Multiple LUN support on some Linux distributions, however, is not a default option and requires users to manually add this option to the running kernel. Perform the following steps to set up and enable multiple LUN support on IA32 architecture:

1. Add one parameter to a boot loader configuration file.
 - For LILO boot loader:
 - a. Add `append="max_scsi_luns=128"` to the `/ect/lilo.conf` file.
 - b. Run `lilo`.
 - For GRUB boot loader:
 - a. Add `max_scsi_luns=128` after the kernel image list at `/etc/grub.conf` file for RedHat distribution.
 - b. Add `max_scsi_luns=128` after the kernel image list at `/boot/grub/menu.1` file for SuSE distribution.
2. Restart the system.

Using Tivoli Storage Manager to perform a ddtrace on Linux or HP-UX

HP-UX

Linux

The passthru device driver can be traced by issuing the **DDTRACE** command.

To enable trace, issue the following commands from the server console or admin client:

- Linux:

```
trace enable lpdd <other server trace class names>
trace begin <file name>
```

- HP-UX:

```
trace enable pvrhppdd <other server trace class names>
trace begin <file name>
```

Select one of the following three options:

- ddtrace start librarydd tapedd (to trace both library and drive)
- ddtrace start librarydd (library trace only)
- ddtrace start tapedd (drive trace only)

Remember: **DDTRACE GET** and **DDTRACE END** are not required.

The IBM Tivoli Storage Manager passthru device driver trace cannot be enabled through the ddtrace utility.

Using the HP-UX passthru driver

HP-UX

If the autoconf utility did not claim your device, ensure that the correct drivers are loaded to the kernel.

For HP-UX 11i v2, the sctl driver, as well as the HP-UX stape, sdisk, and schgr native drivers are required for device configuration for the IBM Tivoli Storage Manager passthru device driver. For HP-UX 11i v3 on IA64, the esctl driver, as well as the HP-UX estape, esdisk, and eschgr native drivers are required for device configuration.

Refer to the HP-UX operating system documentation for information on how to load drivers to the kernel.

Related reference:

“Device driver hints and tips” on page 225

“Adjusting to changes in the HBA or SCSI adapter connecting to the device” on page 226

“Resolving a loose cable connection” on page 227

Updating device information of host systems on a dynamic SAN without restarting

When devices in a SAN environment change, the information about this changed environment is not automatically sent to host systems attached to the SAN.

If the device information has not been updated to host systems attached to the SAN, previously-defined device paths will no longer exist. If you use the existing device information to define device paths, backup, or restore data, these operations might fail. In order to avoid these kinds of failures, use a different method for different platforms to update the device information on the SAN without restarting host systems.

AIX Issue the **CFGMR** command to force the operating system to re-configure itself. Then run SMIT to re-configure your IBM Tivoli Storage Manager devices.

HP-UX Run **autoconf** with the **-f** option to issue the **IOSCAN** command which forces the operating system to re-scan SCSI buses and fibre channels.

Linux There is no system command to re-configure the operating system. In order to re-scan SCSI buses and fibre channels, the adapter drivers corresponding to these SCSI adapters and fibre channel adapters must be unloaded and then reloaded into the Linux kernel. After reloading HBA drivers, run **autoconf** or **TSMSCSI** to re-configure Tivoli Storage Manager devices on Linux. You might issue the **LSPCI** command to find out which SCSI adapter and fibre channel adapter is available on the system. The **RMMOD** command unloads a driver from the kernel and the **MODPROBE** command loads a driver to the kernel.

Table 32. HBA adapters and corresponding drivers for all architectures of Linux

HBA Adapters	HBA Driver Name	Available Architectures
Adaptec 7892	aix7xxx	IA32, AMD64
Qlogic 22xx	qla2200	IA32, AMD64
Qlogic 23xx	qla2300	IA32, AMD64
Qlogic 2362	qla2362	EM64T
Emulex	lpfcdd	IA32, iSeries®, pSeries®

Setting the multiple LUN options to “on” for Adaptec SCSI and Qlogic Fibre-Channel HBA BIOS settings on Linux

By default, Adaptec SCSI adapters set the multiple logical unit number (LUN) option to “off” in their BIOS, which makes the SCSI adapter driver unable to probe a SCSI unit with multi-LUN properly.

The multiple LUN option must be turned on.

Perform the following steps to turn on the multiple LUN options:

1. Press the Ctrl and A keys at the same time.
2. Select **SCSI Device Configuration** in the **Configure/View Host Adapter Setting**.
3. Change No to Yes for Bios Multiple LUN support.

Turning on the tape enable option:

By default, Qlogic Fibre host bus adapters set the tape enable option as off in their BIOS, which affects the running of some SCSI commands on several SCSI tape devices. Therefore, this option must be turned on.

Perform the following steps to turn on the tape enable option:

1. Press the Alt and Q keys at the same time.
2. Select **Advanced Settings**.
3. Change Disable to Enable for Fibre Channel Tape Support.

Hard disk drives and disk subsystems hints and tips

The IBM Tivoli Storage Manager server needs hard disk drives, disk subsystems, vendor-acquired file systems, and remote file systems to perform in a specific way. Performing in a specific way allows Tivoli Storage Manager to appropriately manage and store data by ensuring the integrity of the Tivoli Storage Manager server itself.

The following definitions are provided to help you better understand the hard disk drives and disk subsystems:

Hard disk drive

A hard disk drive storage device is typically installed inside a given computer and used for storage by a Tivoli Storage Manager server on that computer.

Disk subsystem

An external disk subsystem connects to a computer through a SAN (storage area network) or some other mechanism. Generally, disk subsystems are outside of the computer to which they are attached and might be located in close proximity, or they might be located much farther away. These subsystems might also have some method of caching the input/output requests to the disks. If data is cached, despite a bypass cache request which can occur on remote file systems and certain disk subsystems, input/output failures can result. The failures are due to a difference between the Tivoli Storage Manager tracking and what is actually resident in a file system. Remote file systems and disk subsystems exhibiting these characteristics are not supported. Disk subsystems often have their own configuration and management software. A disk subsystem must report the results synchronously.

The Tivoli Storage Manager server might define hard disk drives and disk subsystems used by the computer or operating system on the computer where the Tivoli Storage Manager is installed. Typically, a hard disk drive or disk subsystem is defined to the computer where Tivoli Storage Manager is installed as a drive or file system. After the hard disk drive or disk subsystem is defined to the operating system, Tivoli Storage Manager might use this space by allocating a database, recovery log, or storage pool volume on the device. The Tivoli Storage Manager volume subsequently looks like another file on that drive or file system.

Bypassing cache during write operations

Database, recovery log, and storage pool volumes are opened with the appropriate operating system settings to require data write requests to bypass any cache and be written directly to the device.

By bypassing cache during write operations, IBM Tivoli Storage Manager maintains the integrity of client attributes and data. Bypassing the cache is required. If an external event, such as a power failure, causes the Tivoli Storage Manager server or the computer where the server is installed to halt or break while the server is running, the data in the cache might or might not be written to the disk. If the Tivoli Storage Manager data in the disk cache is not successfully written to the disk, information in the server database or recovery log might not be complete. Also, data that was supposed to be written to the storage pool volumes might be missing.

Hard disk drives installed on the computer where the Tivoli Storage Manager server is installed and running have less of an issue with bypassing cache. In this case, the operating system settings that are used when Tivoli Storage Manager opens volumes on that hard disk drive generally manage the cache behavior appropriately and honor the request to prevent caching of write operations.

Typically, the use and configuration of caching for disk subsystems is a greater issue because disk subsystems often do not receive information from the operating system about bypassing cache for write operations. Disk subsystems also might ignore this information when a volume opens. Therefore, the caching of data write operations might result in corruption of the Tivoli Storage Manager server database or loss of client data, or both. The problems are dependent upon which Tivoli Storage Manager volumes are defined on the disk subsystem and the amount of data lost in the cache. Disk subsystems should be configured to not cache write operations when a Tivoli Storage Manager database, recovery log, or storage pool volume is defined on that disk. Another alternative is to use nonvolatile cache for the disk subsystem. Nonvolatile cache employs a battery backup or some other sort of scheme to allow the contents of the cache to be written to the disk if a failure occurs.

Moving existing data to other volumes prior to altering or moving the database

The size and location of IBM Tivoli Storage Manager storage pool volumes (files) can not change after they are defined and used by the server.

If the size is changed or the file is moved, internal information that Tivoli Storage Manager uses to describe the volume might no longer match the actual attributes of the file. If you need to move or change the size of a Tivoli Storage Manager storage pool volume, move any existing data to other volumes prior to altering or moving the database.

FILE directory mapping between storage agents and servers for shared files

IBM Tivoli Storage Manager servers and storage agents can access the same data in File device classes by defining a set of directories that should be used within a device class definition.

The directory name in a FILE device-class definition identifies the location where the server places the files that represent storage volumes for the device class. When processing the **DEFINE DEVCLASS** command, the server expands the specified directory name into its fully-qualified form, starting from the root directory.

You can specify one or more directories as the location of the files used in the FILE device class. The default is the current working directory of the server at the time that the command is issued. You can specify the directories for AIX, Linux, Oracle Solaris, or HP-UX.

Do not specify multiple directories from the same file system. Doing so can cause incorrect space calculations. For example, if the directories `/usr/dir1` and `/usr/dir2` are in the same file system, the space check will count each directory as a separate file system. The space check does a preliminary evaluation of available space during storage operations. If space calculations are incorrect, the server could commit to a FILE storage pool but not be able to obtain space, causing the operation to fail. If the space check is accurate, the server can skip the FILE pool in the storage hierarchy and use the next storage pool if one is available.

If the server needs to allocate a scratch volume, it creates a new file in the specified directory or directories. (The server can choose any of the directories in which to create new scratch volumes.) To optimize performance, ensure that multiple directories correspond to separate physical volumes.

See Table 33 for the file name extension created by the server for scratch volumes, depending on the type of data that is stored.

Table 33. File name extensions for scratch volumes

For scratch volumes used to store this data:	The file extension is:
Client data	.BFS
Export	.EXP
Database backup	.DBV

For each storage agent that shares **FILE** access, the **PATHs** defined to each **DRIVE** seen by the storage agent must provide access to the same set of directories. When the **PATHs** are defined, the directories for each storage agent must match in number and ordering for the directories as listed in the device class definition on the server. If these definitions are out of sync, the storage agent might be unable to access the FILE volumes, resulting in successful LAN-restores and mount failures for the LAN-free restore operations.

Tape drives and libraries hints and tips

Problems with tape drives and libraries might be with software on the computer trying to use the device, connections to the device, or the device.

Whenever a device problem is encountered, ask, “Has anything been changed?” Suspect anything on the computer trying to use the device. Or look at the device itself, especially if the device worked prior to a given change then stopped working after that change.

- If the adapter firmware changed, a device might exhibit intermittent or persistent failures. Try reverting back to an earlier version of the firmware to see if the problem continues.
- If cabling between the computer and the device was changed, intermittent or persistent failures might occur. Check any cabling changes to verify that they are correct.
- If the device firmware has changed, a device might exhibit intermittent or persistent failures. Try reverting back to an earlier version of the firmware to see if the problem continues.

Adjusting to operating system changes

Operating system maintenance can change kernel levels, device drivers, or other system attributes that can affect a device. Similarly, upgrading the version or release of the operating system can cause device compatibility issues.

If possible, revert the operating system to the state prior to the device failure. If you cannot revert the operating system, check for device driver updates that might be needed based on the fix level, release, or version of the operating system.

Adjusting to device driver changes

A device driver upgrade might result in a tape drive or library device not working. These issues can also occur as a result of the type of driver that you use.

When working with IBM libraries or drives, as opposed to using other vendor libraries and drives, the type of device driver that you choose is important. IBM libraries and drives should use the IBM device driver, while other vendor libraries and drives should use the IBM Tivoli Storage Manager device driver.

Revert to the previous (or earlier) version of the device driver to see if the problem was introduced by the newer version of the driver.

Adjusting to a replaced adapter or other hardware changes

A small computer system interface (SCSI) connection to the device uses a SCSI adapter. A fibre-channel (optical) connection to the device uses a host bus adapter (HBA).

In either case, the cause of the problem might be from a changed adapter or an open computer where other hardware was changed or fixed.

Remember: The connecting point for the device to the computer is known as an adapter. Another term for adapter is *card*.

See the following information to help you adjust to a replaced adapter or hardware:

- If the adapter was changed, revert to the previous adapter to see if the issue is resolved.

- If hardware in the computer was changed or the computer opened, check the computer to ensure that the adapter is properly seated in the bus. By opening and changing other hardware in the computer, the cards and other connections in the computer might have been loosened. Loosening the connections might cause intermittent problems or total failure of the devices or other system resources.

Resolving a loose cable connection

Problems might occur to the device if a connection is loose from the computer to the cable, or from the cable to the device.

Check the connections and verify that the cable connections are correct and secure.

For SCSI devices, check that the SCSI terminators are correct and that there are no bent pins in the terminator itself. An improperly-terminated SCSI bus might result in problems with one or more devices on that bus.

Using error messages to resolve a device malfunction

A device might report an error to a system error log where you can try to find the cause of the problem.

Examples of various system error logs are:

- errpt for AIX
- Event Log for Windows

The system error logs can be useful because the messages and information recorded might help to report the problem, or the messages might include recommendations on how to resolve the problem. Check the appropriate error log and take any recommended actions based on messages issued to the error log.

SAN hints and tips

Problems with a SAN (storage area network) might be with software on the computer trying to use the device, connections to the device, or the device.

Whenever a SAN problem is encountered, ask “Has anything been changed?” Any kind of changes might be suspect, from the computer trying to use the device to the device itself, especially if the device worked prior to a given change, then stopped working after that change.

To better understand how to diagnose problems with a SAN, review the following terminology and typical abbreviations:

Fibre channel

Fibre channel denotes a fibre-optic connection to a device or component.

Host bus adapter

A host bus adapter (HBA) is used by a given computer to access a SAN. An HBA is similar in function to a network adapter in how it provides access for a computer to a LAN (local area network) or WAN (wide area network).

SAN A SAN is a network of shared devices that are typically accessed using fibre. Often, a SAN is used to share devices between many different computers.

Knowing your SAN configuration

Understanding the SAN configuration is critical in SAN environments. Various SAN implementations have limitations or requirements on how the devices are configured and set up.

The three SAN configurations are point-to-point, arbitrated loop, and switched fabric.

Point-to-point

The devices are connected directly to the host bus adapter (HBA).

Arbitrated loop

Arbitrated loop topologies are ring topologies and are limited in terms of the number of devices that are supported on the loop and the number of devices that can be in use at a given time. In an arbitrated loop, only two devices can communicate at the same time. Data being read from a device or written to a device is passed from one device on the loop to another until it reaches the target device. The main limiting factor in an arbitrated loop is that only two devices can be in use at a given time.

Switched fabric

In a switched fabric SAN, all devices in the fabric will be fibre native devices. This topology has the greatest bandwidth and flexibility because all devices are available to all HBAs through some fibre path.

Verifying that your devices are supported by the SAN

Many devices or combinations of devices might not be supported in a given SAN. These limitations arise from the ability of a given vendor to certify their device using Fibre Channel Protocols.

For a given device, verify with the device vendor that it is supported in a SAN. Verification includes whether or not it is supported by the host bus adapters (HBAs) used in your SAN environment, meaning that you must verify with the vendors that this device is supported by the hubs, gateways, and switches that make up the SAN.

Ensuring that your HBA works with your SAN

The host bus adapter (HBA) is a critical device for the proper functioning of a SAN. The problems that might occur with HBAs range from improper configuration to outdated BIOS or device drivers.

For a given HBA, check the following items:

BIOS HBAs have an embedded BIOS that can be updated. The vendor for the HBA has utilities to update the BIOS in the HBA. Periodically, the HBAs in use on your SAN should be checked to see if there are BIOS updates that should be applied.

Device driver

HBAs use device drivers to work with the operating system to provide connectivity to the SAN. The vendor typically provides a device driver for use with their HBA. Similarly, the vendor provides instructions and any necessary tools or utilities for updating the device driver. Periodically the device driver level should be compared to what is available from the vendor and, if needed, updated to pick up the latest fixes and support.

Configuration

HBAs have a number of configurable settings. The settings typically affect how IBM Tivoli Storage Manager functions with a SAN device.

Related reference:

“HBA configuration issues”

HBA configuration issues

Host bus adapters (HBAs) have many different configuration settings and options.

The HBA vendor typically provides information about the settings for your HBA and the appropriate values for these settings. Similarly, the HBA vendor might provide a utility and other instructions on how to configure your HBA. The following settings are those that typically affect using IBM Tivoli Storage Manager with a SAN:

- Storage area network (SAN) topology

The HBA should be set appropriately, based on the currently-used SAN topology. For example, if your SAN is an arbitrated loop, the HBA should be set for this configuration. If the HBA connects to a switch, this HBA port should be set to “point-to-point” and not “loop.”

With Tivoli Storage Manager SAN Device Mapping, you can perform SAN discovery on most of the platforms and the persistent binding of the devices are not required. A Tivoli Storage Manager server can find the device if the device path was changed due to a restart or other reason.

Go to http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager to verify the platform/HBA vendor/driver level support for Tivoli Storage Manager SAN discovery.

- Fibre channel link speed

In many SAN topologies, the SAN is configured with a maximum speed. For example, if the fibre channel switch maximum speed is 1 GB/sec, the HBA should also be set to this same value. Or the HBA should be set for automatic (AUTO) negotiation if the HBA supports this capability.

- Is fibre channel tape support enabled?

Tivoli Storage Manager requires that an HBA is configured with tape support. Tivoli Storage Manager typically uses SANs for access to tape drives and libraries. As such, the HBA setting to support tapes must be enabled.

Fibre-channel switch configuration issues

A fibre-channel switch supports many different configurations. The ports on the switch must be configured appropriately for the type of SAN that is set up and for the attributes of the SAN.

The vendor for the switch usually provides information about the appropriate settings and configuration based upon the SAN topology being deployed. Similarly, the switch vendor should provide a utility and other instructions on how to configure it. The following settings typically affect how IBM Tivoli Storage Manager uses a switched SAN:

Fibre-channel link speed

In many SAN topologies, the SAN is configured with a maximum speed. For example, if the fibre-channel switch maximum speed is 1 GB/sec, the host bus adapter (HBA) should also be set to this same value. Or the HBA should be set for automatic (AUTO) negotiation if the HBA supports this capability.

Port mode

The ports on the switch must be configured appropriately for the type of SAN topology being implemented. For example, if the SAN is an arbitrated

loop, the port should be set to FL_PORT. For another example, if the HBA is connected to a switch, the HBA options should be set to “point-to-point” and not “loop.”

Data gateway port settings

A data gateway in a SAN translates fibre channel to SCSI for SCSI devices attached to the gateway.

Data gateways are popular in SANs because they allow the use of SCSI devices, therefore it is important that the port settings for a data gateway are correct.

The vendor for the data gateway usually provides information about the appropriate settings and configuration based upon the SAN topology being deployed and SCSI devices used. Similarly, the vendor might provide a utility and other instructions on how to configure it. The following settings can be used for the fibre channel port mode on the connected port on a data gateway:

Private target

Only the SCSI devices attached to the data gateway are visible and usable from this port. For the available SCSI devices, the gateway simply passes the frames to a given target device. Private target port settings are typically used for arbitrated loops.

Private target and initiator

Only the SCSI devices attached to the data gateway are visible and usable from this port. For the available SCSI devices, the gateway simply passes the frames to a given target device. As an initiator, this data gateway might also initiate and manage data movement operations. Specifically, there are extended SCSI commands that allow for vendor-acquired data movement. By setting a given port as an initiator, it is eligible to be used for vendor-acquired data movement SCSI requests.

Public target

All SCSI devices attached to the data gateway, as well as other devices available from the fabric, are visible and usable from this port.

Public target and initiator

All SCSI devices attached to the data gateway as well as other devices available from the fabric are visible and usable from this port. As an initiator, this data gateway might also initiate and manage data movement operations. Specifically, there are extended SCSI commands that allow for vendor-acquired data movement. By setting a given port as an initiator, it is eligible to be used for vendor-acquired data movement SCSI requests.

SAN configuration between devices

Devices in a SAN, such as a data gateway or a switch, typically provide utilities that display what that device sees on the SAN.

It is possible to use these utilities to better understand and troubleshoot the configuration of your SAN. The vendor for the data gateway or switch typically provides a utility for configuration. As part of this configuration utility there is usually information about how this device is configured and other information that this device sees in the SAN topology (of which it is a part). You can use these vendor utilities to verify the SAN configuration between devices:

Data gateway

A data gateway reports all the fibre-channel devices as well as the SCSI devices that are available in the SAN.

Switch

A switch reports information about the SAN fabric.

Tivoli Storage Manager Management Console

The IBM Tivoli Storage Manager management console displays device names and the paths to those devices, which can be useful to help verify that the definitions for Tivoli Storage Manager match what is actually available.

The fibre-channel link error report

Most SAN devices provide monitoring tools that can be used to report information about errors and performance statistics.

The vendor for the device should provide a utility for monitoring. If a monitoring tool is available, it typically reports errors. The following errors are experienced more frequently:

CRC error, 8b/10b code error, and other similar symptoms

These errors are recoverable, where the error handling is usually provided by firmware or hardware. In most cases, the method to recover the device is to retransmit the failing frame. The fibre-channel link is still active when these errors are encountered. Applications using a SAN device that encounter this type of link error usually are not aware of the error unless it is a solid error. A solid error is one where the firmware and hardware recovery cannot successfully retransmit the data after repeated attempts. The recovery for these type of errors is typically very fast and will not cause system performance to degrade.

Link failure (loss of signal, loss of synchronization, NOS primitive received)

This error indicates that a link is actually “broken” for a period of time. It is likely due to a faulty gigabit interface connector (GBIC), media interface adapter (MIA), or cable. The recovery for this type of error is disruptive. This error appears in the application using the SAN device that encountered this link failure. The recovery is at the command exchange level and involves the application and device driver having to perform a reset to the firmware and hardware, which causes the system to run degraded until the link recovery is complete. These errors should be monitored closely, as they typically affect multiple SAN devices.

Remember: Fibre-channel link errors are often caused by a customer engineer (CE) action to replace a SAN device. As part of the maintenance performed by the CE to replace or repair a SAN device, the fibre cable might be temporarily disconnected. If the fibre cable is disconnected, the time and duration of the error should correspond to when the service activity was performed.

Common SAN device errors

Several SAN-specific messages can be issued when you experience problems with your storage agent SAN devices.

See Table 34 on page 240 for errors that are generated for SAN devices.

Table 34. Common SAN device errors

Error	Explanation
<p>ANR8302E I/O error on drive <i>TSMDRIVE01 (/dev/mt9)</i> (OP=<i>WRITE</i>, Error Number=<i>5</i>, CC=<i>205</i>, KEY=<i>FF</i>, ASC=<i>FF</i>, ASCQ=<i>FF</i>, SENSE=<i>**NONE**</i>, Description=<i>General SCSI failure</i>). Refer to Appendix D in the 'Messages' manual for recommended action</p>	<p>This message is often issued for SAN device errors. The CC=205 reports that the device driver detects a SCSI adapter error. In the case of a SAN-attached device that encounters a link reset caused by link loss, it will be reported back to the device driver as a SCSI adapter error.</p> <p>The underlying cause of this error is the event that caused the link reset due to the link loss. The path for this device should be updated to ONLINE=NO by issuing the UPDATE PATH command. Do not set the path to ONLINE=YES until the cause for the link reset was isolated and corrected.</p>
<p>ANR8957E: <i>command</i> : Autodetect is OFF and the serial number reported by the library did not match the serial number in the library definition</p>	<p>The IBM Tivoli Storage Manager SAN Device Mapping encountered a path for the library that reports a different serial number than the current Tivoli Storage Manager definition for the library. The AUTODETECT parameter was set to NO for the command which prevented the server from updating the serial number for the library.</p> <p>Determine the new path and issue the UPDATE PATH command to correct this.</p>
<p>ANR8958E: <i>command</i> : Autodetect is OFF and the serial number reported by the drive did not match the serial number in the drive definition</p>	<p>Tivoli Storage Manager SAN Device Mapping encountered a path for a drive that reports a different serial number than the current Tivoli Storage Manager definition for that drive. The AUTODETECT parameter was set to NO for the command, which prevents the server from updating the serial number for this drive.</p> <p>Determine the new path and issue the UPDATE PATH command to correct this.</p>

Table 34. Common SAN device errors (continued)

Error	Explanation
<p>ANR8963E: Unable to find path to match the serial number defined for drive <i>driveName</i> in library <i>libraryName</i></p>	<p>The SAN Device Mapping was not able to find a SAN device that was previously defined to the server. The most likely cause for this is that the device itself has been removed or replaced in the SAN. The following steps might resolve this:</p> <ul style="list-style-type: none"> • Device Removed <p>If the device was removed from the SAN, simply delete the server definitions that refer to this device. Issue the <code>QUERY PATH F=D</code> command to determine any paths that reference the device. Then issue the DELETE PATH command to remove these paths.</p> • Device Replace <p>A SAN Device was replaced with a new device as a result of maintenance or an upgrade. Perform the following procedures:</p> <ul style="list-style-type: none"> – Try not to delete the drive or drive path definition after you replace the drive. – Issue one of the following server commands: <ul style="list-style-type: none"> - <code>UPDate DRive <libraryName> <driveName> SERIAL=AUTODetect</code> <p>This command force-records the new serial number into the server database. Because the drive is replaced, the element number stays the same.</p> - <code>UPDate PATH <sourceName> <driveName> SRCT=SERVER DESTT=DRIVE LIBRARY=<libraryName> DEVIce=xxxxx AUTODetect=Yes</code> <p>This command force-records the new serial number into the database. Because the drive is replaced, the element number stays the same.</p> – If the drive or drive path is deleted, redefine this new, replaced drive. You must restart the Tivoli Storage Manager server so that the element number/serial number map for the library is refreshed. This mapping only occurs at initialization. <p>Issue the QUERY PATH F=D command to find any paths defined on the server that reference this device, then issue the following command to update the path information:</p> <pre>UPDATE PATH AUTODetect=Yes</pre>
<p>ANR8972E: Unable to find element number for drive <i>driveName</i> in library <i>libraryName</i></p>	<p>If the ELEment parameter is set to <code>AUTODetect</code> when defining the drive, Tivoli Storage Manager automatically gets the drive's element number. However, if the library does not provide an element number/serial number map, this message is issued.</p> <p>Perform the following steps to correct this:</p> <ol style="list-style-type: none"> 1. Determine the element number for this tape drive. 2. Issue the UPDATE DRIVE command to update the device element number.

Related concepts:

“SAN device mapping errors” on page 243

SAN device mapping hints and tips

SAN device discovery and device mapping are supported on Windows 2000, Windows 2003 (32 bits), AIX, Oracle Solaris, and Linux (except Linux zSeries®).

The following items illustrate the advantages of IBM Tivoli Storage Manager SAN device discovery and device mapping:

Tivoli Storage Manager can display all the devices on the SAN

The **QUERY SAN** server command shows all the devices seen by the Tivoli Storage Manager server via all the Fibre Channel host bus adapters (HBAs) installed on the system. The parameters shown are device type, vendor name, product model name, serial number, and the device name. If **FORMAT=DETAIL** is specified for the query, additional information such as World Wide Name (WWN), port, bus, target, and LUN are displayed. This information helps identify all the tape, disk, and Data Mover devices on the SAN. For AIX, the data mover is transparent and is not shown.

Tivoli Storage Manager can update the device path automatically when a device's path changes

Tivoli Storage Manager does not require persistent binding for the devices it sees through the HBA. Instead, the server uses the SNIA (Storage Networking Industry Association) HBAAPI to discover and obtain the serial number for all the devices on the SAN. It can also determine each device's path. By comparing a device's serial number recorded in the Tivoli Storage Manager database with the serial number obtained from the device in real time, a change in a device's path is detected. If the path was changed, SAN discovery automatically obtains the new path for the device. The Tivoli Storage Manager database is also updated with the new path information.

The HBAAPI wrapper library is the wrapper used by the Tivoli Storage Manager server to communicate with the SNIA HBAAPI. The HBAAPI wrapper library is installed in the same directory as the Tivoli Storage Manager executable file (unless the full path is given). The following list shows the HBA wrapper files that are shipped with the Tivoli Storage Manager server package (except on AIX):

- Windows: hbaapi.dll
- AIX: /usr/lib/libhbaapi.a (provided by AIX with HBAAPI installation)
- 32-bit Linux: libhbaapi32.so
- 64-bit Linux: libhbaapi64.so
- 32-bit Oracle Solaris: libhbaapi32.so
- 64-bit Oracle Solaris: libhbaapi64.so

If any of these files are missing, the "ANR1791W HBAAPI wrapper library xxxxxxxx failed to load or is missing." message is displayed.

Related concepts:

"SAN device mapping errors" on page 243

Disabling SAN device mapping:

Occasionally you must disable SAN device mapping to circumvent or isolate a problem when you are troubleshooting device problems.

Perform the following step to disable SAN device mapping and device discovery:

Issue the **setopt SANDISCOVER OFF** server command. The **setopt SANDISCOVERY** commands can be issued as many times as needed.

Tip: Another way to disable/enable SAN discovery is to enter the following option in the `dsmserv.opt` file:

SANDISCOVERY OFF disables SAN discovery.

SANDISCOVERY ON enables SAN discovery.

SANDISCOVERY ON is the default for the AIX, Linux, Oracle Solaris, and Windows platforms.

Platform-specific information:

When working on your SAN device mapping, it is important that you know your platform-specific information.

AIX The **QUERY SAN** command will NOT show any Gateway devices because Gateway devices are transparent to AIX.

Linux There are separate libraries, utilities, and other items for RHEL3U3. To run them you must also install an Emulex ioctl kernel module in addition to the Emulex driver. Make sure to load the Emulex driver before loading the ioctl module.

Emulex provides an Application kit for RHEL3. To find the Emulex application kit, go to the Emulex website and click on **Support**. Under **Choose your supplier from the following list**, select vendor **IBM**. A list of drivers and kits are available for you to download.

Tip: See the list of supported HBAs and required driver levels by operating system.

SAN device mapping errors

The errors that are most often generated during SAN device mapping can be related to SAN discovery, SAN device malfunction, libraries that are not valid, and other SAN-related issues.

ANR1745I: Unable to discover SAN devices. Function is busy.

This error message appears if there is another active SAN discovery.

The IBM Tivoli Storage Manager server is not able to perform SAN discovery. Retry again after the other SAN discovery is completed.

ANR1786W, ANR1787W or ANR1788W

You might see error messages ANR1786W, ANR1787W, or ANR1788W due to a problem with SAN discovery. The following three messages usually indicate that the HBA API library is not working in general:

- ANR1786W HBA API not able to get adapter name

- ANR1787W Not able to open adapter *adaperName*
- ANR1788W Not able to get the adapter attributes for *adapterName*

If the result is that the Tivoli Storage Manager server is unable to perform SAN discovery, go to http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager to verify that the host bus adapter (HBA) driver is up-to-date and at a supported level.

ANR1789W Get HBA target mapping failed

Error message ANR1789W is the most common HBAAPIn error on the SAN.

“Get HBA target mapping failed” means that the HBA encountered an error while gathering device mapping information by sending various SCSI commands.

Verify that all SAN devices are working properly (for example, a SAN Data Gateway could be hung and need to be rebooted). If all devices appear functional, verify the firmware of device on the SAN, as well as HBA driver, are at the appropriate levels. If the result is that the Tivoli Storage Manager server is not able to perform SAN discovery, go to http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager to verify that the HBA driver is up-to-date and at a supported level.

Tip: For IBM tape devices, make sure the latest firmware is installed. Firmware prior to 4772 for IBM 3580 tape devices causes problems with Qlogic HBAAPI.

ANR1790W SAN discovery failed

Error message ANR1790W is a general message indicating that the HBAAPI function failed and cannot be performed SAN discovery.

Verify that all SAN devices are working properly (for example, a SAN Data Gateway could be hung and need to be rebooted). If all devices appear functional, verify that the firmware of device on the SAN, as well as the HBA driver, are at the appropriate levels.

Tip: For IBM tape devices, make sure the latest firmware is installed. Firmware prior to 4772 for IBM 3580 tape devices causes problems with Qlogic HBAAPI.

ANR1791W HBAAPI wrapper library xxxxx failed to load or is missing

The HBAAPI wrapper library is used by the Tivoli Storage Manager server to communicate with the SNIA HBAAPI.

The HBAAPI wrapper libraries are in the same directory as the Tivoli Storage Manager executable file (unless full path is given as shown below). The following list shows the HBA wrapper files shipped with the Tivoli Storage Manager server package (except on AIX and Linux zSeries). Error message ANR1791W indicates that the HBAAPI wrapper file is either missing or could not be loaded by the Tivoli Storage Manager. Verify that the wrapper file is in the same directory as the Tivoli Storage Manager executable file. The HBAAPI wrapper library files are shown in the following list:

- Windows: hbaapi.dll
- AIX: /usr/lib/libhbaapi.a (provided by AIX with HBAAPI installation)
- 32-bit Linux: libhbaapi32.so

- 64-bit Linux: libhbaapi64.so
- 64-bit Oracle Solaris: libhbaapi64.so

The result is that the Tivoli Storage Manager server is not able to perform SAN discovery.

ANR1792W HBAAPI vendor library failed to load or is missing

Error message ANR1792W indicates that the vendor's library file failed to load. Verify the validity of the library files.

AIX, HP-UX, Linux, or Oracle Solaris systems (except on Linux zSeries) store their HBAAPI libraries in the location specified by the /etc/hba.conf file. Windows files are stored in the C:\winnt\system32 directory. The following examples are of vendor library files:

- C:\winnt\system32\qlsdm.dll (QLogic's Windows file)
- /usr/lib/libHBAAPI.a (Emulex's AIX file)
- /usr/lib/libqlsdm.so (Qlogic's Linux file)
- /usr/lib/libemulexhbaapi.so (Emulex's Linux 32-bit file)
- /usr/lib64/libemulexhbaapi.so (Emulex's Linux 64-bit file)
- /usr/lib/libqlsdm.so (Qlogic's Oracle Solaris file)
- /opt/JNIsnia/Solaris/Jni/64bit/JniHbaLib.so (JNI's Oracle Solaris file)

The result is that the Tivoli Storage Manager server is not able to perform SAN discovery.

ANR1793W Tivoli Storage Manager SAN discovery is not supported on this platform or this version of OS

Error message ANR1793W is only displayed if the Tivoli Storage Manager attempts a SAN device mapping or device discovery operation on an unsupported operating system. The following platforms are not currently supported by SAN device mapping or device discovery:

- HP-UX
- 64-bit Windows 2003
- AIX versions that are not 52L or 53A. Support for SAN device mapping and device discovery on AIX requires either version 52L (fileset level of 5.2.0.50) or 53A (fileset level of 5.3.0.10) or higher.

The result is that the Tivoli Storage Manager server is not able to perform SAN discovery.

ANR1794W Tivoli Storage Manager SAN discovery is disabled by options

Error message ANR1794W indicates that the SAN discovery on the Tivoli Storage Manager server is disabled.

The SAN discovery can be disabled or enabled by issuing the following server commands:

setopt SANDISCOVERY OFF and setopt SANDISCOVERY PASSIVE

These two commands disable the SAN discovery. The Tivoli Storage Manager server is not able to correct the device path automatically if the path was changed. This command only has to be issued one time.

The difference between these two commands is that **SANDISCOVERY OFF** polls the device and marks the inactive path off-line. **SANDISCOVERY PASSIVE** does not poll the device and does not mark the inactive path off-line.

setopt SANDISCOVERY ON

This command enables the SAN discovery. The **SETOPT SANDISCOVERY ON** command can be issued as many times as necessary.

Another way to disable/enable SAN discovery is to put the following option in the `dsmserv.opt` file:

SANDISCOVERY OFF or SANDISCOVERY PASSIVE

These two commands can disable the SAN discovery.

SANDISCOVERY ON

This command enables the SAN discovery.

SANDISCOVERY is defaulted to ON for AIX, Linux, Oracle Solaris, and Windows platforms.

Go to http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager to verify the platform/HBA vendor/driver level support level prior to setting **SANDISCOVERY ON** to enable the SAN discovery.

ANR2034E QUERY SAN: No match found using this criteria

Error message ANR2034E is issued when the Tivoli Storage Manager server tries to collect configuration information for the SAN and finds nothing.

The result is that the Tivoli Storage Manager server is unable to perform SAN discovery.

The following are possible reasons for not finding information about the SAN:

- The platform or OS level is unsupported.
- This environment is not a SAN environment.
- There may be a problem with the SAN.
- HBA API may return the zero value of the number of HBAs on the system.
- HBA API may return the zero value of the number of devices on the system.

Perform the following tasks to find the SAN configuration information:

- Check the fibre-channel HBA driver and make sure it is installed and enabled.
- Check the HBA driver level to make sure that it is up-to-date.
- Use the HBA vendor's utility to check for any reported fibre-channel link problems.
- Uninstall and then install the HBA driver. If there is an issue with the HBA configuration, device driver, or compatibility, sometimes uninstalling and re-installing it corrects the problem.
- Check the fibre-channel cable connection to the HBA.
- Check the fibre-channel cable connection from the HBA to the SAN device (switch, data gateway, or other device).
- Check the Gigabit Inter-phase Converter (GBIC).
- On the SAN device (switch, data gateway, or other device) try a different target port. Sometimes the SAN devices may have a specific port failure.

- Halt the Tivoli Storage Manager server, restart the machine, and restart the server. If there were configuration changes in the SAN, sometimes the operating system, device driver, or HBA requires a machine restart before they can communicate with the SAN.
- Recycle the destination port on the SAN device.
- Re-seat the HBA card.
- Replace the HBA.

ANR8226E Error detecting version of HBA-API library

Error message ANR8226E is only displayed for AIX.

The server attempted to determine the level of the `devices.common.IBM.fc.hba-api` fileset and encountered an error. Error message ANR8226E indicates that an error occurred while trying to detect the HBA-API libraryFileset version on AIX.

The result is that the Tivoli Storage Manager server is not able to perform SAN discovery.

AIX

ANR8227E Fileset `devices.common.IBM.fc.hba-api` is not at the required level

Due to problems in AIX HBAAPI code, the minimum fileset `devices.common.IBM.fc.hba-api` level needed for successful SAN discovery are shown in the following list:

- AIX52 - Need 5.2.0.50
- AIX53 - Need 5.3.0.10

The server specified that the file set `devices.common.IBM.fc.hba-api` is at a level that is incompatible with Tivoli Storage Manager operations. Install the latest maintenance for this file set if you use SAN devices.

The result is that the Tivoli Storage Manager server is not able to perform SAN discovery.

Related reference:

“SAN device mapping hints and tips” on page 242

SAN devices are missing from the display of `QUERY SAN` server command:

The possible reasons for the `QUERY SAN` server command not displaying all the devices can be due to configuration or vendor support issues.

Refreshing the SAN configuration:

The `QUERY SAN` server command might not be displaying all the devices because of the SAN configuration.

You might have to refresh the SAN because the configuration was changed (add/remove device) and the system configuration needs to be updated.

Update configuration on AIX:

For IBM devices:

Issue the `cfgmgr` command to configure new devices and view the new configuration. The special file name for IBM tape devices (not

the IBM Tivoli Storage Manager devices) is /dev/rmtX for tape drives and /dev/smcX for medium changers.

Tip: Special file name: /dev/rmt0, /dev/smc0

For the Tivoli Storage Manager devices:

To update the special files, use **smitty > devices > Tivoli Storage Manager Devices > remove all defined devices**, then **discover devices supported by Tivoli Storage Manager**. The special file name is /dev/mtX for tape drives and /dev/lbX for medium changers.

Tip: Special file name: /dev/mt0, /dev/lb0

Alternatively, you can reinstall the IBM device driver. Tivoli Storage Manager device driver updates all the current special file name.

Update configuration on Windows:

With the plug and play, the Windows registry is updated and the device name might change without the need to restart the computer or have the device driver's involvement. The Tivoli Storage Manager server detects the change in a special file name and updates the new special file name when it accesses the tape devices (during server initialization or normal operation). The correct device name is updated in the Tivoli Storage Manager database. The special file name is /dev/mtA.B.C.D for both Tivoli Storage Manager devices and IBM devices, and /dev/lbA.B.C.C for both Tivoli Storage Manager devices and IBM medium changers. The special file name TapeX is only for IBM tape drives and ChangerX is only for IBM medium changes.

Tip: Special file name: mt0.1.0.0, lb0.0.1.0, Tape0, and Changer0.

Update configuration on Linux:

The host bus adapter (HBA) gets the most up-to-date configuration information as a result of the RSCN. Sometimes, the computer must be restarted to be able to pick up the configuration changes.

For IBM devices:

Issue the **lin_taped** command to reconfigure devices. The device information can be retrieved from the /proc/scsi/IBMTape file for tape devices and /proc/scsi/IBMchanger file for medium changers. The special file name is /dev/IBMTapeX for tape devices and /dev/IBMChangerX for medium changers.

Tip: Special file name: /dev/IBMTape0, /dev/IBMChanger0

For the Tivoli Storage Manager devices:

Users can issue **autoconf**, the Tivoli Storage Manager device driver auto configure script. This script resides in the /opt/tivoli/tsm/devices/bin directory (or in the same directory as the tsm SCSI file) to be able to configure devices and get all the current special file names and device information. The device special file name is /dev/mtX for tape devices and /dev/lbX for medium changers.

Tip: Special file name: dev/tsmscsi/mt0, /dev/tsmscsi/lb0

Alternatively, you can reinstall the IBM device driver. Tivoli Storage Manager device driver updates all the current special file names.

With the Linux pass-thru device driver for the Tivoli Storage Manager devices, the HBA driver and the generic driver must be reloaded to get all the current special file names. You have to run the autoconf script so that the Tivoli Storage Manager device driver can create configuration files (`/dev/tsmcsci/lbinfo` and `/dev/tsmcsci/mtinfo`). These files are used by the Tivoli Storage Manager server to create the special file names after each SAN discovery.

32 bits (Linux xSeries®)

Ensure that the HBAAPI wrapper library `libhbaapi32.so` is in the same directory as `dsmserv` or in the `/opt/tivoli/tsm/server/bin` directory.

64 bits (Linux pSeries)

Ensure that the HBAAPI wrapper library `libhbaapi64.so` is in the same directory as `dsmserv` or in the `/opt/tivoli/tsm/server/bin` directory.

64 bits (Linux zSeries)

Ensure that the pseudo-HBAAPI wrapper library `libhbaapi64.so` is in the same directory as `dsmserv` or in the `/opt/tivoli/tsm/server/bin` directory. The wrapper library, `libhbaapi64.so`, is a link to the `/usr/lib64/libzfcphbaapi.so` file.

Update the configuration on Oracle Solaris:

The HBA gets the most up-to-date configuration information as result of the RSCN. Most of the time, the computer must be restarted to pick up the configuration changes. Reinstall the IBM device driver package or the Tivoli Storage Manager device driver package and run `autoconf`. Or issue the `rem_drv` and `add_drv` commands to reconfigure devices and update the special file name:

For IBM devices:

The IBM device driver has already completed configuring the devices after installing the device driver package. The `/opt/IBMtape/tapelist -l` command can show all IBM device information on the system. The special file name is `/dev/rmt/Xst` for tape devices and `/dev/rmt/Xsmc` for medium changers.

Tip: Special file name: `/dev/rmt/0st`, `/dev/rmt/0smc`

For Tivoli Storage Manager devices:

Modify `/usr/kernel/drv/mt.conf` and `/usr/kernel/drv/lb.conf` and ensure that `name="mt" parent="pseudo" instance=16383`; and `name="lb" parent="pseudo" instance=16383`; are not commented out in the `mt.conf` and `lb.conf` files. Ensure that each entry for a device in `/usr/kernel/drv/lb.conf` and `/usr/kernel/drv/mt.conf` is correct. After running `autoconf`, device information can be found in the `mtinfo` and `lbinfo` files in the `/opt/tivoli/tsm/devices/bin` directory or the same directory as the `autoconf` script. The special file name is `/dev/rmt/Xmt` for tape drives and `/dev/rmt/Xlb` for medium changers.

Tip: Special file name: `/dev/rmt/0mt`, `/dev/rmt/0lb`

Important:

Ensure that the two pseudo devices, `/devices/pseudo/mt@16383:tsmmtctl` and `/devices/pseudo/lb@16383:tsmlbctl`, are in the `/devices/pseudo` directory.

Ensure that the pseudo device special files, `/dev/tsmmtctl` and `/dev/tsmlbctl` are linked to their corresponding pseudo devices in `/devices/pseudo/mt@16383:tsmmtctl` and `/devices/pseudo/lb@16383:tsmlbctl`.

Resolving configuration problems that cause SAN device absence:

The possible reasons for the **QUERY SAN** server command not displaying all the devices can be due to a configuration problem with the HBA hardware, HBA driver level, or operating system level.

Perform the following steps to resolve configuration issues:

1. Go to http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager. Verify the platform/HBA vendor/driver level support level to make sure the HBA driver level and operating system level are compatible and supported by IBM Tivoli for SAN discovery.
2. Use the HBA vendor utility to check to see if the device can be seen by the HBA. If the device is not seen by the HBA, the device might not be connected. Check the Fibre Channel or SCSI cable. If the device is seen by the HBA, check the HBA driver version. This driver version might have problems with the HBA API.

Verifying vendor support for any particular device in the SAN:

Many devices or combinations of devices might not be supported in a given storage area network (SAN). These limitations arise from the ability of a given vendor to certify their device using Fibre Channel Protocol.

For a given device, verify with the device vendor that it is supported in a SAN environment. Vendor support includes all hardware associated with the SAN, which means verifying that this device is supported with the vendors of the HBAs, hubs, gateways, and switches that make up the SAN environment.

NDMP filer-to-Tivoli Storage Manager server operation hints and tips

IBM Tivoli Storage Manager defaults to the standard network data management protocol (NDMP) control port of 10000. If this port is in use by another application (such as a second Tivoli Storage Manager server), all filer-to-server operations fail.

To avoid conflicts with other applications, use the `NDMPCONTROLPORT` server option to specify a different port for your server.

During filer-to-server operations, Tivoli Storage Manager uses the following items:

- Up to two additional TCP/IP ports.
- A control port used internally by Tivoli Storage Manager during both backup and restore operations.
- A data port during NDMP backup operations to a Tivoli Storage Manager native storage pool

The data port is an ephemeral port that is acquired at the beginning of NDMP backup operations to a Tivoli Storage Manager native storage pool. If a port is not available, an error message is issued and backup of NAS devices to Tivoli Storage Manager native pools is not possible. To avoid conflicts with other applications, you can control which port is acquired for use as the data port during NDMP backup operations by setting the `NDMPPORTRANGELOW` and `NDMPPORTRANGEHIGH` server options. Note that a data port is not needed by the Tivoli Storage Manager server for NAS restores from Tivoli Storage Manager native pools.

Firewall considerations with NDMP filer-to-Tivoli Storage Manager server backup and restore

A firewall might prevent the network-attached storage (NAS) file server from contacting the Tivoli Storage Manager server on the acquired data port during NAS backup operations to a native storage pool. If you must modify the data port selected by the Tivoli Storage Manager server, use the `NDMPPORTRANGELOW` and `NDMPPORTRANGEHIGH` server options.

A firewall might prevent the Tivoli Storage Manager server from contacting the NAS file server on the configured data port during NAS restore operations from a native storage pool. If a firewall prevents Tivoli Storage Manager from accessing the NAS file server, the outbound connection from Tivoli Storage Manager fails.

Resolving SCSI device problems

Tape drives and libraries might report information back to IBM Tivoli Storage Manager about the error encountered. This information is reported in one or more of the messages.

If messages ANR8300, ANR8301, ANR8302, ANR8303, ANR8943, or ANR98944 are issued, the data that Tivoli Storage Manager reports from these devices might help to determine the steps needed to resolve the problem. Generally, when the Tivoli Storage Manager server reports device sense data using these messages, the problem is typically with the device, the connection to the device, or some other related issue outside of Tivoli Storage Manager.

Using the information reported in Tivoli Storage Manager message ANR8300, ANR8301, ANR8302, ANR8303, ANR8943, or ANR8944, refer to the Tivoli Storage Manager Messages manual. This appendix documents information about standard errors that might be reported by any SCSI device. You can also use this information with documentation provided by the vendor for the hardware to help determine the cause and resolution for the problem.

Resolving sequential media volume (tape) errors through messages ANR0542W or ANR8778W

Problems occurring with sequential media volumes can be revealed through error messages ANR0542W and ANR8778W.

ANR0542W Retrieve or restore failed for session *sessionNumber* for node *nodeName* - storage media inaccessible

Error message ANR0542W is often related to an issue with the drive or connection to the drive that was selected to read this tape volume.

Perform the following steps to verify that IBM Tivoli Storage Manager can access this volume:

- Issue the `QUERY LIBVOL libraryName volumeName` command.
- For a 349X library, issue the `mtlib -l /dev/lmcp0 -qV volumeName` command. The device is typically `/dev/lmcp0`, but if it is different, then substitute the correct library manager control point device.

The following steps might possibly resolve this problem:

1. If `mtlib` does not report this volume, then it appears that this volume is out of the library. In this case, put the volume back into the library.
2. If the volume is not reported by `QUERY LIBVOL`, then the server does not know about this volume in the library. Issue the **CHECKIN LIBVOL** command to synchronize the library inventory in the server with the volumes that are actually in the tape library.
3. If both commands successfully report this volume, then the cause is likely a permanent or intermittent hardware error. There might be an error with the drive itself or an error with the connection to the drive. In either case, review the system error logs and contact the vendor of the hardware to resolve the problem.

ANR8778W Scratch volume changed to private status to prevent re-access

Review the activity log messages to determine the cause of the problem involving this scratch volume. Also, review the system error logs and device error logs for an indication that there was a problem with the drive used to try to write to this scratch volume.

If this error was caused by a drive requiring cleaning or some other hardware-specific issue that was resolved, any volumes that were set to private status as a result of this might be reset to scratch by issuing the `AUDIT LIBRARY libraryName` command.

Appendix A. Getting call stack information from a core file

You can use the sample `gt` script shell provided here to get the call stack for each running thread from a core file.

The input parameters are the path/name of the executable file (default `./dsmserv`) and the path/name of the core file (default `./dsmcore`). The output file is `dsm_gdb.info`.

Restriction: Files named `dsm_gdb.cmd` and `dsm_gdb.info` are overwritten when you run this script.

```
#!/bin/ksh
#
# If you see the following error:
# ./dsm_gdb.cmd:9: Error in source command file:
# No symbol table is loaded. Use the "file" command.
# then comment out the line that prints buildStringP
#
# if you see other errors, you're on your own ...
exe=${1:-"./dsmserv"} # get parm 1 (executable file path/name), set default
core=${2:-"./dsmcore"} # get parm 2 (core file path/name), set default
echo " "
# look for the executable file ... quit if not found
if [[ -f $exe ]]; then
echo "using executable file:" $exe
else
echo "didn't find executable file ("$exe") ... exiting"
exit
fi
# look for the core file, if not found, look for ./core ... quit if not found
if [[ -f $core ]]; then
echo "using core file:" $core
else
if [[ -f ./core ]]; then
echo "didn't find core file ("$core") but found ./core ... renaming to" $core
mv ./core $core
echo "using core file:" $core
else
echo "didn't find core file ("$core") ... exiting"
exit
fi
fi
echo " "
# make gdb command file to get thread info
nl="\0134\0156" # octal codes for \n (so echo won't think it's \n)
echo "# dsm gdb command file" >|dsm_gdb.cmd
echo "define doit" >>dsm_gdb.cmd
echo "info registers" >>dsm_gdb.cmd # show register values
echo "echo" $nl >>dsm_gdb.cmd
echo "where" >>dsm_gdb.cmd # show function traceback
echo "echo" $nl"======"$nl >>dsm_gdb.cmd
echo "end" >>dsm_gdb.cmd
echo "echo" $nl"======"$nl$nl >>dsm_gdb.cmd
echo "x/s buildStringP" >>dsm_gdb.cmd
echo "echo" $nl"======"$nl$nl >>dsm_gdb.cmd
echo "info threads" >>dsm_gdb.cmd # show thread info
echo "echo" $nl"======"$nl >>dsm_gdb.cmd
echo "thread apply all doit" >>dsm_gdb.cmd
echo "quit" >>dsm_gdb.cmd
echo "invoking gdb to get thread info (watch for errors) ..."
echo "if you see:"
```

```
echo ". warning: The shared libraries were not privately mapped; setting a"
echo ". breakpoint in a shared library will not work until you rerun the program"
echo "that's ok."
echo "if you see:"
echo ". ./dsm_gdb.cmd:x: Error in source command file:"
echo "then type 'quit', edit this script, and read the comments at the top"
gdb -se $exe -c $core -x ./dsm_gdb.cmd >|dsm_gdb.info
rm dsm_gdb.cmd # done with this now
exit
```

Appendix B. Installing and running the tsmdiag utility

The IBM Tivoli Storage Manager diagnostic utility (tsmddiag) speeds up the process for gathering data that is considered valuable to assist in diagnosing a problem caused by a server component.

The tsmdiag utility collects a set of data by default. This data is specified in a configuration file called `configfile.txt`. If diagnosing a problem from a Tivoli Storage Manager component and more data is available, modify the tsmdiag configuration file so that such data is collected by the tsmdiag utility.

A strong suggestion is that users run tsmdiag on the system where the Tivoli Storage Manager component with a problem is installed *before* submitting a problem report to Tivoli Storage Manager Support. Submit the data collected by tsmdiag along with the problem report.

Attention: The tsmdiag utility can inadvertently delete a directory if you specify a system directory when using the `-RESULTS` option. The deletion occurs without a warning.

Table 35. Supported operating system environments and components

Operating System Environment	Tivoli Storage Manager Component
AIX	Server Storage Agent Device Driver
HP-UX	Server Storage Agent Device Driver
Linux	Server Storage Agent Device Driver
Solaris	Server Storage Agent Device Driver
Windows	Server Storage Agent Device Driver Management Console

Tip: The `tsmdiag` utility is not installed with the Tivoli Storage Manager device driver, or with the Tivoli Storage Manager Management Console. If these components are installed on the system where `tsmdiag` is running, `tsmdiag` collects data for these components.

Where to find and how to use the Tivoli Storage Manager diagnostic utility

Tivoli Storage Manager is equipped with the `tsmdiag` utility, with supported Tivoli Storage Manager components found under the component installation directory. The `tsmdiag` utility can be found in the following locations:

- **Windows** Windows: `C:\Program Files\Tivoli\TSM\server\tsmdiag`
- **AIX** **HP-UX** **Linux** **Solaris** All other platforms:
`/opt/tivoli/tsm/server/bin/tsmdiag`

Usage:

`tsmdiag [options]`

Where options can be any combination of the following:

Option	Description	Default Value
<code>-id adminName</code>	Tivoli Storage Manager server administrator ID	admin
<code>-pa adminPwd</code>	Tivoli Storage Manager server administrator password	admin
<code>-tcpserveraddress ipAddress</code>	Tivoli Storage Manager server TPC/IP name or address	localhost
<code>-tcpport portNumber</code>	Tivoli Storage Manager server TCP/IP port	1500
<code>-results resultsDir</code>	Directory for created files	(\$cwd)/results
<code>-v</code>	Verbosely list activity processed	non-verbose
<code>-?</code>	Display usage information	N/A
<code>-i instanceDir Server</code>	Specifies the Tivoli Storage Manager server instance directory	current directory

Examples:

```
./tsmdiag
./tsmdiag -v
./tsmdiag -id tsmadmin -pa pwd4u -results /home/tsmdiag/results_oct29
```

Where to find the data collected by `tsmdiag`

After the `tsmdiag` utility has completed running, go to either the default results directory, `$CWD/results`, or the directory that you specified through the `-results` command-line option.

There is a file there called `tsmdiag_results.tar` (for AIX, HP-UX, Linux, or Oracle Solaris environments) and `tsmdiag_results.zip` (for Windows environments). Submit this file with the PMR.

Default data collected by `tsmdiag`

The default data collected by the `tsmdiag` utility is specified in the `configfile.txt` configuration file. More data can be added to this file as needed.

Each line in the configuration file consists of five fields, each separated by a comma. A line that begins with an asterisk(*) is considered a comment and is ignored by the `tsmdiag` utility.

The first field of a line in the configuration file denotes the operating system where `tsmdiag` is run. Valid platforms are Linux, AIX, HP, Oracle, and Windows. The second field in the line denotes an action. The valid actions are in the following table:

Action Name	Action Performed
SYSCOMMAND	Run a command on the system.
COPY	Copy a specified file.
COMPCOMMAND	Run a command on a specified Tivoli Storage Manager component.
REGISTRY	Get the value for a given registry key. (Windows only.)

The third field denotes the Tivoli Storage Manager component (as in a component for which `tsmdiag` collects data).

Valid components are:

- SERVER
- STAGENT
- BACLIENT (AIX, HP-UX, Linux, or Oracle Solaris only)
- MGMTCONSOLE (Windows only)
- DEVDRIVER
- SYSTEM
- CLIENAPI (AIX, HP-UX, Linux, or Oracle Solaris only)

The fourth field denotes the data that the action acts upon. The data field can be a file name, a system command, a command for a Tivoli Storage Manager component, a registry key, and so on. The action determines which data is used.

The last field is optional and denotes the location of where the collected data is to be stored.

Appendix C. IBM Global Security Kit return codes

The latest version of the Tivoli Storage Manager server uses the IBM Global Security Kit (GSKit) Version 8.0.14.14 for SSL (Secure Sockets Layer). The V6.4 backup-archive client uses GSKit Version 8.0.14.14 for SSL, which allows processing between the Tivoli Storage Manager server and the backup-archive client. Some messages issued for SSL processing include GSKit return codes.

GSKit is automatically installed or updated during Tivoli Storage Manager installation and provides the following libraries:

- GSKit SSL
- GSKit Key Management API
- IBM Crypto for C (ICC)

The `tsmdiag` utility reports the GSKit level installed on your system, or you can use one of the following methods:

- For Windows, issue the following commands:

```
regedit /e gskitinfo.txt "HKEY_LOCAL_MACHINE\software\ibm\gsk8\  
notepad gskitinfo.txt
```

CAUTION:

You can damage the system registry if you use `regedit` incorrectly.

- For the 64-bit AIX server, issue the following command from the command line:
`gsk8ver_64`

See Table 36 for the GSKit SSL return codes.

The Tivoli Storage Manager server uses the GSKit Key Management API to automatically create the key management database and Tivoli Storage Manager server private and public keys. Some messages issued for this processing might include GSKit Key Management return codes. See Table 37 on page 263 for the key management return codes.

Table 36. IBM Global Security Kit SSL general return codes

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000000	0	GSK_OK	The task completed successfully. Issued by every function call that completes successfully.
0x00000001	1	GSK_INVALID_HANDLE	The environment or SSL handle is not valid. The specified handle was not the result of a successful <code>open()</code> function call.
0x00000002	2	GSK_API_NOT_AVAILABLE	The dynamic link library (DLL) was unloaded and is not available (occurs on Microsoft Windows systems only).
0x00000003	3	GSK_INTERNAL_ERROR	Internal error. Report this error to IBM Software Support.
0x00000004	4	GSK_INSUFFICIENT_STORAGE	Insufficient memory is available to perform the operation.

Table 36. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000005	5	GSK_INVALID_STATE	The handle is not in a valid state for operation, such as performing an <code>init()</code> operation on a handle twice.
0x00000006	6	GSK_KEY_LABEL_NOT_FOUND	Specified key label not found in key file.
0x00000007	7	GSK_CERTIFICATE_NOT_AVAILABLE	Certificate not received from partner.
0x00000008	8	GSK_ERROR_CERT_VALIDATION	Certificate validation error.
0x00000009	9	GSK_ERROR_CRYPTO	Error processing cryptography.
0x0000000a	10	GSK_ERROR_ASN	Error validating ASN fields in certificate.
0x0000000b	11	GSK_ERROR_LDAP	Error connecting to user registry.
0x0000000c	12	GSK_ERROR_UNKNOWN_ERROR	Internal error. Report this error to IBM Software Support.
0x00000065	101	GSK_OPEN_CIPHER_ERROR	Internal error. Report this error to IBM Software Support.
0x00000066	102	GSK_KEYFILE_IO_ERROR	I/O error reading the key file.
0x00000067	103	GSK_KEYFILE_INVALID_FORMAT	The key file does not have a valid internal format. Recreate key file.
0x00000068	104	GSK_KEYFILE_DUPLICATE_KEY	The key file has two entries with the same key.
0x00000069	105	GSK_KEYFILE_DUPLICATE_LABEL	The key file has two entries with the same label.
0x0000006a	106	GSK_BAD_FORMAT_OR_INVALID_PASSWORD	The key file password is used as an integrity check. Either the key file has become corrupted or the password ID is incorrect.
0x0000006b	107	GSK_KEYFILE_CERT_EXPIRED	The default key in the key file has an expired certificate.
0x0000006c	108	GSK_ERROR_LOAD_GSKLIB	An error occurred loading one of the GSK dynamic link libraries. Be sure GSK was installed correctly.
0x0000006d	109	GSK_PENDING_CLOSE_ERROR	Indicates that a connection is trying to be made in a GSK environment after the <code>GSK_ENVIRONMENT_CLOSE_OPTIONS</code> was set to <code>GSK_DELAYED_ENVIRONMENT_CLOSE</code> and <code>gsk_environment_close()</code> function was called.
0x000000c9	201	GSK_NO_KEYFILE_PASSWORD	Neither the password nor the stash-file name was specified, therefore, the key file could not be initialized.
0x000000ca	202	GSK_KEYRING_OPEN_ERROR	Unable to open the key file. Either the path was specified incorrectly or the file permissions did not allow the file to be opened.
0x000000cb	203	GSK_RSA_TEMP_KEY_PAIR	Unable to generate a temporary key pair. Report this error to IBM Software Support.
0x000000cc	204	GSK_ERROR_LDAP_NO_SUCH_OBJECT	A User Name object was specified that is not found.

Table 36. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000cd	205	GSK_ERROR_LDAP_INVALID_CREDENTIALS	A Password used for an LDAP (lightweight directory access protocol) query is not correct.
0x00000ce	206	GSK_ERROR_BAD_INDEX	An index into the Fail Over list of LDAP servers was not correct.
0x00000cf	207	GSK_ERROR_FIPS_NOT_SUPPORTED	This installation of GSKit does not support FIPS mode of operation.
0x0000012d	301	GSK_CLOSE_FAILED	Indicates that the GSK environment close request was not properly managed. Cause is most likely due to a <code>gsk_secure_socket*()</code> command being attempted after a <code>gsk_close_environment()</code> call.
0x00000191	401	GSK_ERROR_BAD_DATE	The system date was not set to a valid value.
0x00000192	402	GSK_ERROR_NO_CIPHERS	Neither SSLv2 nor SSLv3 is enabled.
0x00000193	403	GSK_ERROR_NO_CERTIFICATE	The required certificate was not received from partner.
0x00000194	404	GSK_ERROR_BAD_CERTIFICATE	The received certificate was formatted incorrectly.
0x00000195	405	GSK_ERROR_UNSUPPORTED_CERTIFICATE_TYPE	The received certificate type was not supported.
0x00000196	406	GSK_ERROR_IO	An I/O error occurred on a data read or write operation.
0x00000197	407	GSK_ERROR_BAD_KEYFILE_LABEL	The specified label in the key file could not be found.
0x00000198	408	GSK_ERROR_BAD_KEYFILE_PASSWORD	The specified key file password is incorrect. The key file could not be used. The key file also might be corrupt.
0x00000199	409	GSK_ERROR_BAD_KEY_LEN_FOR_EXPORT	In a restricted cryptography environment, the key size is too long to be supported.
0x0000019a	410	GSK_ERROR_BAD_MESSAGE	An incorrectly formatted SSL message was received from the partner.
0x0000019b	411	GSK_ERROR_BAD_MAC	The message authentication code (MAC) was not successfully verified.
0x0000019c	412	GSK_ERROR_UNSUPPORTED	Unsupported SSL protocol or unsupported certificate type.
0x0000019d	413	GSK_ERROR_BAD_CERT_SIG	The received certificate contained an incorrect signature.
0x0000019e	414	GSK_ERROR_BAD_CERT	Incorrectly formatted certificate received from partner.
0x0000019f	415	GSK_ERROR_BAD_PEER	Did not receive a valid SSL protocol from partner.
0x000001a0	416	GSK_ERROR_PERMISSION_DENIED	Report this error to IBM Software Support.
0x000001a1	417	GSK_ERROR_SELF_SIGNED	The self-signed certificate is not valid.

Table 36. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x000001a2	418	GSK_ERROR_NO_READ_FUNCTION	The read() failed. Report this error to IBM Software Support.
0x000001a3	419	GSK_ERROR_NO_WRITE_FUNCTION	The write() failed. Report this error to IBM Software Support.
0x000001a4	420	GSK_ERROR_SOCKET_CLOSED	The partner closed the socket before the protocol completed.
0x000001a5	421	GSK_ERROR_BAD_V2_CIPHER	The specified V2 cipher is not valid.
0x000001a6	422	GSK_ERROR_BAD_V3_CIPHER	The specified V3 cipher is not valid.
0x000001a7	423	GSK_ERROR_BAD_SEC_TYPE	Report this error to IBM Software Support.
0x000001a8	424	GSK_ERROR_BAD_SEC_TYPE_COMBINATION	Report this error to IBM Software Support.
0x000001a9	425	GSK_ERROR_HANDLE_CREATION_FAILED	The handle could not be created. Report this error to IBM Software Support.
0x000001aa	426	GSK_ERROR_INITIALIZATION_FAILED	Initialization failed. Report this internal error to service.
0x000001ab	427	GSK_ERROR_LDAP_NOT_AVAILABLE	Not able to access the specified user registry when validating a certificate.
0x000001ac	428	GSK_ERROR_NO_PRIVATE_KEY	The specified key did not contain a private key.
0x000001ad	429	GSK_ERROR_PKCS11_LIBRARY_NOTLOADED	A failed attempt was made to load the specified PKCS11 shared library.
0x000001ae	430	GSK_ERROR_PKCS11_TOKEN_LABELMISMATCH	The PKCS #11 driver failed to find the token specified by the caller.
0x000001af	431	GSK_ERROR_PKCS11_TOKEN_NOTPRESENT	A PKCS #11 token is not present in the slot.
0x000001b0	432	GSK_ERROR_PKCS11_TOKEN_BADPASSWORD	The password/pin to access the PKCS #11 token is not valid.
0x000001b1	433	GSK_ERROR_INVALID_V2_HEADER	The SSL header received was not a properly formatted SSLv2 header.
0x000001b2	434	GSK_CSP_OPEN_ERROR	Could not open the hardware-based cryptographic service provider. Either the CSP name is not specified correctly or a failed attempt was made to access the specified CSP certificate store.
0x000001b3	435	GSK_CSP_OPEN_ERROR	Some conflicting attributes for SSL operation were defined.
0x000001b4	436	GSK_CSP_OPEN_ERROR	The Microsoft Crypto API is only supported on Microsoft Windows 2000 with Service Pack 2 applied.
0x000001b5	437	GSK_CSP_OPEN_ERROR	System is running in IPv6 mode without setting a PEERID.
0x000001f5	501	GSK_INVALID_BUFFER_SIZE	The buffer size is negative or zero.
0x000001f6	502	GSK_WOULD_BLOCK	Used with nonblocking I/O. Refer to the nonblocking section for usage.
0x00000259	601	GSK_ERROR_NOT_SSLV3	SSLv3 is required for reset_cipher(), and the connection uses SSLv2.

Table 36. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x0000025a	602	GSK_MISC_INVALID_ID	A valid ID was not specified for the <code>gsk_secure_soc_misc()</code> function call.
0x000002bd	701	GSK_ATTRIBUTE_INVALID_ID	The function call does not have a valid ID. This might also be caused by specifying an environment handle when a handle for a SSL connection should be used.
0x000002be	702	GSK_ATTRIBUTE_INVALID_LENGTH	The attribute has a negative length, which is not valid.
0x000002bf	703	GSK_ATTRIBUTE_INVALID_ENUMERATION	The enumeration value is not valid for the specified enumeration type.
0x000002c0	704	GSK_ATTRIBUTE_INVALID_SID_CACHE	A parameter list that is not valid for replacing the SID cache routines.
0x000002c1	705	GSK_ATTRIBUTE_INVALID_NUMERIC_VALUE	When setting a numeric attribute, the specified value is not valid for the specific attribute being set.
0x000002c2	706	GSK_CONFLICTING_VALIDATION_SETTING	Conflicting parameters were set for additional certificate validation.
0x000002c3	707	GSK_AES_UNSUPPORTED	The AES cryptographic algorithm is not supported.
0x000002c4	708	GSK_PEERID_LENGTH_ERROR	The PEERID does not have the correct length.
0x000005dd	1501	GSK_SC_OK	
0x000005de	1502	GSK_SC_CANCEL	
0x00000641	1601	GSK_TRACE_STARTED	The trace started successfully.
0x00000642	1602	GSK_TRACE_STOPPED	The trace stopped successfully.
0x00000643	1603	GSK_TRACE_NOT_STARTED	No trace file was previously started, so it cannot be stopped.
0x00000644	1604	GSK_TRACE_ALREADY_STARTED	Trace file already started, so it cannot be restarted.
0x00000645	1605	GSK_TRACE_OPEN_FAILED	Trace file cannot be opened. The first parameter of <code>gsk_start_trace()</code> must be a valid full path file name.

Table 37. IBM Global Security Kit key management return codes

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000000	0	GSK_OK	The task completed successfully. This message is issued by every function call that completes successfully.
0x00000001	1	GSK_INVALID_HANDLE	The environment or SSL handle is not valid. The specified handle was not the result of a successful <code>open()</code> function call.
0x00000002	2	GSK_API_NOT_AVAILABLE	The DLL (dynamic link library) was unloaded and is not available (occurs on Microsoft Windows systems only).

Table 37. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000003	3	GSK_INTERNAL_ERROR	Internal error. Report this error to IBM Software Support.
0x00000004	4	GSK_INSUFFICIENT_STORAGE	Insufficient memory is available to perform the operation.
0x00000005	5	GSK_INVALID_STATE	The handle is in an incorrect state for operation, such as performing an init() operation on a handle twice.
0x00000006	6	GSK_KEY_LABEL_NOT_FOUND	Specified key label not found in key file.
0x00000007	7	GSK_CERTIFICATE_NOT_AVAILABLE	Certificate not received from partner.
0x00000008	8	GSK_ERROR_CERT_VALIDATION	Certificate validation error.
0x00000009	9	GSK_ERROR_CRYPTO	Error processing cryptography.
0x0000000a	10	GSK_ERROR_ASN	Error validating ASN fields in certificate.
0x0000000b	11	GSK_ERROR_LDAP	Error connecting to user registry.
0x0000000c	12	GSK_ERROR_UNKNOWN_ERROR	Internal error. Report this error to IBM Software Support.
0x00000065	101	GSK_OPEN_CIPHER_ERROR	Internal error. Report this error to IBM Software Support.
0x00000066	102	GSK_KEYFILE_IO_ERROR	I/O error reading the key file.
0x00000067	103	GSK_KEYFILE_INVALID_FORMAT	The key file has an internal format that is not valid. Recreate key file.
0x00000068	104	GSK_KEYFILE_DUPLICATE_KEY	The key file has two entries with the same key.
0x00000069	105	GSK_KEYFILE_DUPLICATE_LABEL	The key file has two entries with the same label.
0x0000006a	106	GSK_BAD_FORMAT_OR_INVALID_PASSWORD	The key file password is used as an integrity check. Either the key file has become corrupted or the password ID is incorrect.
0x0000006b	107	GSK_KEYFILE_CERT_EXPIRED	The default key in the key file has an expired certificate.
0x0000006c	108	GSK_ERROR_LOAD_GSKLIB	An error occurred while loading one of the GSK dynamic link libraries. Be sure GSK was installed correctly.
0x0000006d	109	GSK_PENDING_CLOSE_ERROR	This message indicates that a connection is trying to be made in a GSK environment after the GSK_ENVIRONMENT_CLOSE_OPTIONS was set to GSK_DELAYED_ENVIRONMENT_CLOSE and gsk_environment_close() function was called.
0x000000c9	201	GSK_NO_KEYFILE_PASSWORD	Neither the password nor the stash-file name was specified, so the key file could not be initialized.

Table 37. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x000000ca	202	GSK_KEYRING_OPEN_ERROR	Unable to open the key file. Either the path was specified incorrectly or the file permissions did not allow the file to be opened.
0x000000cb	203	GSK_RSA_TEMP_KEY_PAIR	Unable to generate a temporary key pair. Report this error to IBM Software Support.
0x000000cc	204	GSK_ERROR_LDAP_NO_SUCH_OBJECT	A User Name object was specified that is not found.
0x000000cd	205	GSK_ERROR_LDAP_INVALID_CREDENTIALS	A Password used for an LDAP query is not correct.
0x000000ce	206	GSK_ERROR_BAD_INDEX	An index into the Fail Over list of LDAP servers was not correct.
0x000000cf	207	GSK_ERROR_FIPS_NOT_SUPPORTED	This installation of GSKit does not support FIPS mode of operation.
0x0000012d	301	GSK_CLOSE_FAILED	Indicates that the GSK environment close request was not properly managed. Cause is most likely due to attempting a <code>gsk_secure_socket*()</code> command after a <code>gsk_close_environment()</code> call.
0x00000191	401	GSK_ERROR_BAD_DATE	The system date was set to a value that is not valid.
0x00000192	402	GSK_ERROR_NO_CIPHERS	Neither SSLv2 nor SSLv3 is enabled.
0x00000193	403	GSK_ERROR_NO_CERTIFICATE	The required certificate was not received from partner.
0x00000194	404	GSK_ERROR_BAD_CERTIFICATE	The received certificate was formatted incorrectly.
0x00000195	405	GSK_ERROR_UNSUPPORTED_CERTIFICATE_TYPE	The received certificate type was not supported.
0x00000196	406	GSK_ERROR_IO	An I/O error occurred on a data read-or-write operation.
0x00000197	407	GSK_ERROR_BAD_KEYFILE_LABEL	The specified label in the key file could not be found.
0x00000198	408	GSK_ERROR_BAD_KEYFILE_PASSWORD	The specified key file password is incorrect. The key file could not be used. The key file also might be corrupt.
0x00000199	409	GSK_ERROR_BAD_KEY_LEN_FOR_EXPORT	In a restricted cryptography environment, the key size is too long to be supported.
0x0000019a	410	GSK_ERROR_BAD_MESSAGE	An incorrectly formatted SSL message was received from the partner.
0x0000019b	411	GSK_ERROR_BAD_MAC	The MAC was not successfully verified.
0x0000019c	412	GSK_ERROR_UNSUPPORTED	Unsupported SSL protocol or unsupported certificate type.
0x0000019d	413	GSK_ERROR_BAD_CERT_SIG	The received certificate contained an incorrect signature.
0x0000019e	414	GSK_ERROR_BAD_CERT	Incorrectly formatted certificate received from partner.

Table 37. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x0000019f	415	GSK_ERROR_BAD_PEER	An SSL protocol that is not valid received from partner.
0x000001a0	416	GSK_ERROR_PERMISSION_DENIED	Report this error to IBM Software Support.
0x000001a1	417	GSK_ERROR_SELF_SIGNED	The self-signed certificate is not valid.
0x000001a2	418	GSK_ERROR_NO_READ_FUNCTION	The read() failed. Report this error to IBM Software Support.
0x000001a3	419	GSK_ERROR_NO_WRITE_FUNCTION	The write() failed. Report this error to IBM Software Support.
0x000001a4	420	GSK_ERROR_SOCKET_CLOSED	The partner closed the socket before the protocol completed.
0x000001a5	421	GSK_ERROR_BAD_V2_CIPHER	The specified V2 cipher is not valid.
0x000001a6	422	GSK_ERROR_BAD_V3_CIPHER	The specified V3 cipher is not valid.
0x000001a7	423	GSK_ERROR_BAD_SEC_TYPE	Report this error to IBM Software Support.
0x000001a8	424	GSK_ERROR_BAD_SEC_TYPE_COMBINATION	Report this error to IBM Software Support.
0x000001a9	425	GSK_ERROR_HANDLE_CREATION_FAILED	The handle could not be created. Report this error to IBM Software Support.
0x000001aa	426	GSK_ERROR_INITIALIZATION_FAILED	Initialization failed. Report this internal error to service.
0x000001ab	427	GSK_ERROR_LDAP_NOT_AVAILABLE	Unable to access the specified user registry when validating a certificate.
0x000001ac	428	GSK_ERROR_NO_PRIVATE_KEY	The specified key did not contain a private key.
0x000001ad	429	GSK_ERROR_PKCS11_LIBRARY_NOTLOADED	A failed attempt was made to load the specified PKCS11 shared library.
0x000001ae	430	GSK_ERROR_PKCS11_TOKEN_LABELMISMATCH	The PKCS #11 driver failed to find the token specified by the caller.
0x000001af	431	GSK_ERROR_PKCS11_TOKEN_NOTPRESENT	A PKCS #11 token is not present in the slot.
0x000001b0	432	GSK_ERROR_PKCS11_TOKEN_BADPASSWORD	The password/pin to access the PKCS #11 token is incorrect.
0x000001b1	433	GSK_ERROR_INVALID_V2_HEADER	The SSL header received was not a properly formatted SSLv2 header.
0x000001b2	434	GSK_CSP_OPEN_ERROR	Could not open the hardware-based cryptographic service provider (CSP). Either the CSP name is not specified correctly or a failed attempt was made to access the specified CSP certificate store.
0x000001b3	435	GSK_CSP_OPEN_ERROR	Some conflicting attributes for SSL operation were defined.
0x000001b4	436	GSK_CSP_OPEN_ERROR	The Microsoft Crypto API is only supported on Microsoft Windows 2000 with Service Pack 2 applied.
0x000001b5	437	GSK_CSP_OPEN_ERROR	System is running in IPv6 mode without setting a PEERID.

Table 37. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x000001f5	501	GSK_INVALID_BUFFER_SIZE	The buffer size is negative or zero.
0x000001f6	502	GSK_WOULD_BLOCK	Used with nonblocking I/O. Refer to the nonblocking section for usage.
0x00000259	601	GSK_ERROR_NOT_SSLV3	SSLv3 is required for reset_cipher(), and the connection uses SSLv2.
0x0000025a	602	GSK_MISC_INVALID_ID	An ID that is not valid was specified for the gsk_secure_soc_misc() function call.
0x000002bd	701	GSK_ATTRIBUTE_INVALID_ID	The function call has an ID that is not valid. This also might be caused by specifying an environment handle when a handle for a SSL connection should be used.
0x000002be	702	GSK_ATTRIBUTE_INVALID_LENGTH	The attribute has a negative length, which is not valid.
0x000002bf	703	GSK_ATTRIBUTE_INVALID_ENUMERATION	The enumeration value is not valid for the specified enumeration type.
0x000002c0	704	GSK_ATTRIBUTE_INVALID_SID_CACHE	A parameter list that is not valid for replacing the SID cache routines.
0x000002c1	705	GSK_ATTRIBUTE_INVALID_NUMERIC_VALUE	When setting a numeric attribute, the specified value is not valid for the specific attribute being set.
0x000002c2	706	GSK_CONFLICTING_VALIDATION_SETTING	Conflicting parameters were set for additional certificate validation.
0x000002c3	707	GSK_AES_UNSUPPORTED	The AES cryptographic algorithm is not supported.
0x000002c4	708	GSK_PEERID_LENGTH_ERROR	The PEERID does not have the correct length.
0x000005dd	1501	GSK_SC_OK	
0x000005de	1502	GSK_SC_CANCEL	
0x00000641	1601	GSK_TRACE_STARTED	The trace started successfully.
0x00000642	1602	GSK_TRACE_STOPPED	The trace stopped successfully.
0x00000643	1603	GSK_TRACE_NOT_STARTED	No trace file was previously started so it cannot be stopped.
0x00000644	1604	GSK_TRACE_ALREADY_STARTED	Trace file already started so it cannot be started again.
0x00000645	1605	GSK_TRACE_OPEN_FAILED	Trace file cannot be opened. The first parameter of gsk_start_trace() must be a valid, full-path file name.

Appendix D. Accessibility features for the Tivoli Storage Manager product family

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 the Tivoli Storage Manager family of products:

- 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

The Tivoli Storage Manager Information Center, and its related publications, are accessibility-enabled. For information about the accessibility features of the information center, see the following topic: http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.help.ic.doc/iehs36_accessibility.html.

Keyboard navigation

On Windows, the Tivoli Storage Manager product family follows Microsoft conventions for all keyboard navigation and access. Drag and Drop support is managed by using the Microsoft Windows Accessibility option known as *MouseKeys*. For more information about MouseKeys and other Windows accessibility options, see the Windows online help, citing the keyword "MouseKeys".

On other operating systems, these products follow the operating-system conventions for keyboard navigation and access.

Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

IBM and accessibility

See the IBM Human Ability and Accessibility Center (<http://www.ibm.com/able>) for information about the commitment that IBM has to accessibility.

Notices

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

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

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

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

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

*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 want 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 information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

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

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol ([®] or [™]), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.



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.

LTO and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

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

Other product and service names might be trademarks of IBM or other companies.

Glossary

A glossary is available with terms and definitions for the IBM Tivoli Storage Manager family of products.

You can view the glossary in the Tivoli Storage Manager information center at <http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3>.

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

Index

Special characters

\$\$_TSMDBMGR_\$\$ hidden user ID 79

A

accessibility features 269
administration center
 configuring the IP address 106
 create a user 102
 create a user ID 103
 establishing a connection with the IBM Tivoli Storage Manager server 99
 messages versus IBM Tivoli Storage Manager messages 114
 support utility 111
 task fails
 check the server activity log 112
 internal errors 114
 portlets 113
 wizard 113
 task fails with message 111
 trace 127
 unable to access the server from a web browser 107
administrators
 locked 14
AIX JFS2
 image backup 30
 snapshot-based backup-archive 30
allocating additional memory 52
ANR1221E
 error message 85
ANR2317W
 error message 86
API
 option file 33
application programming interface (API)
 instrumentation 31
 tracing 176
auto-update 117
automatic database backup errors 70

B

backup application
 files automatically excluded 20
 files excluded by EXCLUDE DIR 22
 files excluded by include/exclude statements 19
 files excluded due to incremental copy frequency 52
 include/exclude due to compression, encryption, and subfile backup statements 23
 include/exclude statements coded wrong 24
 platform-specific include/exclude statements 23

BACKUP DB
 ANR2971E with SQL code 76
 common errors 77
 incorrect environment variables 75
backup-archive client
 automatic deployment 119
 help 1
 restarting 119
 SHOW commands 45
backup-archive client deployment
 unsupported platforms 120
BIRT 204, 209
books
 See publications

C

cache
 bypass during write operations 232
certificate authority 96
client
 authentication failure 10
 can problem be reproduced 6
 error messages
 examining 5
 generating errors
 connected to the server 95
 identifying when and where problems occur 5
 image backup 27
 resolving problems 5
 scheduler 16
 server activity log
 examining 5
 trace classes 162, 166
client option sets
 resolving problems 9
 using 9
client schedule log 17
Cognos 202, 205, 206, 209, 211
collector tool 103
communication errors
 resolving 95
complex password
 LDAP directory server 12
complex passwords
 audit the LDAP directory server 14
compressed data during
 backup-archive 175
configuration 185
copy frequency 52
CTGTRV009E
 error message 90
CTGTRV011E
 error message 90
customer support
 contact xiii

D

Daemon traceflags
 client and journal 161
data
 sent to the IBM Tivoli Storage Manager storage agent or server 34
 unreadable 223
data storage hints and tips
 backup or copy problem with specific node 225
 change the server policies 224
 change the storage hierarchy 224
 HELP 223
 reading or writing to a device 224
 recreate the problem 223
 server activity log 223
 specific volume 225
database error messages 74
database ID file missing or incorrect 73
database manager
 start problems 69
database reorganization 79
database restore errors 73
DB2 log files 71
DB2 memory consumption
 limiting memory 71
DB2 password
 expired 58
DB2 version information 71
DB2CKUPGRADE command
 upgrading the server 59
db2dump directory
 shutdown resolution 67
debug output
 obtaining 40
deployment 117, 119
deployment engine 49, 58
device driver
 32-bit Linux kernel modules 227
 64-bit Linux kernel modules 227
 Adaptec SCSI requirements 230
 error messages in the system error log 227
 HBA changes 226
 HBA drivers on the Linux 2.6.x kernels 228
 Linux server running on x86_64 architecture 227
 loose cable connection 227
 multiple LUN support on Linux kernels 228
 operating system changes 226
 performing ddtrace from version 5.3.2 on Linux 229
 Qlogic fibre-channel HBA BIOS requirements 230
 SCSI adapter changes 226
 updating device information 230
device driver trace
 from a command shell - AIX, Oracle Solaris, Windows 159

- device driver trace (*continued*)
 - from the server console/admin client 158
- diagnostic tips
 - client 5
 - storage agent 121
- diagram 194
- disability 269
- documentation
 - See also* publications
 - to resolve client problems 6

E

- education
 - see Tivoli technical training xi
- enabling 212
- encrypted data during
 - backup-archive 175
- encrypted file system 27
- error messages
 - ANR1330E 87
 - ANR1331E 87
 - ANR2968E 76
 - CTGTRV009E 90
 - CTGTRV011E 90
 - LDAP authenticated passwords 15
- establishing a connection with the IBM Tivoli Storage Manager server 99
- external user repository server stoppage 63

F

- fail 181, 185
- FILE directory mapping 233
- fixes, obtaining xii, xiii
- flowchart 194
- full file system 215

G

- GSKit
 - installation problems 60
 - return codes 259
- gt script 253

H

- health monitor 108
 - how it works 109
 - re-synchronizing the ADMIN_CENTER administrator ID password 110
 - warning or critical database status 110
 - warning or critical storage status 109
- help
 - server or storage agent 2
- help facilities 1
- help system
 - CLI for server or storage agent 3
 - dsmcutil 2
 - GUI and Web GUI clients 3
 - reporting a problem 3

- help system (*continued*)
 - server or storage agent
 - commands 2
 - messages 3
 - Windows 2
- hints and tips
 - device driver 225
 - disk subsystems 231
 - hard disk drives 231
 - NDMP filer-to-Tivoli Storage Manager server operations 250
 - SAN 235
 - SAN configuration 236
 - SAN device mapping 242
 - tape drives and libraries
 - adapter firmware changes 234
 - cabling between the computer and device changes 234
 - device driver changes 234
 - device firmware changes 234
 - error messages in system error log 235
 - loose cable connections 235
 - operating system changes 234
 - other hardware changed or fixed 234
 - replaced adapter 234
 - historical reporting
 - Warehouse Proxy workspace 193
 - history configuration 183

I

- IBM Global Security Kit
 - key management return codes 259
 - return codes 259
- IBM Publications Center vii, x
- IBM Support Assistant xii
- image backup
 - client 27
 - error 27, 29
- INCLEXCL option 19
- InstallAnywhere
 - exit codes 53
- installation 181, 183, 185
- Installation 181
- installation problems 53
- Internet, searching for problem resolution xi, xii
- IP address
 - configuring 106

J

- Japanese language
 - reporting problems 211
- Java archive
 - collector tool 103
- java.net.ConnectException
 - CTGRI0011E error 212
- journal
 - restarting 37
- journal-based backup (JBB)
 - database viewing utility 38
 - determining 37
 - running in foreground 38

K

- key database file
 - out-of-synch 98
 - password recovery 98
- keyboard 269
- knowledge bases, searching xi
- Korean language
 - reporting problems 211

L

- LABEL LIBVOLUME 53
- LAN-free setup
 - storage agent 122
- LDAP directory server
 - password 12
- LDAP-authenticated password
 - problem resolution 11
- Linux image backup error 27
- Linux installation
 - SELinux AVC denial 53
- Linux Snapshot image backup error
 - error message ANS1258E 29
- locked nodes and administrators 14
- log files
 - DB2 upgrade 73
 - installation 53
- logon 190
- logs.zip 181
- LVSA
 - examining Windows system event log 39
 - forcing a memory dump 41
 - full memory dump 40
 - problem determination 39

M

- maintenance
 - client deployment 117
- Management Console
 - server service start/stop 66
- manuals
 - See* publications
- message definitions
 - IBM Tivoli Storage Manager 115
- Microsoft diagnostic information
 - VSS 43
- Microsoft tuning
 - VSS 43
- monitoring agent 191
- monitoring agents
 - activating tracing 177, 179
- moving data to other volumes 232

N

- nodes
 - locked 14
- non-root user
 - running applications using the API 35
- ntbackup.exe 45

O

- open file support
 - best practices 41
 - problem determination 39

P

- Passport Advantage xiii
- password authentication
 - client setup 12
- performing ddtrace on HP-UX 229
- portlets
 - errors caused by starting or stopping 113
- problem determination
 - describing problem for IBM Software Support xiv
 - determining business impact for IBM Software Support xiv
 - submitting a problem to IBM Software xiv
- process ended 83
- process started 83
- process symptoms
 - files not expired 91
 - migration does not run 91
 - migration only uses one process 91
- pruning 215
- publications
 - download vii
 - order vii
 - search vii
 - Tivoli Storage FlashCopy Manager x
 - Tivoli Storage Manager viii

Q

- queries 188

R

- reconfiguring services 216
- RELABEL 53
- remote registry service 212
- reorganization
 - database 79
- reports 194, 199, 201, 202, 204, 205, 206, 209, 211, 219
- Resetting passwords 216
- resolving problems 181
- RESTORE DB
 - ANR2971E with SQL code 76
 - common errors 77
 - incorrect environment variables 75

S

- SAN
 - configuration 247
 - configuration between devices 238
 - configuration problems 250
 - devices supported 236
 - fibre channel switch
 - configuration 237
 - fibre-channel link error report 239

- SAN (*continued*)
 - gateway port settings 238
 - host bus adapter configuration 237
 - host bus adapters 236
 - vendor support 250
- SAN device mapping
 - disabling 243
 - errors 243
 - missing from the display of QUERY SAN 247
- SAN devices
 - storage agent 239
- scheduled event
 - status 17
- scheduler
 - client service restart 18
- scheduling
 - automatic deployment 119
- script 183
- SCSI devices 251
- Secure Sockets Layer (SSL)
 - determining errors 96
 - general return codes 259
- sequential media volume
 - tape 251
- server
 - database 68
 - diagnostic tips
 - change server options or the settings create errors 51
 - checking the server activity log 50
 - code page conversion failure 160
 - failing a scheduled client operation 51
 - recreating the problem 49
 - resolving errors from reading or writing to a device 50
 - resolving failed connections by client or administrators 95
 - resolving server space issues 51
 - process 80
 - process messages 80
 - stoppage or loop errors 62
 - storage pool
 - ANR0522W error message 92
 - collocation 93
 - COPY ACTIVATEDATA command 94
 - high volume usage 93
 - resolving problems 92
 - simultaneous write 94
 - unable to store data 94
- server activity log
 - checking 112
 - checking for errors 17
- server command-definition file 116
- server or storage agent
 - trace class 125
 - trace classes 131
- server stoppage
 - activity log 66
 - library files 65
 - resolving general problems 61
 - server error file (dsmserv.err) 64
 - system image 64
 - system logs 66

- SHOW commands
 - server or storage agent 145
- Snapshot Difference
 - resolving problems 24
- snapshot directory 26
- software distribution 117
- software support
 - describing problem for IBM Software Support xiv
 - determining business impact for IBM Software Support xiv
 - submitting a problem xiv
- Software Support
 - contact xiii
- SSL (Secure Sockets Layer)
 - determining errors 96
 - general return codes 259
- stack trace
 - server or storage agent 130
- starting-up problems
 - dsmdmc 7
 - dsmc 7
 - dsmj 7
- status
 - scheduled event 17
- storage agent
 - diagnostic tips
 - check the server activity log 121
 - error caused by reading or writing to a device 121
 - problems caused by changing server options 122
 - problems from changing storage agent options 122
 - LAN-free setup
 - data sent directly to server 122
 - storage pool configured for simultaneous write 123
 - testing LAN-free configuration 124
 - SAN devices 239
 - support contract xiii
 - support for API
 - before calling IBM
 - files to gather 32
 - information to gather 31
 - support information x
 - support subscription xiii
- SUSE 185

T

- table reorganization 79
- test flags
 - VSS 42
- Tivoli Common Reporting 194, 199, 201, 202, 204, 205, 206, 209, 211
- Tivoli Enterprise Portal 185, 186, 188, 190, 191, 192
- Tivoli Integrated Portal
 - error messages 115
 - excessive memory consumption 106
 - server stoppage
 - log analyzer tool 105
 - user authority problems 101
- Tivoli Monitoring for Tivoli Storage Manager 181

- Tivoli technical training xi
- trace
 - administration center 127
 - client
 - backup-archive client 166
 - device driver 158
 - enable client trace on command line 167
 - enable client trace while client is running 168
 - known problems and limitations 172
 - options 173
 - server or storage agent 129
- trace classes
 - administration center 125
 - client 162
 - server or storage agent 131
- trace data
 - is it compressed during backup-archive 175
 - is it encrypted during backup-archive 175
- tracing
 - agents 177, 179
 - application programming interface (API) 176
 - client 160
 - User ID/Password plug-in 70
- Traditional Chinese language reporting problems 211
- training, Tivoli technical xi
- transient errors
 - VSS 42
- troubleshooting 181, 186, 194
- tsmdiag utility 255
- web browser issues
 - Tivoli Monitoring for Tivoli Storage Manager 213
- Windows
 - VSS 41
- Windows fixes
 - VSS 43
- Windows Services
 - server service start/stop 66
- wizards
 - errors caused by starting or stopping 113
- workspaces 186, 188, 192

U

- uninstall stoppage 58
- unsupported platforms
 - backup-archive client deployment 120
- upgrade
 - server manually 58
- upgrade problems 53

V

- Volume Shadow Copy Services
 - Windows 41
- vsreq.exe sample program 44
- VSS
 - Microsoft diagnostic information 43
 - Microsoft tuning 43
 - ntbackup.exe 45
 - test flags 42
 - trace 44
 - transient errors 42
 - vsreq.exe sample program 44
 - Windows 41
 - Windows fixes 43

W

- WAREHOUS 219
- Warehouse Proxy agent 215



Product Number: 5608-E01
5608-E02
5608-E03

Printed in USA

GC23-9789-04

