

Tivoli Tivoli Storage Manager
Version 6.2

Problem Determination Guide



Tivoli Tivoli Storage Manager
Version 6.2

Problem Determination Guide



Note:

Before using this information and the product it supports, read the information in "Notices" on page 233.

| This edition applies to Version 6.2 of IBM Tivoli Storage Manager (product number 5608-E01, 5608-E02, 5608-E03,
| 5608-E07, 5608-E012), and to all subsequent releases and modifications until otherwise indicated in new editions or
| technical newsletters. This edition replaces GC23-9789-01.

© Copyright IBM Corporation 1993, 2010.

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

Contents

Preface vii

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

New for Tivoli Storage Manager

Version 6.2 xv

New for the server in Version 6.2	xv
Client-side data deduplication	xv
Automatic backup-archive client deployment	xv
Simultaneous-write operations during storage pool migration	xvi
In-flight data encryption using SSL	xvi
New for the Tivoli Storage Manager reporting and monitoring feature in version 6.2	xvi
SCSI passthru support for Windows	xvii
Concurrent read-and-write access to Centera volumes	xvii
The Tivoli Integrated Portal GUI	xvii
The Administration Center not installable on HP-UX	xviii
Sun StorageTek T10000B drive encryption	xviii
MOVESIZETHRESH server option	xviii
CHECKTAPEPOS server option to validate data position on tape	xviii

Chapter 1. Resolving Tivoli Storage Manager client problems 1

Examining error messages	1
Examining the server activity log messages	1
Identifying when and where the problem can occur	1
Reproducing the problem	2
Collecting documentation to resolve problems with the client application	2
Determining why the dsmdc or dsmdmc or dsmdj does not start	3
Resolving problems with client option sets	4
Working with client option sets	5
Resolving password authentication problems	6
Resolving client scheduling problems	7
Determining the status of a scheduled event.	7
Checking for errors in the server activity log	7
Starting and stopping the client service	8
Resolving errors when including or excluding client files during backup processing	9
Identifying files that are included or excluded by the server client option set.	9
Excluding files automatically from backup processing	10
Excluding files with the EXCLUDE.DIR statement	12

Determining whether compression, encryption, and subfile backup statements include or exclude	13
Using delimiters to include or exclude files.	13
Resolving errors due to the incorrectly coded include or exclude list	14
Resolving Snapshot Difference problems.	14
Resolving snapshot directory problems for NetApp or N-Series filer volumes	16
Resolving problems with EFS on AIX.	17
Resolving image backup errors	17
Resolving Linux image backup errors.	17
Resolving image backup failures when using Linux snapshot	19
Resolving errors during AIX JFS2 snapshot-based backup-archive and image backup.	20
API problem determination	21
Support solutions for the Tivoli Storage Manager API	21
Determining if data is sent to the Tivoli Storage Manager storage agent rather than the server	24
Running applications that use the API as a non-root user.	25
Journal Based Backup problem determination	26
Determining if a backup will be journal-based.	27
Running the journal daemon in the foreground	28
Using the Journal Database Viewing utility.	28
Using open file support and the logical volume snapshot agent	29
Examining the Windows system event log	29
Obtaining debug print output	29
Configuring the system for a full memory dump	30
Forcing a memory dump for a system stoppage when you suspect a logical volume snapshot agent problem	30
Best practices for open file support	31
Using Windows Volume Shadow Copy Services	31
Defining VSS transient errors	31
Defining Windows VSS test flags	32
Windows 2003 VSS fixes	32
Volume Shadow Copy Services tuning	32
Gathering VSS diagnostic information for Microsoft assistance	33
Troubleshooting errors using a VSS trace	33
Running VSS API calls with the vsreq.exe sample program	33
Comparing Tivoli Storage Manager and Ntbackup.exe interaction with VSS	34
Show commands for the backup-archive client	34

Chapter 2. Resolving server problems 37

Installation log files.	37
Checking the server activity log	41
Recreating the problem	41
Checking system error log files for device errors	41
Reverting server options or settings	41

Restarting the scheduling service	42	Creating a user ID with access to the Administration Center.	76
Resolving server space issues	42	Resolving a user ID access problem with the Administration Center.	76
Allocating additional server memory	42	Resolving a stopped Tivoli Integrated Portal server	77
Changing the copy frequency	43	Running the collector tool to obtain problem-analysis information	77
Tracing to detect a code page conversion failure	43	Diagnosing log-entry problems by using the log analyzer tool (showlog)	79
Resolving server stoppages	43	Resolving excessive memory consumption problems with the Tivoli Integrated Portal server	80
Resolving a stoppage or loop	44	Configuring the IP address to align with the Administration Center.	80
Resolving wait-state problems with NIS servers	45	Resolving server access problems	81
Finding the server error file (dsmserv.err)	45	Resolving Administration Center health monitor problems	82
Finding the system image (core file)	46	Health monitor conditions that can cause an unknown server status	83
Retrieving library files	46	Resolving health monitor conditions that can cause a warning or critical storage status	83
Retrieving system log files	48	Health monitor conditions that can cause a warning or critical database status.	84
Retrieving the activity log	48	Determining when to resynchronize the ADMIN_CENTER administrator ID password	85
LAN-free restore operations: setting the IDLETIMEOUT option.	48	Administration Center Support utility	85
Resolving database errors.	48	Responding to Administration Center task failure messages	86
Locating DB2 log files after an upgrade	49	Responding to Administration Center messages about unexpected results	87
Resolving a missing or incorrect database ID file problem	49	Checking the server activity log to resolve Administration Center problems	87
Resolving problems with the BACKUP DB and the RESTORE DB commands	50	Resolving errors caused by starting or stopping a wizard or portlet	87
Resolving a stopped uninstallation process	54	Resolving problems caused by internal errors	88
Analyzing the process symptoms to resolve problems	54	Determining the source of a message	89
Reviewing process messages to determine the state of server operations	54	Defining Tivoli Storage Manager messages	90
Analyzing the ANR1221E error message.	60	Resolving Tivoli Storage Manager server command-definition file problems	92
Analyzing the ANR2317W error message	60	Resolving backup-archive client deployment problems	93
Analyzing error messages ANR1330E and ANR1331E.	61	Configuring the server for automatic backup-archive client deployments	94
Resolving error messages CTGTRV009E and CTGTRV011E.	64	Restarting the client operating system during a deployment	95
Files are not expired after reducing versions	65		
Process symptoms indicate migration errors	66		
Resolving storage pool issues	66		
“ANR0522W Transaction failed...” message received	67		
Storage pool experiences high volume usage after increasing MAXSCRATCH value	67		
Storage pool has “Collocate?=Yes” but volumes still contain data for many	67		
Resolving active data pool, storage issues	68		
Chapter 3. Resolving communication problems.	69	Chapter 5. Resolving Data Protection problems.	97
Resolving errors created when connecting to the server	69	Troubleshooting IBM Tivoli Storage Manager for Enterprise Resource Planning	97
Resolving failed connections by clients or administrators	69	Troubleshooting IBM Tivoli Storage Manager for Enterprise Resource Planning common problems	97
Determining Secure Sockets Layer errors	70	Troubleshooting Data Protection for SAP for DB2 problems	102
Recovering the key database file password	72	Troubleshooting Data Protection for SAP for Oracle problems	105
Chapter 4. Resolving Administration Center problems	73	Data Protection for Exchange with VSS backup-restore support	113
Re-establishing a connection between the Administration Center and a Tivoli Storage Manager server	73	Determining if the problem is a Data Protection for Exchange issue or a general VSS issue	114
Resolving Tivoli Integrated Portal user authority problems	75	Tracing the Data Protection client when using VSS technology.	116

Gathering Exchange with VSS information before calling IBM	116		Resolving historical data reporting problems in the Warehouse Proxy workspace	179
Gathering Exchange with VSS files before calling IBM	117		Resolving a reporting and monitoring agent installation hang	180
Troubleshooting Data Protection for Exchange VSS and SAN Volume Controller	119			
Chapter 6. Resolving storage agent problems	121		Chapter 10. Help facilities	181
Checking the server activity log for storage agent information	121		Backup-archive client help	181
Resolving an error caused by reading or writing to a device	121		Accessing help for the Windows service configuration utility (dsmcutil)	182
Resolving problems caused by changing storage agent options	122		Server or storage agent help	182
Resolving problems caused by changing server options or settings.	122		Accessing server or storage agent help for commands	182
			Accessing help for messages	183
			Command-line interface help for the client	183
			Reporting a problem with a help topic	183
Chapter 7. Storage agent LAN-free setup	123		Chapter 11. Determining data storage problems	185
Resolving the issue of data being sent directly to the server	123		Using data storage diagnostic tips	185
Resolving a disqualified LAN-free-enabled storage pool	124		Checking the server activity log to resolve data storage issues	185
Ensuring that data is transferred using a LAN-free environment.	124		Checking HELP for messages issued for a data storage problem	185
			Recreating the data storage problem.	185
			Resolving data storage errors related to reading or writing to a device	186
			Changing the storage hierarchy to resolve data storage problems	186
			Changing the server policies to resolve data storage problems	186
Chapter 8. Using trace to resolve problems	125		Resolving a data storage backup or copy problem that occurs only with a specific node	187
Trace classes for the Administration Center	125		Resolving a data storage problem that occurs only for a specific volume	187
Enabling Administration Center trace	127		Hints and tips for storage	187
Enabling a trace for the server or storage agent	129		Device driver hints and tips	187
Enabling a stack trace for specific messages for the server or storage agent	130		Hard disk drives and disk subsystems hints and tips.	193
Trace classes for a server or storage agent	131		Tape drives and libraries hints and tips.	195
Show commands for the server or storage agent	144		Storage area network hints and tips	197
Enabling a trace for the Tivoli Storage Manager device driver	157		NDMP filer-to-Tivoli Storage Manager server operation hints and tips	213
Tracing from the server console	157		Resolving SCSI device problems	213
Tracing data from a command shell for AIX, Sun Solaris, and Windows	158		Resolving sequential media volume (tape) errors through messages ANR0542W or ANR8778W	214
Tracing data for the client	159			
Client and Journal Daemon traceflags	160		Appendix A. Accessibility features for Tivoli Storage Manager	215
Client trace classes	160			
Enabling a backup-archive client trace	165		Appendix B. Using gt script	217
Determining if data is encrypted or compressed during backup-archive through trace	174			
Tracing the Tivoli Storage Manager reporting and monitoring agent for AIX and Linux.	175		Appendix C. Installing and running the tsmdiag utility	219
Tracing the Tivoli Storage Manager reporting and monitoring agent for Windows	176			
Tracing data for the API.	177		Appendix D. IBM Global Security Kit return codes.	223
Chapter 9. Resolving problems with the Tivoli Storage Manager reporting and monitoring feature	179		Notices	233
Resolving Warehouse Proxy workspace reporting problems	179			

Trademarks 235

Index 239

Glossary 237

Preface

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, you should be familiar with the following areas:

- The operating systems on which your Tivoli Storage Manager servers and clients reside
- The communication protocols installed on your client and server computers

Who should read this guide

This guide was written for anyone administering or managing IBM Tivoli Storage Manager. Similarly, information provided by this guide might 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 configured 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. The 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 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

IBM Tivoli Storage Manager publications and other related publications are available online.

You can search all publications in the Tivoli Storage Manager Information Center: <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r2>.

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 Tivoli Storage Manager products at <http://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Tivoli+Storage+Manager>.

You can also order some related publications from the IBM Publications Center Web site. The Web site 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

Publications are available for the server, storage agent, client, and Data Protection.

Table 1. IBM Tivoli Storage Manager troubleshooting and tuning publications

Publication title	Order number
<i>IBM Tivoli Storage Manager Client Messages and Application Programming Interface Return Codes</i>	SC27-2877
<i>IBM Tivoli Storage Manager Server Messages and Error Codes</i>	SC27-2878
<i>IBM Tivoli Storage Manager Performance Tuning Guide</i>	GC23-9788
<i>IBM Tivoli Storage Manager Problem Determination Guide</i>	GC23-9789

Table 2. 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 Sun Solaris Installation Guide</i>	GC23-9784
<i>IBM Tivoli Storage Manager for Sun Solaris Administrator's Guide</i>	SC23-9772
<i>IBM Tivoli Storage Manager for Sun 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 Server Upgrade Guide</i>	SC23-9554
<i>IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack</i>	SC27-2828

Table 3. 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
<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 Sun Solaris Storage Agent User's Guide</i>	SC23-9800

Table 3. Tivoli Storage Manager storage agent publications (continued)

Publication title	Order number
<i>IBM Tivoli Storage Manager for SAN for Windows Storage Agent User's Guide</i>	SC23-9553

Table 4. 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 for Space Management for UNIX and Linux: User's Guide</i>	SC23-9794
<i>IBM Tivoli Storage Manager Using the Application Programming Interface</i>	SC23-9793

Table 5. Tivoli Storage Manager Data Protection publications

Publication title	Order number
<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

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 these Web sites for training information:

Tivoli software training and certification

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

Tivoli Support Technical Exchange

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

Searching knowledge bases

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

Begin by searching the Tivoli Storage Manager Information Center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r2>. From this Web site, you can search the current Tivoli Storage Manager documentation.

Searching the Internet

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

To search multiple Internet resources, go to the support Web site for Tivoli Storage Manager at http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager.

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 Web site, 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 resolve your problem.

An independent user discussion list, ADSM-L, is hosted by Marist College. You can subscribe by sending an e-mail 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 user community, go to the Tivoli Storage Manager wiki at <http://www.ibm.com/developerworks/wikis/display/tivolistoragemanager>.

Using IBM Support Assistant

IBM Support Assistant is a complimentary software product that helps you with problem determination. You can 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.

IBM Support Assistant helps you gather support information when you need to open a problem management record (PMR), which you can then use to track the problem. For more information, see the IBM Support Assistant Web site at <http://www.ibm.com/software/support/isa/>.

The product-specific plug-in modules provide you with the following resources:

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

Find add-ons for specific products here: <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 Web site.

You can determine what fixes are available by checking the IBM Software Support Web site 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 Tivoli Storage Manager product, or one of the other Tivoli Storage Manager components that you want to find a fix for.
 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 **My notifications** in the notifications module.
2. Sign in using your IBM ID and password. If you do not have an ID and password, click **register now** above the IBM ID and password.
3. Click the **Subscribe** tab to select your product family and click **Continue**.
4. Select the type of information that you want to receive, and add your personal preferences. You can specify how you want to be notified, how often, and you can also optionally select a folder for the notifications.
5. Click **Submit**.
6. For notifications for other products, repeat steps 4 and 5.

Tip: You can also pick a product first, from the main support portal site, and then click in the **Notifications** section to create or update your subscription for that product.

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.

Before you contact IBM Software Support, follow these steps:

1. Set up a subscription and support contract.
2. Determine the business impact of your problem.
3. Describe your problem and gather background information.

Then see “Submitting the problem to IBM Software Support” on page xiii for information on contacting IBM Software Support.

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 UNIX® operating systems), enroll in IBM Passport Advantage® in one of the following ways:

- **Online:** Go to the Passport Advantage Web page at <http://www.ibm.com/software/lotus/passportadvantage/>, click **How to enroll**, and follow the instructions.
- **By Phone:** You can call 1-800-IBMSERV (1-800-426-7378) in the United States, or for the phone 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 gather 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 recreated? 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 phone.

Online

|
|
|
|

Go to the IBM Software Support Web site 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 phone

For the phone number to call in your country, go to the contacts page of the IBM Software Support Handbook at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html>.

New for Tivoli Storage Manager Version 6.2

IBM Tivoli Storage Manager version 6.2 includes many new features.

New for the server in Version 6.2

Tivoli Storage Manager server Version 6.2 contains many new features and changes. Any updates that have been made to the information since the previous edition are marked with a vertical bar (|) in the left margin.

Client-side data deduplication

| In client-side data deduplication, the Tivoli Storage Manager backup-archive client and the server work together to identify duplicate data.

| Data deduplication is a method of reducing storage needs by eliminating
| redundant data. In Tivoli Storage Manager V6.1, only the server could identify and
| remove redundant data. In V6.2, you have the option of identifying and removing
| redundant data during backup and archive processing before data is sent to the
| server. This method of data deduplication is called *client-side data deduplication*. It is
| available with V6.2 backup-archive clients and the V6.2 Tivoli Storage Manager
| application programming interface (API).

| Client-side data deduplication provides several advantages to server-side data
| deduplication. Client-side data deduplication reduces the amount of data sent over
| the local area network (LAN). In addition, the processing power that is required to
| identify duplicate data is offloaded from the server to client nodes. The processing
| that is required to remove duplicate data on the server is eliminated. Space savings
| occur immediately.

| If you used server-side data deduplication, V6.2 client nodes can access existing
| deduplicated data and storage pools that are already set up for data deduplication.
| When restoring or retrieving files, the client node queries for and displays files as
| it normally does. If a user selects a file that exists in a deduplicated storage pool,
| the server manages the work of reconstructing the file.

| You enable client-side data deduplication using a combination of settings on the
| client node and the server. The primary storage pool that is specified by the copy
| group of the management class associated with the client data must be a
| sequential-access disk (FILE) storage pool that is enabled for data deduplication.

Automatic backup-archive client deployment

| IBM Tivoli Storage Manager V6.2 can deploy backup-archive client code to
| workstations that already have the backup-archive client installed.

| You can now deploy backup-archive client code to candidate client workstations
| from the Tivoli Storage Manager V6.2 Administration Center. From the
| Administration Center, you can coordinate the client updates to each workstation
| that is at release 5.4 and later to V6.2. You are helped through the process by
| wizards that configure your workstation and schedule the deployments. The
| backup-archive client deployment feature is available for Windows backup-archive
| clients only.

Simultaneous-write operations during storage pool migration

With Tivoli Storage Manager, you can now write data simultaneously to copy storage pools and active-data pools during server data-migration processes.

The simultaneous-write function during migration can reduce the amount of time required to back up storage pools or copy active data. Data that is simultaneously written to copy storage pools or active-data pools during migration is not copied again to the copy storage pools or active-data pools. For example, suppose that you migrate all the data in your primary random-access disk storage pool nightly and then back up your primary storage pools. By using the simultaneous-write function during migration, you can significantly reduce the amount of time required for backup operations.

You can also use the simultaneous-write function during migration if you have many client nodes and the number of mount points that are required to perform the simultaneous-write function during client store operations is unacceptable. If mounting and demounting tapes when writing data simultaneously during client store operations is taking too much time, consider writing data simultaneously during migration.

With Tivoli Storage Manager V6.2, you can specify the simultaneous-write function for a primary storage pool if it is the target for *any* of the eligible operations (client store sessions, server import processes, and server data-migration processes).

In-flight data encryption using SSL

Support for Secure Sockets Layer (SSL) is available on HP-UX, Linux[®], Solaris, AIX[®], and Windows platforms.

With SSL industry-standard communications, you can encrypt all traffic between the backup-archive client, the administrative command-line clients, and the IBM Tivoli Storage Manager server. You can use either self-signed or vendor-acquired SSL certificates.

New for the Tivoli Storage Manager reporting and monitoring feature in version 6.2

The Tivoli Storage Manager reporting and monitoring feature, Version 6.2 has a few new changes.

The Tivoli Storage Manager reporting and monitoring feature, Version 6.2, has been integrated into a new user interface called the Tivoli Integrated Portal. This move affects the reporting and monitoring reports that are run from the Administration Center. The Administration Center moved from the Integrated Solutions Console to the Tivoli Integrated Portal. The Tivoli Integrated Portal provides all the functions that were available in the Integrated Solutions Console, but with a new look-and-feel.

The Administration Center is installed separately and is not included in the reporting and monitoring installation.

There is a new information roadmap for the Tivoli Storage Manager reporting and monitoring feature on the Tivoli Storage Manager Wiki. This roadmap has detailed information on planning, installing, configuring, customizing, and trouble shooting. Reporting and monitoring feature information roadmap

SCSI passthru support for Windows

Windows

SCSI passthru support is available for Windows in Tivoli Storage Manager Version 6.2.

With this support, you can choose to use a Windows Hardware Qualification Lab certified native device driver instead of the Tivoli Storage Manager device driver to control devices. Devices currently controlled by the Tivoli Storage Manager device driver can be switched to a native driver without updating drive or device class definitions.

Concurrent read-and-write access to Centera volumes

AIX

HP-UX

Solaris

Windows

In previous versions of Tivoli Storage Manager, a client session or server process had to wait for a Centera volume if the volume was in use by another session or process. In V6.2, server read-access and write-access to a Centera volume are available concurrently.

Concurrent access improves restore performance. Two or more clients can read the same volume at the same time. One client can also write to the volume while it is being read. In addition, multiple client sessions and server processes (for example, a client restore operation and an export node operation) can read the same volume concurrently.

The following server processes can share read access to Centera volumes:

- Exporting client node definitions or file data to sequential media or directly to another server for immediate import
- Exporting all or part of server control information and client file data (if specified) from the server to sequential media
- Generating a backup set for a backup-archive client node

The following server processes cannot share read access to Centera volumes:

- Checking for inconsistencies between a storage pool volume and database information
- Deleting a storage pool volume and, optionally, the files stored in the volume

A Centera volume can appear as the current volume for more than one session and as the target of concurrent read and write operations. There are no command changes associated with this feature.

The Tivoli Integrated Portal GUI

AIX

Linux

Solaris

Windows

The IBM Tivoli Integrated Portal is a graphical user interface (GUI) that is included with Tivoli Storage Manager V6.2. The Tivoli Integrated Portal provides all the functions that were available in the Integrated Solutions Console.

The Administration Center, Tivoli Storage Manager reporting and monitoring feature, and other applications are integrated into this new graphical user interface. The Administration Center can be moved to the Tivoli Integrated Portal if the

servers being managed are at version 5.5 or later. By deploying the Tivoli Integrated Portal early, you can prepare your system for an upgrade to Tivoli Storage Manager V6.2. Servers at versions earlier than 6.2 that are managed using the V6.2 Administration Center cannot use the version V6.2 features.

The Administration Center not installable on HP-UX

The Administration Center, a Web-based interface for centrally configuring and managing Tivoli Storage Manager servers, cannot be installed on an HP-UX server.

In IBM Tivoli Storage Manager Version 6.2, the Administration Center cannot be installed on an HP-UX server. However, when installed on a supported server platform, the Administration Center can be used to manage HP-UX servers. For Administration Center system requirements, see the following Web site:
<http://www.ibm.com/support/docview.wss?uid=swg21410467>

Sun StorageTek T1000B drive encryption

You can now use tape device encryption with Sun StorageTek T1000B drives. Encryption provides security for data on individual tapes and protects sensitive information that is transported off-site. When enabled, Tivoli Storage Manager handles encrypting and decrypting data on tapes according to specifications set when defining an ECARTRIDGE device class.

MOVESIZETHRESH server option

The MOVESIZETHRESH server option default and maximum values have been increased.

The MOVESIZETHRESH option specifies, in megabytes, a threshold for the amount of data moved as a batch, within the same server transaction. When this threshold is reached, no more files are added to the current batch, and a new transaction is started after the current batch is moved. The default value for MOVESIZETHRESH has been increased from 2048 to 4096; and the maximum value has also been increased from 2048 to 32768.

CHECKTAPEPOS server option to validate data position on tape

With the new CHECKTAPEPOS server option, you can determine the validity and consistency of the position of data blocks on tape.

The CHECKTAPEPOS option applies to only operations using tape drives. It does not apply to non-tape, sequential-access device classes such as FILE or OPTICAL. If the server information about position does not match the position detected by the drive, an error message is displayed, the transaction is rolled back, and the data is not committed to the database.

Chapter 1. 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

One way to resolve problems is to examine the error messages that are generated during operation.

The QUIET processing option for the 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 dsmserror.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

Check for server activity log using QUERY ACTLOG for messages issued for this 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

If the problem can be reproduced, try to minimize the circumstances under which it can occur.

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 TCPSERVERADDRESS, TCPPORT, 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 IBM Tivoli Storage Manager client creates valuable information in a number of different sources. If any of the information in the documentation relates to your problem, provide it to Support.

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 be used to 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 client might use a combination of files for its configuration. These files are `dsm.opt`, `dsm.sys` for AIX, HP-UX, Linux, or Sun 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 client stops running, many platforms will generate an application dump. The operating system provides the application dump.
- Memory dump. If the Tivoli Storage Manager 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 will collect 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 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. These can be queried from the server using the QUERY SCHEDULE command.
- The policy information configured for this node on the IBM Tivoli Storage Manager server. This information can be queried from the server using the QUERY DOMAIN, QUERY POLICYSET, QUERY MANAGEMENTCLASS, and QUERY COPYGROUP commands.

Determining why the dsmc or dsmadmc or dsmej does not start

IBM Tivoli Storage Manager uses `dsmc/dsmadmc/dsmej` in its startup procedures. The `dsmc/dsmadmc/dsmej/` occasionally does not start, keeping you from using Tivoli Storage Manager.

Processing stops and the following message is displayed if the `dsmc/dsmadmc/dsmej/` will 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: The `dsmerror.log` file is used only as an example file in the message.

Client applications do not run without a log to which you can write and the system will deny “write” access to the log file named in the message. If the log does not exist, it will be 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 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 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 logs 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 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 ('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 attempt to run Tivoli Storage Manager applications using an error log 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/mgs/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/mgs/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 client's option file. The 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 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 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.

Working with client option sets

There are several situations where you can use client option sets.

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

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

In this scenario you have a critical environment and restoring is a high priority.

You want to use `collocatebyfilespec` so that all `filespec` data is stored on as few tapes as possible, which will enhance restore processing because fewer tape mounts are needed.

You do not want the client to be able to override this option. 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
```

In this scenario you have workstations on a slow network with limited space on the Tivoli Storage Manager server for data. Use the compression option to limit the amount of data sent and stored.

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
```

In this scenario you use a database which cannot be stopped. Because the files are open, the server cannot back them up. You want to exclude all files and subdirectories from Tivoli Storage Manager backups and add the files and subdirectories to the existing "space_rest" client option set. Use the `EXCLUDE DIR` command and specify the directory path to be excluded. Use the following Tivoli Storage Manager server command:

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

In this scenario you have a fast network and want to make the best possible use of Tivoli Storage Manager client resources to complete backups. Set `RESOURCEUTILIZATION` to the maximum amount. Use 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 generally occurs when the password expires. It 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. You can change the password using either of the following two methods:

- The administrator can reset the node's password by issuing the UPDATE NODE command.
- You can issue a DSMC QUERY SESSION command and, when prompted, enter the new password.

If you receive the ANS1025E error message during a noninteractive session, such as central scheduling, ensure that the client option is PASSWORDACCESS GENERATE. This option causes the client to store the password locally. The password is encrypted and stored either in the registry for Windows clients or in a file named TSM.PWD for Macintosh clients. You should 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: If this field is blank, the default password expiration period of 90 days is in effect.

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 will disable 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 instead.

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 “Changing the processing options used by the scheduler service” section of the *Windows Backup-Archive Clients Installation and User's Guide*.

Resolving client scheduling problems

The IBM Tivoli Storage Manager administrator can schedule Tivoli Storage Manager to perform tasks automatically.

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

- Changes to the Tivoli Storage Manager client options are not recognized by the client scheduler until the scheduler is restarted. Stop and restart the client scheduler.
- Use the SHOW PENDING diagnostic tool to display schedules, nodes, and when they should next run.
- View the dsm.sys stanza for the node and the MANAGEDSERVICES, PRESCHEDCMD, and POSTSCHEDCMD option values from the client options file for information when a node misses a scheduled event.

Determining the status of a scheduled event

The server maintains a record of all scheduled events, which is useful when managing 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 for more information.

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

The 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. The schedule log file is typically stored in the same directory as the `dsmerror.log` in a file named `dsmsched.log`. 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 `SCHEMLOGNAME` 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 has been 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 what you might search:

- 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

To resolve client scheduling problems, you can opt for starting and stopping the client service.

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 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 Sun Solaris, you can write a shell script to search for and stop running 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.

The following shell script is an example of 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. If you implicitly or explicitly indicate that a file should be included or excluded during backup processing and it was not processed correctly, there are several possible causes.

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

The 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 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 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 have been 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 incl excl
*** 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 desired 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

Two reasons for the backup application to not back up some files might be that there are files that were identified by the operating system as not necessary for backup, or that there are files that IBM Tivoli Storage Manager uses for internal processing.

If these files must be included in the backup processing, Tivoli Storage Manager can include them by putting *INCLUDE* statements in the client options set on the server.

Important: Because these 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's DSMC QUERY INCLEXCL command to identify these files. The output from this 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 6 on page 11 for the files that are automatically excluded.

Table 6. Files automatically excluded during backup processing

Platform	Files Excluded
Windows	<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 • RSM database files (these files are processed in the system object or system state backup) • IIS metabytes (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 file • 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 will not be processed during backup of the system drive but 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>
AIX, HP-UX, Linux, Sun Solaris	Client trace file
Macintosh	<ul style="list-style-type: none"> • Volatile, temporary, and device files used by the operating system • Client trace file

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 the 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) will not be filtered.

If you want to perform a selective backup of a single file from the command-line client, it will not be 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 which is used in subsequent command-line actions:

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

The following selective backup will always result 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 will inactivate the file “/home/spike/my.file.”

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

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

```
exclude.dir /usr/ (AIX, HP-UX, Linux, Sun 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, Sun 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 desired results.

Consider the following example:

AIX

HP-UX

Linux

Solaris

```
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

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

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

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, Sun 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 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 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
- **Windows** 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 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. It can only back up names containing ASCII characters. Two Snapshot Difference behaviors have been 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 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 log 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 `snapdiff` 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 using the `DSMC DELETE BACKUP` command. The Snapshot Difference API does not detect that a file has been 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 snapdiff 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
- snapdiffdb

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 filer volumes

When a Network File System (NFS) mounted or a Common Internet File System (CIFS) mapped volume is backed up, all snapshots within the .snapshot directory are also backed up.

Remember: The NFS mounted or CIFS mapped volumes can be either NetApp or N-Series.

To avoid backing up unwanted snapshots, use NDMP (Network Data Management Protocol) backup. Or you can back up your data with the SNAPSHOTROOT option or run an incremental backup using the SNAPDIFF option. The SNAPDIFF option uses the snapshot difference feature. Alternatively, exclude the .snapshot directory from any backup.

Resolving problems with EFS on AIX

AIX

The encrypted file system (EFS) keystore opens automatically during login when the keystore password matches the user login password. This default behavior occurs during the initial user keystore creation for AIX 6.1 and later releases.

When the login password for AIX is different from the EFS keystore password, the user must open the keystore manually before starting the IBM Tivoli Storage Manager client. Issue the following command to open the keystore:

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

For example, to start the command-line client, issue the following command:

```
efskeymgr -o ./dsmc
```

To start the Java GUI client, issue the following command:

```
efskeymgr -o ./dsmj
```

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

```
efskeymgr -n
```

Resolving image backup errors

Errors that might occur with image backup processes are a Linux image backup failure, a Linux Snapshot image backup failure, or errors occurring 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 will display similar to the following output:

```

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

```

or

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

```
dumas:/develop/rem/peterman # rpm -q -a |grep device-mapper
```

The output will display 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 a lower version, please upgrade the device mapper rpm file at the following Web site:
<http://www.novell.com>.

2. Verify the installation after the upgrade.

```

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

```

You can also check the /lib directory to see the correct versions installed. A system with the right levels will 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 this command fails with error “Snapshot: Required device-mapper target(s) not detected in your kernel,” the **:dm_snapshot** kernel module is not loaded. This command could also fail for other reasons as well, 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:
suzie:/home2/rat # lsmod |grep dm_
dm_mod 112104 6
2. Load the module. See the following example:
suzie:/home2/rat # modprobe dm_snapshot
3. Verify that the previous step is successful. See the following example:
suzie:/home2/rat # lsmod |grep dm_
dm_snapshot 44024 0
dm_mod 112104 6 dm_snapshot
suzie:/home2/rat #
4. Create a snapshot from the shell prompt. See the following example:
suzie:/etc # /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:
suzie:/etc # lvremove /dev/system/tsmSnap
Do you really want to remove active logical volume "tsmSnap"? [y/n]: y
Logical volume "tsmSnap" successfully removed
suzie:/etc #

If you followed all the steps, you now might 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 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 in which 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. `UNMOUNT -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 `- UNMOUNT -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 show 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, Tivoli Storage Manager produces an error message 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 snapshot monitoring, will not perform the snapshot deletion feature described above, and will not perform used block image backup unless AIX is at the AIX 5.3.0.70 or later and AIX 6.1 or later levels. 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.

API problem determination

There are several methods you can use to determine the source of problems with the application programming interface (API).

Support solutions for the Tivoli Storage Manager API

A number of resources are available to learn about or to diagnose the 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

By collecting all the necessary information about your environment, you can significantly help to determine the problem.

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 Tivoli Storage Manager application program interface (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)?
- List all applications running on the system.
- List the steps required to recreate the problem (if the problem can be recreated).
- If you cannot recreate the problem, list the steps that caused the problem.

Gathering API files before calling IBM support

A number of log files and other data might be created by the Tivoli Storage Manager application programming interface (API).

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

- 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

This 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 this same directory.

DSMI_LOG

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` will 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 will be used as the error log name instead of the default value `dsierror.log`.

- Tivoli Storage Manager API error log file. The default API error log 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. The Tivoli Storage Manager administrator can view this log 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.

Verifying that the API uses the correct option file:

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

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 'ERRORNEOUS_OPTION 12345' line into the srvr1.cmpron server stanza of the dsm.sys file. See the following example:

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

SERVERNAME srvr1.cmpron
COMPRESSION YES
ERRORNEOUS_OPTION 12345
TCPSEVERADDRESS computer.company.com

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

2. Verify that the API detects this 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.
```

The wrong options file was updated if no error is reported.

3. Check the environment variable values that were previously mentioned or repeat steps 1 and step 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 need to know if your data is being sent to the IBM Tivoli Storage Manager storage agent, rather than a server because you will not be able to recover the data if it is sent to the storage agent.

Perform the following steps to verify that the data is being sent to the Tivoli Storage Manager storage agent, rather than 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>'
```

The statement above is followed by the 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

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

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` will fail 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's permissions need to 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 will result in a **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 will 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 `dapismp` 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 files systems with small or moderate amounts of change activity between backup cycles.

Related information

 [JBB FAQ](#)

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 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 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 might want to restart a valid journal for better performance.

The reasons for restarting a valid journal are displayed in the following list:

- 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:
 - The server filespace no longer exists
 - The server filespace 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.

The journal daemon might be started from a Windows command prompt as follows: `tsmjbbd.exe i`

Using the Journal Database Viewing utility

The Journal Database Viewing utility can provide you with valuable information.

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: You must have Version 5.4 of the utility (`dbviewb.exe`) in order to view the 5.4 journal databases, which can be found at the IBM software ftp site: <ftp://ftp.software.ibm.com/storage/tivoli-storage-management>.

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:\tsmjournal\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:\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
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

```

Using open file support and the logical volume snapshot agent

There are several methods in determining problems with the open file support (OFS) and logical volume snapshot agent (LVSA).

Examining the Windows system event log

Windows

In the 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 debug print output

Many problems require that you obtain trace data from the `tsmlvsa.sys` driver to supplement what is obtainable from a client service trace.

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

Perform the following steps to install the DebugView tool:

1. Download the latest version of the DebugView tool. Make sure you use the version that is labeled "you plan on using DebugView on WinNT/2K/XP."
2. Install the DebugView tool. The installation is a simple extraction of the files from the `dbgvnt.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 many GB of free space.

4. Run the `dbgview.exe` executable file prior to running the Tivoli Storage Manager client.
5. Configure `dbgview.exe` to log the output to a file through the **File** → **Log to File** option.
6. Perform the failing operation.

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 rename the previous dump file (if any) to a new name.

When a bug check occurs, note the contents of the bug check screen (commonly referred to as the Blue Screen Of Death or BSOD). 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

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 244138 “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. If the system is stopped and you cannot afford a reboot and recreate:

- a. Download and install BANG! from the Windows Driver Developers Web site. Follow instructions provided in the BANG! package/website.
- b. Run BANG! and press the “Crash Now” button. The system normally 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 Web site.

Best practices for open file support

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

Related information

 <http://www-1.ibm.com/support/docview.wss?uid=swg21248971>

Using Windows Volume Shadow Copy Services

The 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

The 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

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

Microsoft issued several fixes for Volume Shadow Copy Services (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

Microsoft issued several tuning fixes for Volume Shadow Copy Services (VSS).

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 aborted because the diff area file could not grow in time.”

Reduce the I/O load on this system to avoid this problem. If these events still occur, you can 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 below 64 MB.

Gathering VSS diagnostic information for Microsoft assistance

Microsoft Support can assist you with diagnostic information for Volume Shadow Copy Services (VSS) failures.

If the VSS failure is outside of the scope of 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 logs focusing on the error events created by the VolSnap and VSS sources at the time of failure. You might want to extract the germane events from the log to isolate the problem and have a more productive interaction with MS support.

Troubleshooting errors using a VSS trace

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 using the contents shown and change the TraceFile entry to point to a volume that is not going to have a shadow copy created. Note the double-backslash delimiter usage; you must enter “\\” as the delimiter for each backslash in the path 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:ffffffff
"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

The Volume Shadow Copy Services (VSS) Software Developer's Kit (SDK) contains the vsreq (VSS requester) sample program which performs a sequence of VSS API calls similar to 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 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

Using the Ntbackup.exe executable file does not fully utilize Volume Shadow Copy Services (VSS) and cannot always be considered as a benchmark for 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 evidence 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 7 are not all of the available SHOW commands.

Table 7. SHOW commands for the backup-archive client

SHOW Command	Description	Information
CACHE	Displays information about the subfile cache.	For Microsoft Windows clients that have 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 7. 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 DEFINE CLOPTSET 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 that 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 by a client to 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.	It 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.	It 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.	It is helpful in determining which trace classes and aggregate trace classes could be used for this client.
VERSION	Displays the version and build date for this client.	It is helpful in determining which client is running and when it was built.

Chapter 2. 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

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 For Linux users, 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 8 on page 38 describes the files that are created when you install or uninstall IBM Tivoli Storage Manager and recommends which files to check when looking for information that might help you troubleshoot problems:

Table 8. Contents of the log.zip file

File Name	Description	Location
The log.txt file contains information about the following Environment Checks: <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 Sun Solaris, the file is located in /var/tivoli/tsm.</p> <p>The InstallAnywhere exit codes are in the log.txt file and can also be summoned by command. You can only retrieve the exit codes after the installer is finished installing. The exit codes are for both the Tivoli Storage Manager installer as well as the Tivoli Storage Manager reporting and monitoring installer.</p> <p>AIX HP-UX Linux Solaris To use the command line from an AIX, HP-UX, Linux, or Sun Solaris server, issue the following command: echo \$?</p> <p>Windows To use the command line from a Windows server, issue the following command: echo %ERRORLEVEL%</p> <p>See Table 9 on page 39 for all of the InstallAnywhere exit codes.</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 Linux Solaris</p> db2setup.log	<p>Contains information about the DB2 installation.</p> <p>Review this log file if the DB2 installation fails</p>	coi\plan\tmp
<p>Windows</p> db2setup.log	<p>Contains information about the DB2 installation</p>	coi\plan\logs
db2_uninst.log	<p>Contains information about the DB2 uninstallation</p>	coi\plan\logs

Table 8. 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; align-items: center;"> <div style="margin-right: 10px;"> <div style="background-color: #800000; color: white; padding: 2px 5px; margin-bottom: 2px;">AIX</div> <div style="background-color: #800000; color: white; padding: 2px 5px; margin-bottom: 2px;">Linux</div> <div style="background-color: #800000; color: white; padding: 2px 5px; margin-bottom: 2px;">Solaris</div> <div style="background-color: #800000; color: white; padding: 2px 5px;">Windows</div> </div> <div> Administration Center installation log files </div> </div>	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; align-items: center;"> <div style="margin-right: 10px;"> <div style="background-color: #800000; color: white; padding: 2px 5px; margin-bottom: 2px;">AIX</div> <div style="background-color: #800000; color: white; padding: 2px 5px; margin-bottom: 2px;">Linux</div> <div style="background-color: #800000; color: white; padding: 2px 5px;">Solaris</div> </div> <div> Files with the following file extensions are available for troubleshooting: .log .out .err </div> </div>	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 9. InstallAnywhere exit codes

Code	Description
0	Success: The installation completed successfully without any warnings or errors.

Table 9. InstallAnywhere exit codes (continued)

Code	Description
1	The installation completed successfully, but one or more of the actions from the installation sequence caused a warning or a non-fatal error.
-1	One or more of the actions from the installation sequence caused a fatal error.
1000	The installation was cancelled by the user.
1001	The installation includes an invalid command line option.
2000	Unhandled error.
2001	The installation failed the authorization check, may 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 Install 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 launched 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 could not be found at the directory specified by the java.home system property.
4001	An incorrect path to the installer jar caused the relauncher to launch incorrectly.

Checking the server activity log

Check the server activity log 30 minutes before and 30 minutes after the time of the error for other messages.

To review the messages in the server activity log, issue the `QUERY ACTLOG` command (see *IBM Tivoli Storage Manager Administrator's Reference* for more details).

Often, other messages can offer additional information about the cause of the problem and how to resolve it.

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 twenty clients. In some cases, by changing the timing or order of implementation of events might prevent the problem from reoccurring. For example, one way to change the timing would be to run expiration at a time when the nightly scheduled backups for twenty clients was not running.

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.

Some examples of system error logs are shown in the following list:

- `errpt` for AIX
- Event Log for Windows

Reverting server options or settings

If there were more 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 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 server messages being issued.

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. Typically, if operations fail because a storage pool is out of space, this situation is indicated by other server messages being issued. To add more space to a DISK storage pool, allocate one or more new storage pools and define them to the server using DEFINE VOLUME. 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.

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

- **AIX** For AIX, ensure that there is sufficient paging space. You can also use SMIT to determine if the number of applications is causing a memory shortage.
- **HP-UX** For 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** For Sun Solaris, ensure that there is sufficient paging space. Consult your Sun Solaris system documentation for details.
- **Windows** For 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

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 following administrative client command, `QUERY COPYGROUP`, to determine the setting of the copy frequency parameter:

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

Tracing to detect a code page conversion failure

The 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 UNICODER 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.

Resolving server stoppages

If you know the source of your server stoppage, you can resolve several problems.

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 C, "Installing and running the tsmdiag utility," on page 219

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. 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)
- 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** For AIX, HP-UX, Linux, or Sun Solaris operating systems, 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** For the Windows operating system, 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 NIS servers

AIX **Solaris**

When using a network information service (NIS) server that has many user groups defined, IBM Tivoli Storage Manager might seem stopped for a prolonged period.

For example, Tivoli Storage Manager sometimes takes a long time to connect to the IBM DB2 server. The Tivoli Storage Manager server does not respond to administrative requests and seems to be stopped.

Complete the following steps to resolve a wait-state problem that occurs with AIX or Sun Solaris servers when using an NIS 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 Sun Solaris servers:

- a. db2set DB2_ALTERNATE_GROUP_LOOKUP=FASTGROUPS
 - 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** For Linux and AIX, the trap handler is disabled to prevent the function traceback from printing on the console and in dsmserv.err. 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 For the Windows platform, if the server is running as a service, the file is named `dsmsvc.err`.

HP-UX For 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 will be in file `dsm_gdb.info` (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, this file should be renamed 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** For AIX, HP-UX, Linux, or Sun Solaris systems, 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** For Windows systems, 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 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

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

For all platforms, 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 our system, we need all of the following files which are located in the directory where the server is installed:

- `dsmserv`
- `dsmlicense`
- `ndmpspi`
- `dsmcored`
- `dsmaio`
- `centera`

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

- **AIX** For AIX systems, 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 DSMSEV 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** For HP-UX, issue the CHATR DSMSEV 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/libc1.2
 - /usr/lib/libc.2
 - /usr/lib/libdl.2
- **Linux** For Linux, issue the LDD DSMSEV 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** For 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 problems with 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.

Collect the activity log at least two hours prior to the stoppage and 30 minutes after the stoppage using the `QUERY ACTLOG` command. See the *Tivoli Storage Manager Administrator's Reference* for more details.

LAN-free restore operations: setting the IDLETIMEOUT option

When performing a LAN-free restore using an xLinux server, you must set the `IDLETIMEOUT` option on the server and storage agent. The option value must be greater than the expected restore time.

If the main session on the server times out before the restore operation completes, a restartable restore session might be left over. This leftover session can prevent backups from running on the restore directory specification. As a workaround, use the `cancel restore` command to cancel the restartable restore.

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.

IBM Tivoli Storage Manager version 6.2 is installed with the IBM DB2 database application. You can find the most information for DB2 on the DB2 Information Center and within the DB2 documentation.

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

Related information

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

Locating DB2 log files after an upgrade

When upgrading the server from IBM Tivoli Storage Manager V6.1 to V6.2, 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 and DB2 V9.7. 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.

The following link contains more details about possible error messages found in the listed log files: <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.qb.upgrade.doc/doc/t0007187.html>

If the error indicates that the database is in an inconsistent state, see the following link: <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.messages.sql.doc/doc/msql01015n.html>

Upgrade again after you fix 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 IBM 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 dsmserve.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 dsmserve.dbid file. Use the -S parameter when you start the server after restoring.
- When the dsmserve.dbid file gets damaged or lost.

After the initial use of the -S parameter in a restore scenario, the server creates a dsmserve.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 problems are as a result of incorrectly set DSMI_CONFIG or DSMI_DIR,DSMI_LOG environment variables.

The environment variables are used by the IBM 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 Sun Solaris, the DSMI_* variables are set in the instance's userprofile file. For example: `/home/tsminst1/sql1lib/userprofile`

Windows On Windows, 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 `/home/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 The following procedure is an example process for debugging DSMI environment variable errors where the SQL error code (SQLERRMC) 409 is displayed when using the DSMAPIPW command:

1. Open the `/home/tsminst1/sql1lib/userprofile` file and examine the statements for accuracy. If you make any changes to this file, stop and restart the DB2 instance so that the changes are recognized. The userprofile file has statements similar to the following example text:

```
export DSMI_CONFIG=/home/tsminst1/tsminst1/tsmdbmgr.opt
export DSMI_DIR=/usr/tivoli/tsm/client/api/bin64/dsm.sys
export DSMI_LOG=/home/tsminst1/tsminst1
```

As an example, 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
```

2. Verify that the SERVERNAME entry in the `tsmdbmgr.opt` file matches the SERVERNAME entry in the `dsm.sys` file.
3. Run the DSMAPIPW command. You must be logged on using the root user ID.
4. If you can run the DSMAPIPW command, remove the `/home/tsminst1/tsminst1/tsmdbmgr.log` file while still using the root user ID to eliminate permission problems between the root user ID and the `tsminst1` user ID.

Troubleshooting error message ANR2971E using the SQL code

The message ANR2971E is an error message that might be displayed 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 faster.

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 DSMSEV 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   : BIJRPM01
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      sqlcabc: 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 API writes.

AIX **HP-UX** **Linux** **Solaris** On AIX, HP-UX, Linux, or Sun Solaris servers, this error is flagged because the DSMAPIPW 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 typically 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 => ? sq12071
```

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 DSMAPIPW 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.

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.

Analyzing the process symptoms to resolve problems

The process symptoms sometimes can indicate where your errors occur.

The following process symptoms might be encountered:

- 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 10 for the descriptions of the server processes.

Table 10. 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.	
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 EXPINTERVAL=<i>n</i> 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 WAIT=YES.

Table 10. IBM Tivoli Storage Manager server processes (continued)

Process or command	Description	Runs in the foreground or as a multiple process
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: IMPORT ADMIN, IMPORT NODE, IMPORT POLICY, and IMPORT SERVER.</p>	
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 MIGPROCESS 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 WAIT=YES.
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.

Table 10. IBM Tivoli Storage Manager server processes (continued)

Process or command	Description	Runs in the foreground or as a multiple process
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

The server might run tasks as processes. 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

BACKGROUND or FOREGROUND. 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

The “process ended” message that is issued varies, depending on whether or not the process needs to report information about no items or bytes processed, items and 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 might 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 might 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 might 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 FOREGROUND. 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.

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

The ANR1221E message appears in the form ANR1221E COMMAND: Process *processID* terminated - insufficient space in 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 by issuing 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 appears in the form "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 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`, 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` then an undamaged file might exist. If an undamaged file exists, issue the `UPDATE VOLUME volumeName ACCESS=DESTROYED` command and then issue `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 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

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 can recognize that an error occurred and can recover 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. 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

You might receive error messages CTGTRV009E and CTGTRV011E. These error messages depend on the data that is returned for IBM Tivoli Storage Manager server reports, client reports, or both under the Tivoli Integrated Portal reporting and monitoring feature. The cause for these errors is that the memory heap size must be increased 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

To increase the maximum memory heap size used by the Tivoli Integrated Portal, complete the following steps:

1. Open a shell window and change to the Tivoli Integrated Portal default installation directory, `<TIP_home>/products/tsm/bin`.

AIX **HP-UX** **Linux** **Solaris** For AIX, HP-UX, Linux, or Sun Solaris, issue the following command: `# cd /opt/tivoli/tsm/AC/products/tsm/bin`.

Windows For Windows, issue the following command: `> cd C:\Program Files\tivoli\tsm\AC\products\tsm\bin`

2. Run the supportUtil 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 the option **2. Manage the maximum memory size the Administration Center can use**.
 - b. Select the option **1. Update the maximum memory size the Administration Center can use**.
 - c. Enter **1024** for the value.
 - d. Select the option **99**.
 - e. Return to the main menu.
 - f. Select the option **9. Exit**.
5. Stop and restart Tivoli Integrated Portal to update the memory heap size.

Files are not expired after reducing versions

You might encounter an error 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 reasons show 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. In order 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 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 ACTIVE DATA 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 different from Read/Write, check that volume. A volume might be marked Read/Only or Unavailable because of a device error. If the device error was resolved, issue the `UPDATE VOLUME volname ACCESS=READWRITE` command to allow the server to 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.

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 will group 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 has “Collocate?=Yes” but volumes still contain data for many

The two possible reasons for this message are that data was stored on volumes in this storage pool prior to setting “Collocate=Yes” or the storage pool ran out of scratch tapes and stored data on the best possible volume, even though it ignored the request to collocate.

If data for multiple nodes ends up on the same volume for a storage pool with “Collocate=Yes,” this situation can be corrected by one of the following actions:

- Issue the `MOVE DATA` command for the volume or volumes affected. If scratch volumes are available or volumes with sufficient space are assigned to this client

node for collocating their data, the process reads the data from the specified volume and moves it to a different volume or volumes in the same storage pool.

- Allow migration to move all the data from that storage pool by setting the HIGHMIG and LOWMIG thresholds to accomplish this outcome. By allowing migration to migrate all data to the NEXT storage pool, the collocation requirements are honored, provided that the NEXT storage pool is set to "Collocate=Yes," it has sufficient scratch volumes, and it is assigned volumes to satisfy the collocation requirements.
- Issue MOVE NODEDATA for the client nodes whose data resides 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 MOVE NODEDATA process reads the data from the volumes that this node has data on and move it 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. Selecting an appropriate volume is significantly influenced by the number of scratch volumes in a storage pool. Another alternative is to explicitly define more volumes to the storage pool by issuing the DEFINE VOLUME command. Again, the key is to have candidate empty volumes for collocation to use rather than using a volume that already has data on it for a different node.

Resolving active data pool, storage issues

You might experience difficulty in storing data in an active data pool by using the simultaneous-write function or by issuing the COPY ACTIVEDATA 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 3. Resolving communication problems

The need for connectivity in 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 the server is configured for HTTP communications for administrative session and successfully binds to a port on startup, the following message is issued:
ANR8280I HTTP driver ready for connection with clients on port 1501.
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 TCPSPORT setting in the server options file is correct. If the code setting is inadvertently changed, the clients to 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 TCPSPORT and TCPADMINPORT for each server are unique. For example, there are two servers at the same TCP/IP address. The first server has a TCPSPORT of 1500 and a TCPADMINPORT of 1500. The second server has a TCPSPORT 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 that the server is running on 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 Tivoli Storage Manager server because the memory and resource allocation problem might indicate an error in either.

Determining 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 problems are shown in the following list:

Missing key ring database file (**cert.kdb**)

The server creates the cert.kdb key-ring database file if it does not exist. One or both of the SSLTCPSPORT 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 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 the following command:

```
gsk7capicmd -cert -add -db cert.kdb -pw password
-label name -file .der_file -format ascii -trust enable
```

The CA root certificate was not added to the client

Add the root certificate as trusted into the client key database:

```
gsk7capicmd -cert -add -db dsmcert.kdb -pw password  
-label my CA -file ca.arm -format ascii -trust enable
```

Unable to run gsk7capicmd.exe (IBM Global Security Kit [GSKit])

In most cases, this Windows error is generated by an incorrect environment setup. Please refer to the *IBM Tivoli Storage Manager Administrator's Guide* and set up the PATH variable as directed, before running the gsk7capicmd utility.

ANS1595E Bad server certificate

This error is reported when the server certificate is not known to the client. The "bad server certificate" error can occur under the following conditions:

- The certificate was never imported
- The certificate file (cert.arm) was corrupted before being imported
- The command to import the certificate was entered incorrectly
- The DSM_DIR variable points to wrong directory, which contains an incorrect client key database (dsmcert.kdb)

Repeat all the steps necessary for importing the server certificate and check the DSM_DIR variable. Refer to the client error log 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 general error occurs on the client and indicates that the SSL connection could not be 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 is trying to connect. Determine if the client points to the correct server port (TCPPort), which in most cases is a different port 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 appears to stop but will timeout through the server IDLEWAIT option, or end when the server administrator issues the CANCEL SESSION command to manually cancel it. The following 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 could cause a valid handshake process to take some time, do not assume that the output 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.

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 D, “IBM Global Security Kit return codes,” on page 223

Recovering the key database file password

If you forgot the current key ring 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 the key database file password and update it:

1. Issue the `QUERY SSLKEYRINGPW` command to display the current key ring database password.
2. Issue the following command to use the server record of the key ring 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 4. Resolving Administration Center problems

The Administration Center is a Web-based interface that can be used to centrally configure and manage 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 the 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 DSMADMC 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  
-----  
TCPPOINT          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 indicated by the address is accessible.
 - b. Issue the QUERY OPT TCPPOINT 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 DSMADMC 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 TCPPOINT
```

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.

```
dsmdm -id=<admin id> -pass=<admin password> -tcpport=<TCP/IP port>
-tcps=<TCP/IP address>
```

```
tsm: SERVER1>QUERY ACTLOG
```

10. Check the server's database file for corruption. The server's database file is located in the Tivoli Integrated Portal installation directory under *tip_installation_location\products\tsm\tsmservers.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 this 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.

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 Web site: <http://www.ibm.com/support/docview.wss?uid=swg21410467>.

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 of which that user ID is a member. If a user ID has multiple roles, all of the authorizations associated with all of the roles apply (a role can only grant authorization, not deny it).

The Administration Center defines two user roles. Both of the following roles have access to all of the components of the Administration Center:

- tsmAdministrator

- tsmUser

When you install the Administration Center, this creates a new user group 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 will have access to the Administration Center.

For more information about users, roles, and groups, refer to 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. 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 Web site: <http://www.ibm.com/support/docview.wss?uid=swg21410467>.

Creating 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 new user with the proper permissions.

Perform the following steps to create an 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. Determine the group to which the user ID belongs.
 - a. In the navigation tree, expand **Users and Groups**.
 - b. Click **Administrative Group Roles**.
 - c. Click the User name in the table.
 - d. 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 the problem can be reported.

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 81

Running the collector tool to obtain problem-analysis information

The collector tool gathers information about your WebSphere server installation. The tool then packages it in a Java archive (JAR) file that you can send to the WebSphere Customer Center. The center is designed to assist you in determining and analyzing your problem.

Information in the JAR file includes logs, 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. The collector tool can be run to collect summary data or to traverse the system to gather relevant files and command results. The collector tool produces a Java archive (JAR) file of information needed to determine and solve a problem. The collector summary option produces a lightweight collection of version and other information that is useful when first reporting the problem to the WebSphere Customer Center.

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.

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.

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 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 Java Development Kit (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, HP-UX, Linux, or Sun 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.

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 Web site: <http://www.ibm.com/support/docview.wss?uid=swg21410467>.

- 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.
 6. Make the work directory, the current 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.

7. Run the collector program by issuing the COLLECTOR command from the command line.

For example: c:\work>collector

Issuing the COLLECTOR command with no additional parameters gathers one copy of the node data and data from each server in the node. 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 output of the collector tool.

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, enabling it to download the latest symptom database from the IBM Web site.

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.

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 Web site: <http://www.ibm.com/support/docview.wss?uid=swg21410467>.

Using the log analyzer

You cannot view the service or activity log with a text editor. The log analyzer tool lets you view the file.

The Websphere server on which 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
 - **AIX** **HP-UX** **Linux** **Solaris** On all other 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, there are a few actions that can be taken.

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 use.

Log off from Tivoli Integrated Portal on a regular basis or if you determine that the memory usage of the Websphere Server is high. 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. The value can be changed with the Administration Center support utility.

Related reference

“Administration Center Support utility” on page 85

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 system 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 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 Web site: <http://www.ibm.com/support/docview.wss?uid=swg21410467>.

2. Locate the TCP/IP hosts file on your system.
For Windows systems, look in `WINNT\system32\drivers\etc`.
For AIX, Linux, and Solaris 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 shows you the 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.

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 Web site: <http://www.ibm.com/support/docview.wss?uid=swg21410467>.

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*\profiles\TIProfile\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\TIProfile\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. If you do not specify the user ID and password, you are prompted to enter them.

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 *dsmhealthmon.xml* and *dsmhealthmon_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 *dsmhealthmon.xml* or *dsmhealthmon_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 that 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 those messages in the trace file, it indicates that the health monitor is running. You might need to wait some time before seeing 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 from 0 - 10.
5. Repeat the above 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 of 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 all of the drives in the library offline
- If you take all of the paths for all of the drives in a single library offline
- If you take the only path to the library offline

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. Points of between 5 - 9 cause a warning, and points ten and higher indicate a critical status.

The databaseScore is calculated by a summation of the following rules:

For Tivoli Storage Manager Version 6.2 servers

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

For Tivoli Storage Manager Version 5.5 servers

Do not have a space trigger - 2 points

Extension of database is 0 and maximum utilization is at 90%: 5 points

No space trigger, no extendable space for the DB, and current utilization of DB over 90%: 10 points

Extension space exists and utilization is over 85% but less than 95%: 5 points

Extension space exists and utilization is over 94%: 10 points

No database backup in the last 24 hours: 3 points

Cache hit ratio is less than 98%: 3 points

Log utilization is currently over 80% but less than 90%: 3 points

Log utilization is currently over 89% but less than 95%: 5 points

Log utilization is currently over 94%: 5 points

Log maximum utilization is over 90%: 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.

As long as the password that is being updated is 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 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 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 Tivoli Storage Manager installation directory under `\products\tsm\bin`.

Windows On Windows, start the tool by issuing `supportUtil.bat`.

AIX **Linux** **Solaris** On AIX, Linux, or Sun Solaris, 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 cases where the failure message does not seem related to the behavior of the task, it does not include server messages at all, or 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 used to access the Administration Center might be running in different time zones. Activity log information is displayed using the server's time.

If errors are encountered during command routing, check the activity log of the source server first. If it 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 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 89

“Administration Center Support utility” on page 85

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 86

Checking the server activity log to resolve Administration Center problems

The Administration Center processes its work by issuing the 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

Some problems 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

The Administration Center allows you to open multiple work pages and multiple portlets within work pages. Updates to a server made in one portlet do not necessarily propagate to all other open portlets. In particular, the settings in a properties notebook are gathered when the properties notebook is first launched. If server objects in the properties notebook are created in a different portlet while the properties notebook is active, the newer objects will not be automatically added to the properties notebook. Likewise, a wizard that offers a selection of objects will only list those objects that existed at the start of the wizard. However, in some cases, object lists can be manually refreshed.

Some portlets contain many tabs and list many server objects and their attributes. It can be difficult to ascertain whether changes made in another table, notebook, or wizard affected objects in a particular portlet. Furthermore, commands issued directly to the server through the server console or command line client can also add objects that will not appear in a notebook or wizard started before the commands were issued. In general, it is best to limit the number of portlets running at once, and to not use the command line client or server console while using the Administration Center. If a missing or changed object becomes apparent, cancel and restart the current task.

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 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, 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 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 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 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 will typically be 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 will be 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 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, where *message_number* is the message that you want to
| see more details about.

Related reference

“Responding to Administration Center task failure messages” on page 86

Defining Tivoli Storage Manager messages

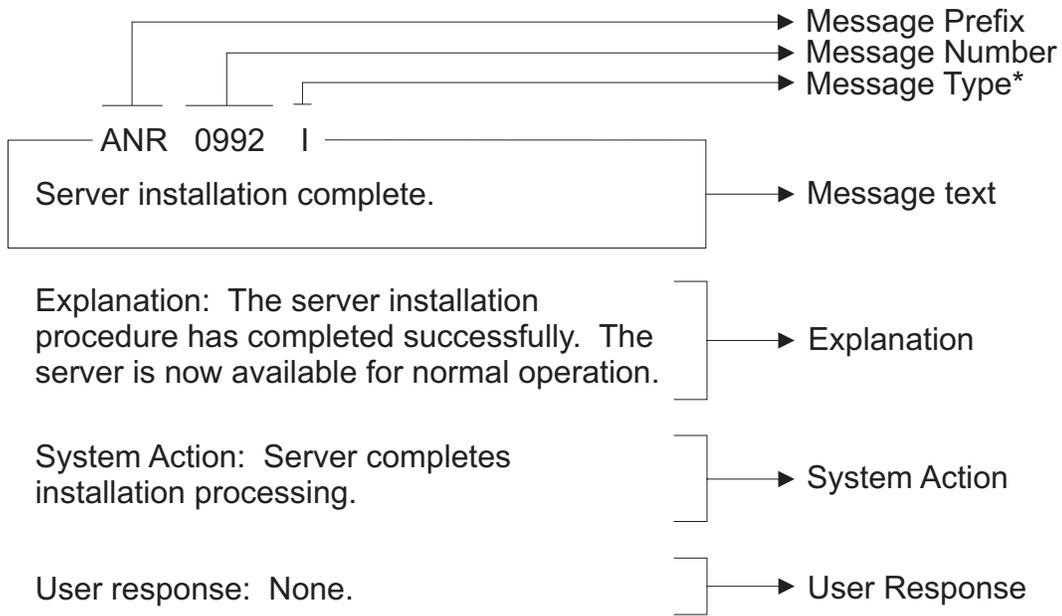
You can determine the kind of Tivoli Storage Manager message being displayed by looking at the message format. The first three characters of each message define what kind of message you are receiving.

The following examples illustrate the format used to describe the Tivoli Storage Manager messages:

- Messages that begin with an **ANE** prefix and are in the 4000-4999 range originate from the backup-archive client. These messages (or events) are sent to the server for distribution to various event-logging receivers.
- The client might send statistics to the server that is providing information about a backup or restore. These statistics are informational messages that might be enabled or disabled to the various event-logging receivers. These messages are not published here.
- Messages that begin with an **ANR** prefix originate from the server or storage agent.
- Messages that begin with an **ANS** prefix are from one of the following clients:
 - Administrative clients
 - Application program interface clients
 - Backup-archive clients
 - Space Manager (HSM) clients
- Messages that begin with an **ACD** prefix are from Data Protection for Lotus Domino®.
- Messages that begin with an **ACN** prefix are from Data Protection for Microsoft Exchange Server.
- Messages that begin with an **ACO** prefix are from Data Protection for Microsoft SQL Server.
- Messages that begin with an **ANU** prefix are from Data Protection for Oracle.
- Messages that begin with a **BKI** prefix are from Data Protection for SAP for DB2 UDB and Data Protection for SAP for Oracle.
- Messages that begin with a **DKP** prefix and are in the 0001-9999 range are from Data Protection for WebSphere Application Server.

Defining the message format

The following example displays the Tivoli Storage Manager message format:



- * I = Information
- E = Error
- S = Severe Warning
- W = Warning
- K = Kernel message that originates from the hierarchical storage management (HSM) client

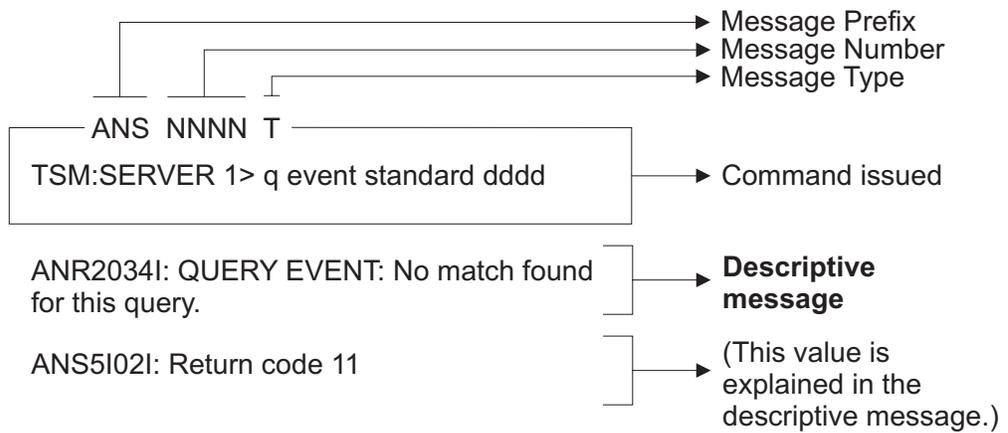
Message variables in the message text appear in italics. The server and client messages fall into the following categories:

- Common messages that pertain to all Tivoli Storage Manager server platforms
- Platform-specific messages that pertain to each operating environment for the server and the client
- Messages that pertain to application clients

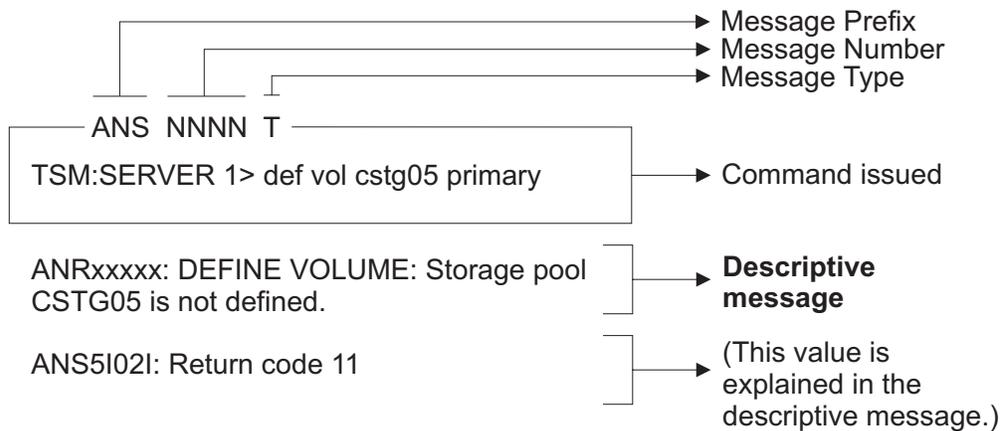
Reading a return code message

Many different commands can generate the same *return code*. The following examples are illustrations of two different commands that are issued that result in the same return code; therefore, you must read the *descriptive message* for the command.

The QUERY EVENT command:



The DEFINE VOLUME command:



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 API and that it has the platform-specific command file and `dsmcmd.xml` 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` 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:

```
\products\tsm\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.

Related reference

“Using data storage diagnostic tips” on page 185

Resolving backup-archive client deployment problems

Some Administration Center backup-archive client deployment feature problems are attributed to missing prerequisites.

Ensure that the following items are completed:

- Install the IBM Tivoli Storage Manager server V6.2 or later.
- Install the Tivoli Storage Manager V6.2 or later 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. 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 Web site:
<http://www.ibm.com/support/docview.wss?uid=swg21410467>.

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 Web site:
<http://www.ibm.com/support/docview.wss?uid=swg21410467>.

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 host or IP address of the server so you can see the deployment results.
2. 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.
3. 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.
4. 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 Manger** → **Manage Servers**. Select a server and select **Manage Client Auto Deployments** from the action menu.

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 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 any of the following conditions is true:

- The client installation program detects that files that are in use must be removed or replaced.
- The Tivoli Storage Manager Logical Volume Snapshot Agent (LVSA) feature is present.

| **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 Web site:
| <http://www.ibm.com/support/docview.wss?uid=swg21410467>.

Chapter 5. Resolving Data Protection problems

The Data Protection client uses the Tivoli Storage Manager API to communicate with the Tivoli Storage Manager server to provide data management functions.

Troubleshooting IBM Tivoli Storage Manager for Enterprise Resource Planning

Information on how to resolve errors that might occur during IBM Tivoli Storage Manager for Enterprise Resource Planning operations is provided.

Troubleshooting IBM Tivoli Storage Manager for Enterprise Resource Planning common problems

Information on how to resolve errors that might occur during IBM Tivoli Storage Manager for Enterprise Resource Planning operations is provided.

Random problems

If a problem occurs inconsistently, try to determine what the difference is when the problem occurs, if any. Compare the log files of the application in question to find out the differences between successful and unsuccessful operations. Look for one of these patterns when the problem occurs:

- The problem always occurs at the same time. If this is true, view the appropriate log files to determine review if there are any scheduled processes occurring simultaneously such as virus checker, automatic updates, or batch jobs.
- The problem always occurs after another operation is performed or the same operation is performed.
- The problem occurs when another application or process is performed in parallel.

Data Protection for SAP for Oracle log files:

- brbackup / brrestore log
- sbtio.log

Data Protection for SAP for DB2 log files:

- tdpdb2.SID.nodename.log
- db2diag.log

Reproducible (repeatable) problems

When encountering a problem that occurs during an operation that has previously performed successfully, consider these possible causes:

- The IBM Tivoli Storage Manager for Enterprise Resource Planning setup changed.
- One (or more) of the Oracle, DB2, SAP, Tivoli Storage Manager, operating system, network, or hardware components changed.
- Patches or updates to one (or more) of the components were applied.
- Changes originated by the system have occurred such as these:
 - Check if the disks are running full with the UNIX or Linux df command.

- If network performance has decreased, check if additional hosts, additional applications, or defects in software or hardware occurred. Compare operation runs in the Administration Assistant Performance Monitor history view or compare the Data Protection for SAP for Oracle brbackup / brrestore log files or Data Protection for SAP for DB2 tdpdb2.SID.nodename.log file.
- If Tivoli Storage Manager server processing has decreased, check if additional clients or additional operations were added. Information is also available in the Tivoli Storage Manager server activity log.

When none of these possible causes has occurred, view the last modified time stamp of the configuration files for Data Protection for SAP for Oracle (`initSID.utl`, `initSID.sap`, `dsm.sys`, `dsm.opt`, `/etc/services`, `/etc/inittab`, ...) or Data Protection for SAP for DB2 (`vendor.env`, `initSID.utl`, `dsm.sys`, `dsm.opt`, `/etc/services`, `/etc/inittab`, ...). This UNIX or Linux command lists all files in the `/etc` directory which have been modified during the previous five days:

```
find /etc -type f -ctime 5 -print
```

If you are able to identify changes made to the system, roll them back one at a time and try to reproduce the problem. This method frequently reveals which change or set of changes caused the problem.

Internet Protocol version 6 (IPv6) support

Data Protection for SAP supports both IPv4 and IPv6 for internal communication in that it will run in IPv4, IPv6, and mixed environments on AIX and Linux. However, these products do not exploit new IPv6 functionality. In a mixed environment, the communication depends on the adapter network settings. There is no option to enforce the use of a specific protocol other than by network configuration. Specifically, the ProLE or acsd service will listen for both IPv4 and IPv6 connection requests if the system is configured accordingly. Connection requests to ProLE are made for the addresses returned by the system for the respective port on the local host. Connection requests to other machines such as the Administration Assistant function for Data Protection for SAP are made for the addresses specified by the user. IPv6 addresses are supported when TCP/IP addresses are specified in a command line or in a profile parameter such as `TCP_ADDRESS`. However, when the IP address and port are specified in the `IPv4 address:service or port` format, then the format needs to be changed to `service or port@<IP address>` if the IP address is specified in the IPv6 notation. In the case of a dotted decimal IPv4 address, the traditional format can still be used.

The specification of IPv6 addresses assumes that Data Protection for SAP is used in an environment in which IPv6 is supported by all hardware and software components involved and has been adequately tested in this environment.

Understanding the Setup

Review these considerations to better understand the installation setup on UNIX or Linux systems:

- Make sure all files are installed as described in the *User's Guide*.
- Make sure an entry similar to this example is defined in the `/etc/inittab` file:

Data Protection for SAP for Oracle:

```
po64:2:respawn:/usr/tivoli/tsm/tdp_r3/ora64/prole -p tdpr3ora64
Server component hostname 5126
```

Data Protection for SAP for DB2:

```
pd64:2:respawn:/usr/tivoli/tsm/tdp_r3/db264/prole
-p tdpr3db264 Server component hostname 5126
```

The <Server component hostname> specifies the name or IP address of the host on which the Administration Assistant Server component is running. 5126 is the default port to which the Server component listens. These Server arguments are only needed when using the Administration Assistant. The purpose of this entry is to start a daemon process for ProLE. This process listens on the Data Protection for SAP for Oracle 64-bit port (tdpr3ora64) for backint and RMAN connections and Data Protection for SAP for DB2 (tdpr3db264) connections with the shared library and sends performance-related information to the Administration Assistant Server component. The port can have a different name; however, the name must match the name in the /etc/services file as shown in this example:

Data Protection for SAP for Oracle:

```
tdpr3ora64      57323/tcp
```

Data Protection for SAP for DB2:

```
tdpr3db264     57324/tcp
```

These lines are added to the /etc/services file by the installer.

- Make sure the Data Protection for SAP for Oracle configuration file `initSID.utl` is located in the `$ORACLE_HOME/dbs` directory.
- When using the BR*Tools with Data Protection for SAP for Oracle, modify the `initSID.sap` file by setting `backup_dev_type = util_file` and variable `util_par_file` to the fully qualified path and file name of `initSID.utl`.

For an overview of the configuration files on a UNIX or Linux system, see the following sections:

- (Data Protection for SAP for Oracle) Figure 1 on page 100
- (Data Protection for SAP for DB2) Figure 2 on page 101

Review these considerations to better understand the installation setup on Windows systems:

- Make sure all files are installed as described in the *User's Guide*.
- Verify that service ProLE Service is running and set to automatic startup. If this service is not running, Data Protection for SAP does not function properly.
- The installer adds lines to the `%SYSTEMROOT%\system32\drivers\etc\services` file similar to these:

(Data Protection for SAP for Oracle)

```
tdpr3ora64      57323/tcp
```

(Data Protection for SAP for DB2)

```
tdpr3db264     57324/tcp
```

The lines `tdpr3ora64` and `tdpr3db264` are the Data Protection for SAP 64-bit ports. For Data Protection for SAP for Oracle, this port name is also needed for the `initSID.sap` file when RMAN is configured.

- Make sure the Data Protection for SAP configuration file `initSID.utl` is located in the `%ORACLE_HOME%\database` directory (on Oracle) and in the directory pointed to by the `TDP_DIR` environment variable (on DB2).
- When using the BR*Tools with Data Protection for SAP for Oracle, modify the `initSID.sap` file by setting `backup_dev_type = util_file` and variable `util_par_file` to the fully qualified path and file name of `initSID.utl`.
- The vendor environment file `vendor.env` must contain the fully qualified path and file name of the `initSID.utl` file for Data Protection for SAP for DB2.
- The vendor environment file `vendor.env` should contain the path of the location where the Data Protection for SAP for DB2 run logs are written. If this location is not specified, temporary directory of the machine is used.

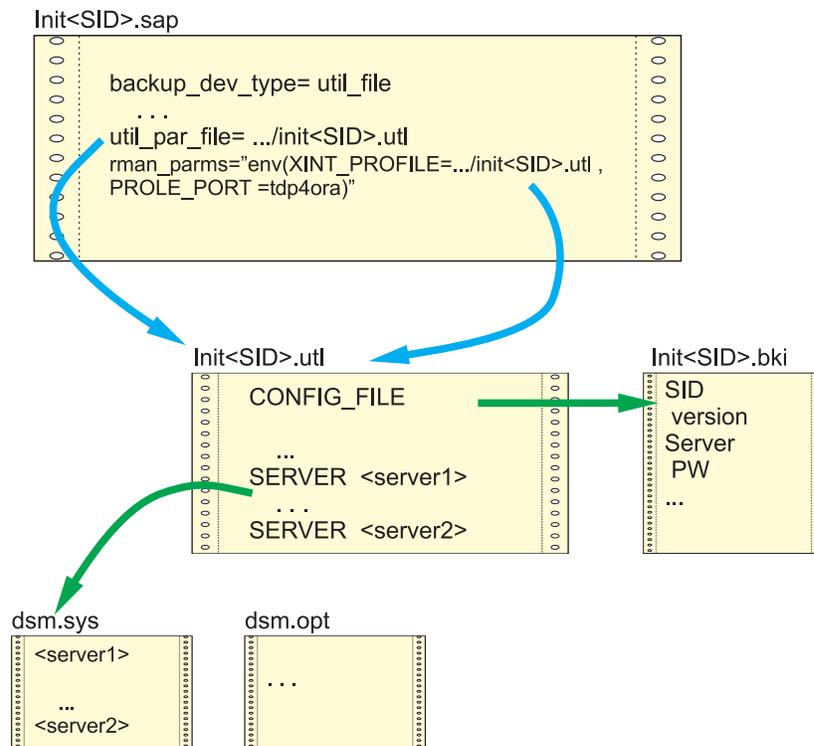


Figure 1. SAP® and Data Protection for SAP for Oracle configuration files on UNIX or Linux

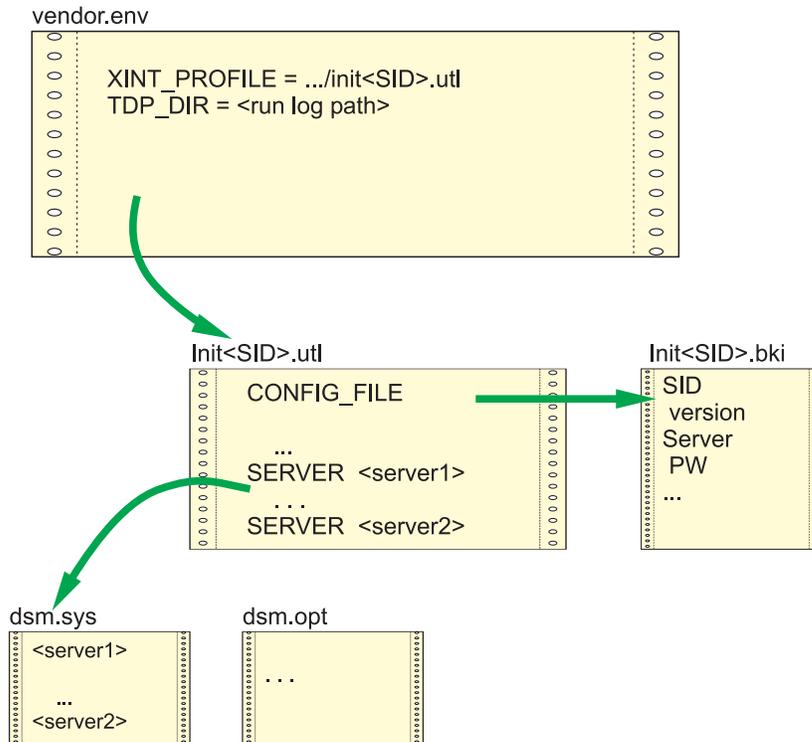


Figure 2. SAP® and Data Protection for SAP for DB2 configuration files on UNIX or Linux

On UNIX or Linux systems, the names of the Tivoli Storage Manager servers specified in `initSID.utl` must match the names in the `dsm.sys` file. If the Tivoli Storage Manager API or Tivoli Storage Manager Backup Archive Client were installed into their default locations, then the `DSMI_*` variables do not need to be set. If the variables are set, however, make sure they specify the correct directories and files. The user ID that runs the backups must have the correct permissions to access all of files and directories specified by these variables. Also verify that write permissions exist for the `initSID.bki` file as this is the only file to which Data Protection for SAP writes persistent information.

On Windows systems, the `dsm.opt` file is used instead of the `dsm.sys` file. However, the content of this file is not relevant to Data Protection for SAP. The directory that contains the `dsm.opt` file must also contain a `server.opt` file for each server specified in the `initSID.utl` file. The environment variable `DSMI_CONFIG` must specify an option file within this directory. `DSMI_CONFIG` should specify the `dsm.opt` file in this directory. The `DSMI_DIR` environment variable must also specify the directory where the Tivoli Storage Manager API message text file resides. This is typically the `c:\Program Files\Tivoli\tsm\api64` directory.

Providing information to IBM or Tivoli support

Provide this information when contacting IBM or Tivoli support:

- The Data Protection for SAP version.
- The operating system level and patches that were applied.
- The Oracle or DB2 database version
- The Tivoli Storage Manager server version.
- The Tivoli Storage Manager server operating system level.

- Data Protection for SAP for DB2 configuration files (vendor.env, initSID.utl or Data Protection for SAP for Oracle configuration file (initSID.utl), including Tivoli Storage Manager client configuration files (dsm.sys, dsm.opt).
- On Data Protection for SAP for DB2, the SAP®-DB2 Administration Tools log file. On Data Protection for SAP for Oracle, the BR*Tools output for the operation in question (brarchive, brrestore)
- The change history of the system components (if the process worked previously).

Additional information might also be requested from the service representative.

Troubleshooting Data Protection for SAP for DB2 problems

Information on how to resolve errors that might occur during Data Protection for SAP for DB2 operations is provided.

General problem resolution

The following graphic (Figure 3) will help you to isolate problems that occur when backing up or restoring of your DB2 database.

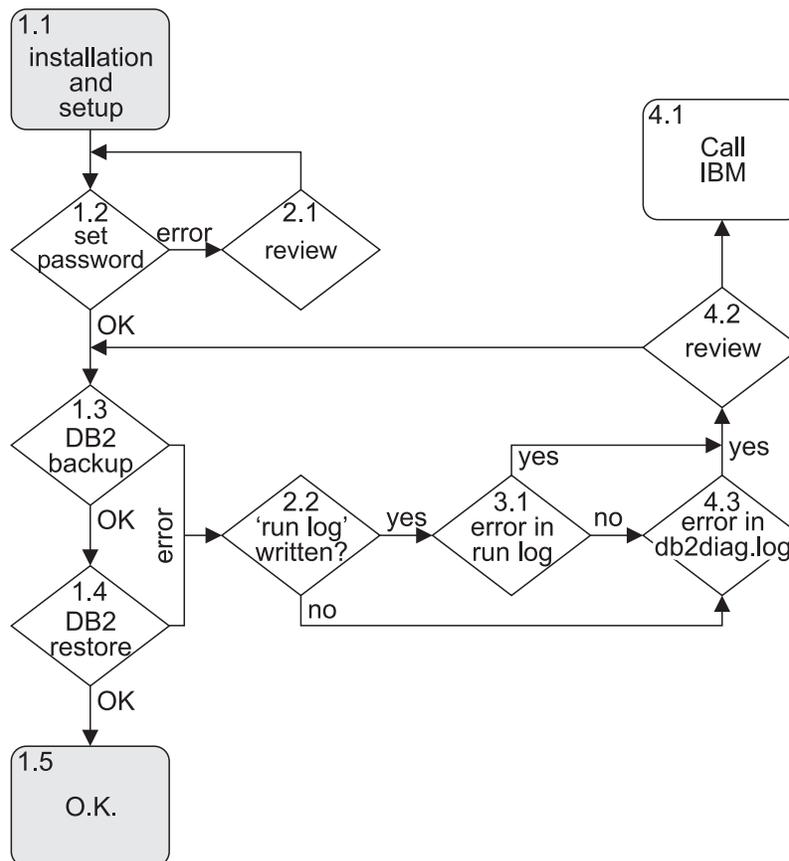


Figure 3. General Problem Isolation

After installation is completed (Step 1.1) and manual password handling is specified, set the password (Step 1.2) as described in the *User's Guide*. When the operation completes successfully, the informational messages BKI0051I: Password successfully verified for node *NODENAME* on server *SERVERNAME* and BKI0024I: Return code is: 0. display for each server configured within the *initSID.utl* file.

An error message displays when a problem occurred. The Administration Assistant can also be used. The Configurator feature loads the configuration of the node on which problems are encountered and allows the Administration Assistant to check the configuration.

These errors are frequently encountered at Step 1.2:

**BKI2001E: Socket error while connecting to ProLE at IP-Address:PORT:
Connection refused**

On Windows, verify that the ProLE Service is running by viewing the Computer Management Services screen or issue this command:

```
net start
```

A list of all running services displays. On UNIX or Linux, verify that the background daemon is running by issuing this command:

```
ps -ef | grep prole
```

Check the entry in `/etc/services` (UNIX or Linux) and `%SYSTEMROOT%\system32\drivers\etc\services` (Windows). Compare the port number from the error message with the port number within `/etc/services`. Also check the entry in `/etc/inittab` (UNIX or Linux). If another port was set using the option `-pPORT`, check this as well. If all of this will not help, start the ProLE from another shell on UNIX or Linux with this command:

```
prole -p PORT
```

Issue this command on Windows:

```
prole -console -p PORT
```

Attempt to start backom again.

BKI5001E: Tivoli Storage Manager Error: Server not found in configuration file On UNIX or Linux, the Tivoli Storage Manager server defined in the `initSID.utl` file does not match the server specified in the `dsm.sys` file. On Windows, the `server.opt` file might be missing.

BKI5001E: Tivoli Storage Manager Error: ANS1353E (RC53) Session rejected: Unknown or incorrect ID entered

This message can display when the node in the server stanza of the UTL file is not valid on the server.

HANG If backom hangs after the password is entered, the server IP address specified in the UNIX or Linux `dsm.sys` file might be incorrect.

When Step 1.2 (setting the password) is successful, proceed to Step 1.3 and perform a backup using the DB2 backup command to verify the settings are correct as described in the *User's Guide*. If the backup was successful you will see a message from DB2:

```
Backup successful. The timestamp for this backup image is: timestamp
```

If an error message displays, view the error description in the *User's Guide* for information regarding how to resolve it.

When an error occurs, always view the Data Protection for SAP run log `tdpdb2.SID.nodename.log` first. This log file is located in the directory specified by the `TDP_DIR` environment variable. If the variable is not specified, the log file resides in the system temporary directory. If the `tdpdb2.SID.nodename.log` file does not exist (Step 2.2), then either DB2 was unable to load the shared library that contains the DB2 connector for Data Protection for SAP or an error was encountered before calling the Data Protection for SAP library. In both cases, a DB2 error message should display on the command line that begins with the SQL prefix and is also written in the DB2 diagnostic log `db2diag.log` (Step 4.3). DB2 provides detailed error descriptions by issuing this command:

```
db2 ? SQLnumber
```

Replace *number* with the appropriate message number. Try to resolve this problem using the DB2 documentation.

If the `tdpdb2.SID.nodename.log` file exists, search for a message beginning with `BKIXXXXY` where `XXXX` is a four digit number and `Y` is the letter I, W, or E. When such a message occurs, the DB2 connector for Data Protection for SAP loaded correctly was called by DB2. In Step 3.1, the `tdpdb2.SID.nodename.log` file is created and an error message starting with `BKI` is recorded.

Location of log files

Text displayed on the screen during DB2 backup, DB2 restore, and BackOM operations are typically written to log files. DB2 also writes messages of internal operations, events, or status in the administration notification log file (`db2SID.nfy`) and diagnostic log file (`db2diag.log`). These log files reside in the directory specified with the DB2 database management configuration parameter `DIAGPATH`. Query the DB2 database management configuration with this command:

```
db2 get dbm cfg
```

Information about how to locate these log files is available in the *User's Guide*.

DB2 vendor reason codes

Data Protection for SAP for DB2 uses these reason codes which might also be displayed or logged by DB2 in the case of problems.

Table 11. DB2 Vendor Reason Codes

Reason Code	Explanation	User Response
1	The library specified could not be loaded.	Check the DB2 diagnostic log for further details.
2	Communication error between shared library and ProLE	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
6	Object specified cannot be found on Tivoli Storage Manager.	There is no backup image on Tivoli Storage Manager matching the given search criteria.

Table 11. DB2 Vendor Reason Codes (continued)

Reason Code	Explanation	User Response
10	Invalid options specified with the options parameter of the DB2 backup/restore command.	Check the options string specified and check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
11	Initialization procedure for shared library failed.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> and the DB2 diagnostic log file for further details.
17	During end processing of either backup/archive or restore/retrieve session(s), an error occurred.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
18	An error occurred during reading or writing data from or to Tivoli Storage Manager.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
26	An error occurred during deleting data from Tivoli Storage Manager.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
29	An abort request from DB2 could not be handled correctly.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> and the DB2 diagnostic log file for further details.
30	A severe error occurred.	Check the DB2 diagnostic log file for further details.

Troubleshooting Data Protection for SAP for Oracle problems

Information on how to resolve errors that might occur during Data Protection for SAP for Oracle operations is provided.

Location of log files

Text displayed on the screen during `brbackup`, `brrestore`, and SAP® Tools operations are typically written to a log file. Oracle also writes internal operations in the alert log and core files that reside in the directory specified in the Oracle control files, for example `$SAPDATA_HOME/saptrace/background/alert_SID.log`. Information about how to locate these log files is available in the *User's Guide*.

Messages

During BR*Tools processing, logs that contain all issued messages are written to paths `/oracle/SID/sapbackup` (for `BRBACKUP`) or `/oracle/SID/saparch` (for `BRARCHIVE`). The message prefix indicates the issuing components. Refer to the documentation for the component that issued the message for detailed information. However, the following prefixes are used when employing BR*Tools with Data Protection for SAP for Oracle:

Table 12. Prefixes when using BR*Tools

Prefix	Issuing Component
ANS / ANR	Tivoli Storage Manager
BKI	Data Protection for SAP
BR	BR*Tools
ORA	Oracle database kernel
RMAN	RMAN

File Manager

The most important requirement for File Manager is that Data Protection for SAP for Oracle is set up correctly. This is especially true in regard to the backint executable file, as this file must be able to connect to the Tivoli Storage Manager server without errors. If this call fails, the File Manager displays an error message but does not analyze the reason for the failure. To analyze the error, invoke backint manually with the inquire function and check the output for error messages.

BACKINT problem resolution

Figure 4 on page 107 displays how to isolate the problem once the settings performed by the installer are verified. Make sure the BACKINT interface is working correctly before examining the RMAN interface.

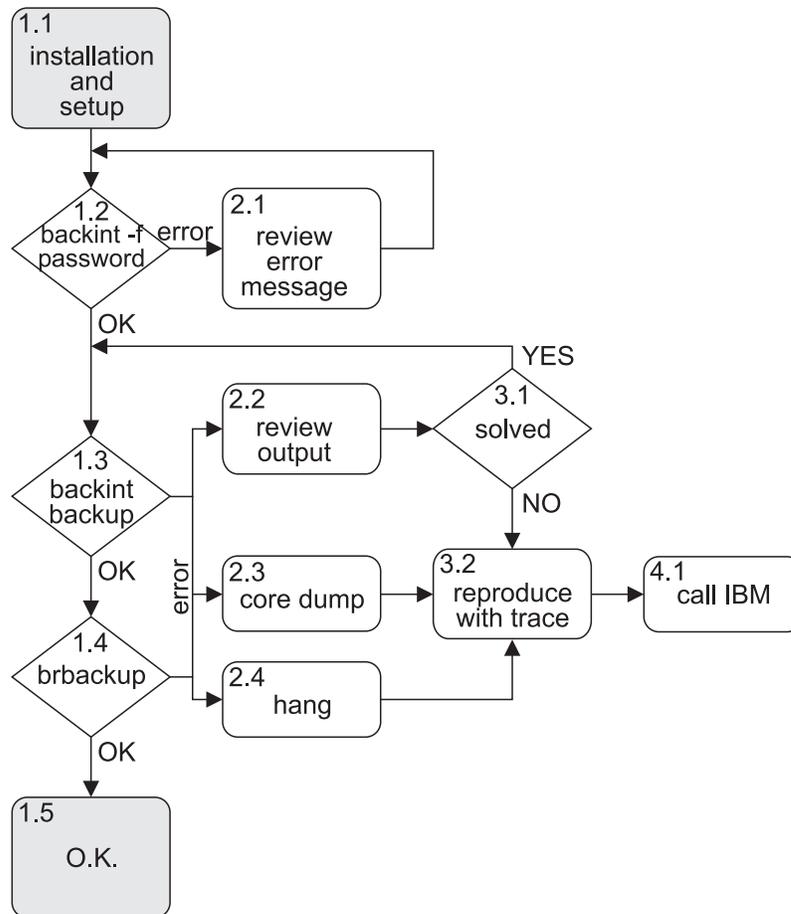


Figure 4. Problem Isolation for Backint

After installation is completed (Step 1.1) and manual password handling is specified, set the password (Step 1.2) as described in the *User's Guide*. When the operation completes successfully, the informational messages BKI0051I: Password successfully verified for node *NODENAME* on server *SERVERNAME* and BKI0024I: Return code is: 0. display for each server configured within the *initSID.utl* file. An error message displays when a problem occurred. The Administration Assistant can also be used. The Configurator feature loads the configuration of the node on which problems are encountered and allows the Administration Assistant to check the configuration.

These errors are frequently encountered at Step 1.2:

BKI2001E: Socket error while connecting to ProLE at IP-Address:PORT: Connection refused

On Windows, verify that the ProLE Service is running by viewing the Computer Management Services screen or issue this command:

```
net start
```

A list of all running services displays. On UNIX or Linux, verify that the background daemon is running by issuing this command:

```
ps -ef | grep prole
```

Check the entry in `/etc/services` (UNIX or Linux) and `%SYSTEMROOT%\system32\drivers\etc\services` (Windows). Compare the port number from the error message with the port number within `/etc/services`. Also check the entry in `/etc/inittab` (UNIX or Linux). If another port was set using the option `-pPORT`, check this as well. If all of this will not help, start the ProLE from another shell on UNIX or Linux with this command:

```
prole -p PORT
```

Issue this command on Windows:

```
prole -console -p PORT
```

Attempt to start backint again.

BKI5001E: Tivoli Storage Manager Error: Server not found in configuration file On UNIX or Linux, the Tivoli Storage Manager server defined in the `initSID.utl` file does not match the server specified in the `dsm.sys` file. On Windows, the `server.opt` file might be missing.

BKI5001E: Tivoli Storage Manager Error: ANS1353E (RC53) Session rejected: Unknown or incorrect ID entered

This message can display when the node in the server stanza of the UTL file is not valid on the server.

HANG If backint hangs after the password is entered, the server IP address specified in the UNIX or Linux `dsm.sys` file might be incorrect.

When Step 1.2 (setting the password) is successful, proceed to Step 1.3 and perform a backup using backint to verify the settings are correct as described in "Backup function" on page 111. When the backup completes successfully, the message `#SAVEDBIDFILENAME` displays for each saved file and `BKI0024I: Return code is: 0` also displays. If an error message displays, view the error description in the *User's Guide* for information regarding how to resolve it. At this point, the primary Data Protection for SAP for Oracle setup is almost complete. The BR*Tools and Oracle (when using RMAN) must also be configured correctly. Proceed to Step 1.3 and start brbackup as described in the *User's Guide*.

RMAN problem resolution

The following graphic (Figure 5 on page 109) will help you to isolate problems that occur when using RMAN.

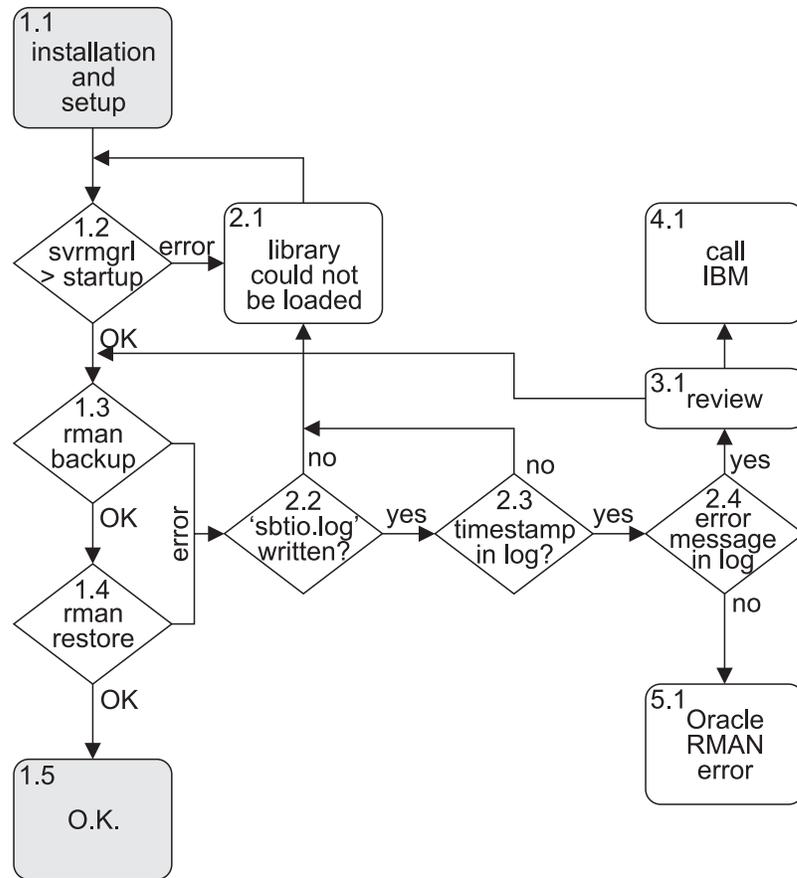


Figure 5. Problem Isolation for RMAN

After Data Protection for SAP for Oracle and Oracle are configured to work together (Step 1.1 in Figure 5), attempt to start Oracle using the server manager svrmgr1 (on UNIX or Linux) or svrmgr30 (on Windows) with Oracle 8.x. Use SQL Plus (sqlplus) with Oracle 9.x. When an error occurs while using RMAN, always view the sbtio.log file first. This file is located in the directory specified by the user_dump_dest keyword in the Oracle initSID.ora profile (located at \$ORACLE_HOME/saptrace/usertrace/sbtio.log by default). If the sbtio.log file does not exist (Step 2.2), then either Oracle was unable to load the shared library that contains the RMAN connector for Data Protection for SAP or an error was encountered before the Data Protection for SAP library was called. In both cases, an Oracle error message should exist in the brbackup log file that begins with ORA-, PLS-, or RMAN-. Try to resolve this problem using the Oracle and SAP® documentation. If the sbtio.log file exists, search for a message beginning with BKIXXXXY where XXXX is a four digit number and Y is the letter I, W, or E. When such a message occurs, the RMAN connector for Data Protection for SAP loaded correctly and was called by RMAN. This should be the first message for every new session in Step 2.3:

```
BKI7060I: Data Protection for SAP version and build number session: 764
```

If this message is not available, Oracle loaded an incorrect library.

Perform these tasks on Windows when an incorrect library is loaded by Oracle:

1. Remove or rename all occurrences of the file `orasbt.dll` except the one in the Data Protection for SAP installation directory. Then copy this one to `%ORACLE_HOME%\bin`.
2. Stop the `OracleServiceSID` and restart it.

Several factors must be considered when an incorrect library is loaded by Oracle on UNIX or Linux. For example, the RMAN library `libtdp_r3.ext` is not located by the Oracle executable. Oracle suggests to use the `SBT_LIBRARY` variable to specify the library. However, do not use this variable for a version of Oracle prior to Oracle 9.2. Be aware that Oracle recommends not to issue the `make` command as described in the *User's Guide*. However, this recommendation is not applicable for all combinations of operating system and Oracle combinations. As a result, issuing the `make` command on any UNIX or Linux system with Data Protection for SAP is acceptable. When issued correctly, this command can confirm that the library and the Oracle executable are compatible. Also, make sure the library and soft link entered during the command exists and that the soft link is valid:

```
make -f ins_rdbms.mk ioracle LLIBMM=lib or link
```

It might be helpful to add the location of the link or library to the `LIBPATH` environment variable (on AIX) or to the `LD_LIBRARY_PATH` environment variable (on other UNIX or Linux systems).

On Windows based systems, the location of `orasbt.dll` must be in the `PATH`. Also, ensure that you have only one `orasbt.dll` in your system's `PATH`. It will be helpful as well to review the setup procedure according to the *User's Guide* and the information given in that chapter. Also, check if a core file is written or if Oracle has written a trace within the `saptrace/usertrace` directory.

In (Step 2.4) the file `sbtio.log` is written and you find an error message starting with `BKI` see the *User's Guide*. Using the `backint` executable file to determine any problems may make it easier because you can see all messages on the screen. Also, you will not disturb Oracle if something goes wrong. If `backint` is working as expected, return to problem determination with RMAN.

When isolating a problem with Data Protection for SAP and RMAN, you can follow the same steps as in "BACKINT problem resolution" on page 106. There must be a connection established to ProLE and the Tivoli Storage Manager server, and a password must be set (using `backint`) as well. If some of these steps fail, you will get exactly the same error messages with RMAN as you get with `backint`. Use the *User's Guide* to review these messages.

Manually invoke Data Protection for SAP for Oracle

Information about how to invoke Data Protection for SAP for Oracle from the command line to assist with troubleshooting efforts is provided.

Data Protection for SAP for Oracle is typically invoked by the BR*Tools utilities with a set of appropriate parameters. For troubleshooting purposes, call Data Protection for SAP directly from the command line:

```
backint -?
```

This command displays a list of all possible Data Protection for SAP command line options. It enables you to manually perform data protection operations which can assist with correcting errors. For the C shell, enclose the option string in quotes (`backint '-?'`).

Backup function:

Data Protection for SAP for Oracle is typically invoked by the BR*Tools utilities with a set of appropriate parameters. For troubleshooting purposes, call Data Protection for SAP directly from the command line:

```
backint -?
```

This command displays a list of all possible Data Protection for SAP command line options. It enables you to manually perform data protection operations which can assist with correcting errors. For the C shell, enclose the option string in quotes (backint '-?').

The backup function is typically invoked by the SAP® database utilities BRBACKUP and BRARCHIVE. These programs provide an input file (in the case of backup and inquire) to Data Protection for SAP that contain the names and paths of the database files to be processed. For troubleshooting purposes, however, it might be necessary to directly call this Data Protection for SAP function directly in order to back up individual files as shown in these examples. Issue this command on UNIX or Linux systems:

```
backint -p /oracle/SID/dbs/initSID.utl -f backup
```

Issue this command on Windows systems:

```
backint -p drive: or UNC name\orant\database\initSID.utl -f backup
```

The Data Protection for SAP profile `initSID.utl` has to be specified with the path and file name statement as shown in the examples. The program prompts you to enter one (or more) file names. Every successful backup operation (collection of one or more files) is allocated its own backup ID within Tivoli Storage Manager. Remember to press CTRL + D (on a UNIX or Linux system) or CTRL + Z (on a Windows system) after the file name to backup has been entered.

Delete function:

Data Protection for SAP for Oracle is typically invoked by the BR*Tools utilities with a set of appropriate parameters. For troubleshooting purposes, call Data Protection for SAP directly from the command line:

```
backint -?
```

This command displays a list of all possible Data Protection for SAP command line options. It enables you to manually perform data protection operations which can assist with correcting errors. For the C shell, enclose the option string in quotes (backint '-?').

The delete function is used as part of the Data Protection for SAP version control mechanism and can only be called by Data Protection for SAP or by a user. The delete function allows you to delete individual files, which might be necessary for recovering from an error. This function can be invoked from the command line as shown in these examples. Issue this command on UNIX or Linux systems:

```
backint -p /oracle/SID/dbs/initSID.utl -f delete
```

Issue this command on Windows systems:

```
backint -p drive: or UNC name\orant\database\initSID.utl -f delete
```

You will be prompted to enter the backup ID and the full file names of the files to be deleted.

Inquire function:

Data Protection for SAP for Oracle is typically invoked by the BR*Tools utilities with a set of appropriate parameters. For troubleshooting purposes, call Data Protection for SAP directly from the command line:

```
backint -?
```

This command displays a list of all possible Data Protection for SAP command line options. It enables you to manually perform data protection operations which can assist with correcting errors. For the C shell, enclose the option string in quotes (backint '-?').

The inquire function (typically invoked by BR*Tools and BRRESTORE) is used to query the Tivoli Storage Manager server for backup IDs or files which belong to a particular backup ID. For troubleshooting purposes, however, it might be necessary to invoke this function manually as shown in these examples. Issue this command on UNIX or Linux systems:

```
backint -p /oracle/SID/dbs/initSID.utl -f inquire
```

Issue this command on Windows systems:

```
backint -p drive: or UNC name\orant\database\initSID.utl -f inquire
```

Data Protection for SAP prompts you to enter the inquiry in one of these four formats:

- **#NULL:** Display all backup IDs that have been saved to this point. A typical line of the response could be #BACKUP JE0__A0DNE9Z74C. The backup ID in this case is JE0__A0DNE9Z74C (#BACKUP does not belong to the backup ID). The first six characters are the user defined prefix (see BACKIDPREFIX as described in the *User's Guide*). The next ten characters after this represent a unique ID of the backup.
- **BackupID:** Display all of the files which belong to that backup ID. A typical result could be ##BACKUP JE0__A0DNE9Z74C /oracle/C21/dbs/initC21.utl.
- **#NULL filename:** Display all of the backup IDs corresponding to the specified file. *Filename* requires an input consisting of path and name of the file.
- **BackupID filename:** Verify whether a particular file has been saved under a certain backup ID. *Filename* requires an input consisting of path and name of the file.

Restore function:

Data Protection for SAP for Oracle is typically invoked by the BR*Tools utilities with a set of appropriate parameters. For troubleshooting purposes, call Data Protection for SAP directly from the command line:

```
backint -?
```

This command displays a list of all possible Data Protection for SAP command line options. It enables you to manually perform data protection operations which can assist with correcting errors. For the C shell, enclose the option string in quotes (`backint '-?'`).

The restore function is typically started by the SAP® database utility BRRESTORE. For troubleshooting purposes, however, it might be necessary to directly call this Data Protection for SAP function directly in order to restore individual files as shown in these examples. Issue this command on UNIX or Linux systems:

```
backint -p /oracle/SID/dbs/initSID.utl -f restore
```

Issue this command on Windows systems:

```
backint -p drive: or UNC name\orant\database\initSID.utl -f restore
```

You will be prompted to enter the backup ID and the full file names of the files to be restored. If the files are to be restored to another directory, it is necessary to specify the path to the input files. If a file is restored directly, any existing file with the same name will be overwritten without warning. Thus, it is recommended that you restore database files directly only in a very controlled manner, when it is absolutely necessary in order to remove an error. In normal operation, a database should never be restored directly because this could corrupt the SAP database.

Data Protection for Exchange with VSS backup-restore support

Data Protection for Exchange provides support for protecting Microsoft Exchange databases through two different methods. The most common method is through the Microsoft Legacy API. Data Protection for Exchange also can use the Microsoft Virtual Shadow Copy Service (VSS).

If you encounter a problem during Data Protection for Exchange processing using VSS for backup and restore, perform the following steps as your first attempt to resolve the problem:

1. Retry the operation that failed.
2. If the problem still exists, close other applications, especially those applications that interact with Exchange (antivirus applications, for example). Retry the operation that failed.
3. If the problem still exists:
 - a. Shut down the Exchange server.
 - b. Restart the Exchange server.
 - c. Run the operation that failed.
4. If the problem still exists:
 - a. Shut down the entire computer.

- b. Restart the computer.
 - c. Run the operation that failed.
5. If the problem still exists, determine if it is occurring on other Exchange servers.

Determining if the problem is a Data Protection for Exchange issue or a general VSS issue

The Data Protection client interacts closely with the backup-archive client (DSMAGENT), which performs all of the Virtual Shadow Copy Service (VSS) operations. Determine first if the problem is with the Microsoft VSS service or with the Tivoli Storage Manager.

Perform the following steps to try to isolate the source of the error:

1. Test the connectivity between the Data Protection client and the Tivoli Storage Manager dsmagent. Issue the TDPEXCC QUERY EXCHANGE command on the computer where the Exchange server is installed to verify that your installation and configuration is correct. This command returns information on the following items:
 - Exchange Server status
 - Storage groups
 - Circular logging
 - VSS components

The following example shows the output generated by the TDPEXCC QUERY EXCHANGE command:

Volume Shadow Copy Service (VSS) Information

```
-----
Writer Name           : Microsoft Exchange Writer
Local DSMAGENT Node   : SERVERA
Remote DSMAGENT Node  : SERVERX
Writer Status         : Online
Selectable Components : 4
```

If the TDPEXCC QUERY EXCHANGE command does not return all of this information, you might have a proxy configuration problem. Refer to your User's manual for instructions on setting up proxy configurations. If all of the information returned to you seems correct, proceed to the next step.

2. Use the vssadmin or vshadow utility to recreate the VSS operation without the Tivoli Storage Manager intervening. When VSS operations are failing, use these programs to recreate the error to determine if this is a general VSS problem or a problem within the Tivoli Storage Manager code.

vssadmin

A utility that is installed with your operating system. It can display current volume shadow copy backups and all installed shadow copy writers and providers in the command window. The following commands are examples of possible VSSADMIN commands:

```
VSSADMIN LIST WRITERS
VSSADMIN LIST PROVIDERS
VSSADMIN LIST SHADOWS
```

vshadow

A utility included with the Microsoft Volume Shadow Copy Services SDK that can be used to exercise most of the VSS infrastructure, such as creating/querying/deleting shadow copies. You can also use

vshadow to create both persistent and nonpersistent shadow copies, transportable snapshots, as well as assign a drive letter or mount point to a shadow copy.

- The following items can be determined by using the vssadmin or vshadow utility:
 - Verify VSS Provider configurations
 - Rule out any possible VSS problems before running the Tivoli Storage Manager VSS functions
 - That you might have a VSS configuration problem or a real hardware problem if an operation does not work with vshadow/vssadmin
 - That you might have a Tivoli Storage Manager problem if an operation works with vshadow/vssadmin but not with the Tivoli Storage Manager

- Perform the following tests to ensure that VSS is working correctly:

Test nonpersistent shadow copy creation and deletion

- a. Run “VSHADOW *k*: *l*:” where *k*: and *l*: are the Exchange Server database and log volumes.
- b. Repeat the step above four times.
- c. Inspect the Windows Event Log to ensure that things look appropriate.

Test persistent shadow copy creation and deletion

- a. Run “VSHADOW -p *k*: *l*:” (where *k*: and *l*: are the Exchange Server database and log volumes. You might have to run “VSHADOW -da” to remove this if you do not have enough space.
- b. Repeat the previous step four times.
- c. Inspect the Windows Event Log to ensure that things look appropriate.

Test nonpersistent transportable shadow copy creation and deletion (VSS Hardware Provider environments only)

- a. Run “VSHADOW -p -t=export.xml *k*: *l*:” where *k*: and *l*: are the Exchange Server database and log volumes.
- b. Copy the resultant “export.xml” file from computer 1 to computer 2 before performing the next step.
- c. On the computer you have set aside for offload, run “VSHADOW -i=export.xml”
- d. Inspect the Windows Event Log to ensure that things look appropriate.

If any of these tests fail repeatedly, there is hardware configuration problem or a real VSS Problem. Consult your hardware documentation for known problems or search Microsoft Knowledge Database for any information.

If all tests pass, continue to Step 3.

3. Recreate your specific problem by using vshadow. If you can only recreate your problem through a series of steps (for example: a backup fails only when you perform two consecutive local backups), try to perform those same tests by using vshadow.
 - Exchange VSS backups to Local are simulated by running a vshadow persistent snapshot.
 - Exchange VSS backups to the Tivoli Storage Manager are simulated by running a vshadow nonpersistent snapshot.
 - Exchange VSS backups to Local and to the Tivoli Storage Manager are simulated by running a vshadow persistent snapshot.
 - Offloaded Exchange VSS backups to the Tivoli Storage Manager are simulated by running a vshadow nonpersistent, transportable snapshot.

Refer to the VSHADOW documentation for the specific commands for performing backups.

If you can recreate the problem, it most likely is a general VSS issue. Refer to Microsoft Knowledge Database for information. If your operation passes successfully with vshadow, it most likely is a Tivoli Storage Manager/Data Protection for Exchange client problem.

Tracing the Data Protection client when using VSS technology

You must gather traces for Data Protection for Exchange, the Tivoli Storage Manager application programming interface (API), and the DSMAGENT processes to ensure proper diagnosis of the Virtual Shadow Copy Service (VSS) operation.

The following traces are the different traces to gather when you diagnose Data Protection for Exchange VSS operational problems:

- **Data Protection for Exchange trace**

To create the trace flag, issue the “/TRACEFILE” and “/TRACEFLAGS” command-line options with the following example command:

```
TDPEXCC BACKUP SG1 FULL /TRACEFILE=DPTRACE.TXT /TRACEFLAG=SERVICE
TDPEXC /TRACEFILE=DPTRACE.TXT /TRACEFLAG=SERVICE
```

- **Tivoli Storage Manager API trace**

Enable tracing with the DP/Exchange DSM.OPT file and the “TRACEFILE” and “TRACEFLAGS” keywords. The following text is an example of the entry in the DP/Exchange DSM.OPT file:

```
TRACEFILE APITRACE.TXT
TRACEFLAG SERVICE
```

- **DSMAGENT trace**

Enable tracing with the DSMAGENT DSM.OPT file and the “TRACEFILE” and “TRACEFLAGS” keywords. The following text is an example of the entry in the DSMAGENT DSM.OPT file:

```
TRACEFILE AGTTRACE.TXT
TRACEFLAG ALL_VSS
```

The trace flag, in this instance, is ALL_VSS (you might need different traceflags, depending on the circumstance).

Gathering Exchange with VSS information before calling IBM

The Data Protection client is dependent upon the operating system and the Exchange application. Collecting all the necessary information about the environment can significantly assist in determining the problem.

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

- The exact level of the Windows operating system, including all service packs and hotfixes that were applied.
- The exact level of the Exchange Server, including all service packs and hotfixes that were applied.
- The exact level of Data Protection for Exchange with Virtual Shadow Copy Service (VSS) Backup/Restore support.
- The exact level of the Tivoli Storage Manager API.
- The exact level of the Tivoli Storage Manager server.
- The exact level of the Tivoli Storage Manager backup-archive client.

- The exact level of the Tivoli Storage Manager storage agent (if LAN-free environment).
- The Tivoli Storage Manager server platform and operating system level.
- The output from the Tivoli Storage Manager server QUERY SYSTEM command.
- The output from the Data Protection for Exchange TDPEXCC QUERY EXCHANGE command.
- The device type (and connectivity path) of the Exchange databases and logs.
- The specific hardware that is being used (for example: HBA, driver levels, microcode levels, SAN Volume Controller levels, DS6000™, DS8000® hardware details).
- Permissions and the name of the user ID being used to run backup and restore operations.
- The name and version of antivirus software.
- The IBM VSS hardware provider level.
- The IBM CIM agent level for DS6000 or DS8000 or the SAN Volume Controller.
- A list of vendor-acquired Exchange applications running on the system.
- A list of other applications running on the system.
- A list of the steps needed to recreate the problem (if the problem can be recreated).
- If the problem can not be recreated, list the steps that caused the problem.
- Is Data Protection for Exchange running in a Microsoft Cluster Server (MSCS) or Microsoft Failover Cluster environment?
- Is the problem occurring on other Exchange servers?

Some operating system information that is listed above can be gathered by using the MPS Reports Tool. This is a Microsoft Reporting tool which lets you gather levels of operating system, software installed, driver versions, and other data. You can download the MPS Reports Tool at Microsoft.com.

Gathering Exchange with VSS files before calling IBM

Several log files and other data can be collected for Data Protection for Exchange server diagnosis.

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

- The Data Protection for Exchange configuration file. The default configuration file is `tdpexc.cfg`.
- The Data Protection for Exchange Tivoli Storage Manager API options file. The default options file is `dsm.opt`.
- The Tivoli Storage Manager registry hive export.
- The Exchange Server registry hive export.
- The Tivoli Storage Manager Server activity log. The Data Protection client logs information to the server activity log. A Tivoli Storage Manager administrator can view this log for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- If the Data Protection client is configured for LAN-free data movement, also collect the options file for the Tivoli Storage Manager storage agent. The default name for this file is `dsmsta.opt`.
- Any screen captures or command-line output of failures or problems.

Log files can indicate the date and time of a backup, the data that is backed up, and any error messages or completion codes that could help to determine your problem. The following files are the Tivoli Storage Manager log files that you can gather:

- The Data Protection for Exchange log file. The default location of this file is C:\Program Files\Tivoli\TSM\TDPEExchange\tdpexc.log
- The Tivoli Storage Manager API Error log file. The default location of this file is C:\Program Files\Tivoli\TSM\TDPEExchange\dsierror.log
- The DSMAGENT error log file. The default location of this file is C:\Program Files\Tivoli\TSM\baclient\dsmerror.log

The Windows event log receives information from the Exchange Server and many different components involved during a Virtual Shadow Copy Service (VSS) operation. Microsoft provides a tool in the system32 directory on Windows XP and Windows 2003, called eventquery.vbs (it is a VBS script, not a binary executable file) to retrieve event log data. You might use the tool as an alternative method of getting Windows event log data. To invoke the utility, issue the following CSCRIPT command:

```
cscript c:\windows\system32\eventquery.vbs parameters > outfile
```

The utility, by default, produces a tabular listing of all event log records in three sections (one section per event log type). To get all log data in default format, issue the following command:

```
cscript c:\windows\system32\eventquery.vbs >eq.out
```

Specify the type of event log you require by using one of the following /L parameters:

- /L Application
- /L Security
- /L System

The following example generates output only for the application and system event logs:

```
cscript c:\windows\system32\eventquery.vbs /L Application >eq_app.out  
cscript c:\windows\system32\eventquery.vbs /L System >eq_sys.out
```

You can utilize the /V parameter to receive detailed (verbose) output:

```
cscript c:\windows\system32\eventquery.vbs /V >eq.out  
cscript c:\windows\system32\eventquery.vbs /L System /V >eq_sys.out
```

You can utilize the /FO parameter to specify tabular, list, or comma-separated (CSV) output. The following are the different methods of specifying the output:

- /FO TABLE
- /FO LIST
- /FO CSV

The default format is TABLE. The LIST output puts each column of the record on a separate line, similar to how the Tivoli Storage Manager administrator's command-line interface (CLI) displays output when it is too wide for tabular display. The CSV output can be loaded into a spreadsheet or database tool for easier viewing. The following example generates a detailed CSV file of the application log:

```
cscript c:\windows\system32\eventquery.vbs /L Application /FO CSV /V >eq_app.out
```

You can get additional help information for the tool by using the following example:

```
cscript c:\windows\system32\eventquery.vbs /?
```

The following VSS provider log files can also be helpful, if applicable:

- System Provider - (Windows Event Log)
- IBM System Storage™ SAN Volume Controller, DS6000, DS8000 - D:\Program Files\IBM\Hardware Provider for VSS\IBMvss.log
- NetApp - D:\Program Files\SnapDrive*.log

Troubleshooting Data Protection for Exchange VSS and SAN Volume Controller

The troubleshooting tips included here are designed to help you accelerate your problem determination task.

Important: If you plan to perform Data Protection for Exchange Virtual Shadow Copy Service (VSS) Instant Restore operations using SAN Volume Controller or DS6000/DS8000 disk subsystems, you must use non-SSL communications with the CIMOM. By default, the CIMOM is configured to use SSL communications. The Tivoli Storage Manager only supports non-SSL communications with the CIMOM. See the *Data Protection for Exchange User's Guide* for information on enabling non-SSL communications.

The following areas are where you can troubleshoot when you are having VSS and SAN Volume Controller or DS6000/DS8000 problems:

- CIMOM (Common Information Model Object Manager) Connectivity issues

To verify connectivity to the CIMOM, perform the following steps:

1. Refer to your SAN Volume Controller documentation.
2. Run the IBMVCFG LIST command. The default location is D:\Program Files\IBM\Hardware Provider for VSS-VDS.
3. Issue the IBMVCFG SHOWCFG command to view the provider configuration information.
4. Check that the CIMOM is properly configured. Run `verifyconfig.bat -u username -p password` on the SAN Volume Controller Master Console.
5. Check the username and password. If there is a problem with the truststore, follow the procedure in the documentation to generate a new truststore.
6. Set up the CIMOM properties file in non-SSL mode if you are using the SAN Volume Controller or DS6000/DS8000 and you plan to use Instant Restore.

- CIMOM operational issues

If your backup or restore fails, check the IBMVSS.log file. If the failure is due to a CIMOM failure, the log displays output similar to the following example:

```
Wed Jan 11 17:34:34.793 - Calling AttachReplicas
Wed Jan 11 17:34:35.702 - AttachReplicas: 909ms
Wed Jan 11 17:34:35.702 - returnValue: 34561
Wed Jan 11 17:34:35.718 - AttachReplicas returned: 34561
java.util.MissingResourceException: Can't find resource for
bundle java.util.PropertyResourceBundle, key 1793
at java.util.ResourceBundle.getObject(ResourceBundle.java:329)
at java.util.ResourceBundle.getString(ResourceBundle.java:289)
at com.ibm.cim.CIMException.<init>(CIMException.java:472)
at ESSService.executeFlashCopy(ESSService.java:3168)
Wed Jan 11 17:34:35.779 - IBMVSS: AbortSnapshots
```

A return value of 0 means that it was successful. To determine why it failed, look at the log files generated by the CLI or graphical user interface (GUI), depending on how you run your operation. These might provide more information on the failure.

- Host configuration issues

If the failure seems to be for a different reason than a CIMOM failure, verify your configuration. Run the latest support levels of the software for the SAN Volume Controller or DS6000/DS8000. Check the IBM Storage web site for details.

- Collecting logs in this environment

If you are unable to resolve these problems, provide the following information to IBM Support:

- Information listed in the Tivoli Storage Manager diagnostic information section
- HBA type, firmware and driver levels
- SDD version
- SAN Volume Controller microcode version (if applicable)
- DS6000/DS8000 microcode version (if applicable)
- SAN Volume Controller Master Console version (if applicable)
- For DS6000/DS8000, the CIM Agent version (if applicable)
- IBMVSS.log
- Application Event Log
- System Event Log

If the problem appears related to CIMOM, you also need the CIMOM logs. Run `CollectLogs.bat` and send the file that is created (`CollectedLogs.zip`) to IBM Support. The default location for the SAN Volume Controller is `C:\Program Files\IBM\svconconsole\support`, and the default location for the DS6000/DS8000 is `C:\Program Files\IBM\cimagent`.

Chapter 6. Resolving storage agent problems

The storage agent is a program that enables Tivoli Storage Manager to back up and restore client data directly to and from SAN-attached storage.

Checking the server activity log for storage agent information

Check the server activity log for other messages, 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 for messages from the storage agent. To review the activity log messages, issue the QUERY ACTLOG command. See the *Tivoli Storage Manager Administrator's Reference* for more details.

If no messages are seen for this storage agent in the server activity log, 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.

If messages are seen on the server for this storage agent, do any error messages give an indication about why this storage agent is not working with the server?

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.

The system error log for AIX is errpt, and for Windows it is the Event Log.

If a device or volume that is used by Tivoli Storage Manager is reporting an 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.

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 with 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.

Chapter 7. 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 “ANE497II 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 using 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 use 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, use `DEFINE PATH` 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 as 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 the server which 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 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 will send 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 store 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 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 8. 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 classes are a way to enable many related trace classes by specifying the aggregate trace class name.

The trace classes documented in table Table 13 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 13. Trace classes for the Administration Center

Trace class	Description	Uses
ADMCNTR	Displays processing information for the portlet and servlet classes. An aggregate 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 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. An aggregate 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 13. 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 FastBackserver 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 13. 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. An aggregate 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 85

Enabling Administration Center trace

Tracing is available for the Tivoli Storage Manager Administration Center.

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 Web site: <http://www.ibm.com/support/docview.wss?uid=swg21410467>.

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 under OPT\IBM\AC\products\tsm\bin. Start the utility by typing supportUtil.bat on Windows or supportUtil.sh on AIX, HP-UX, Linux, or Sun 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:\OPT\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
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 mange tracing menu
99. Return to main menu

Enter Selection:
```

3. Perform the operation that is causing the problem to create a trace of the problem, that is then written to the trace file. The trace file is located in the Tivoli Storage Manager installation directory under C:\IBM\AC\profiles\TIPProfile\logs\server1\trace.log.
4. Disable tracing. Leaving tracing “enabled” might cause performance problems.
5. Run Collect trace files, logs, and system information to send to support. This procedure packages up the trace information collected and log files into a single file called AdminCenterSupport.zip. This file can then be sent to support.

Related reference

“Administration Center Support utility” on page 85

“Trace classes for the Administration Center” on page 125

Enabling a trace for the server or storage agent

Trace commands can be issued from 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 are two trace classes that are always enabled by default, whether they appear on the options file or not. These are ADDMSG and SYSTIME. 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 were 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 has begun.
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 will be 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 was being done to a file, ending the trace will write any remaining trace messages to the file and then close 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 LOGTRACE BEGIN MYTRACE.OUT
```

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, depending on the frequency of the failure. You might not be able to view the activity log, 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.

Restriction: *<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 wanted 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 are a shortcut for enabling many related trace classes by simply specifying the aggregate trace class name for the TRACE ENABLE command.

The trace classes listed in Table 14 are those 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 should be used with the TRACE ENABLE and TRACE DISABLE commands.

Table 14. Server or storage agent trace classes

Trace Classes	Description	Uses
ADDMSG	Issues console messages (ANR, ANE, and so on) 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	Displays information about user data stored on sequential media devices. This is an aggregate trace class that enables 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.
AFCREATE	Displays information about storing user data on sequential media volumes.	Use this trace class to diagnose writing user data on sequential media volumes.

Table 14. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
AFMOVE	Displays operations that move user data using 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	Displays information volume selection and assignment, coordination of drives (mount points), and management of data placement on volumes. This is an aggregate trace class that enables ASALLOC, ASRTRV, ASDEALLOC, ASMOUNT, ASVOL, ASTXN, and ASSD. Typically recommended 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	Displays information about reserving and allocating space on sequential media volumes for the purpose of storing data. This might be for storing data on behalf of a client session or for server data movement operations such as MIGRATION, RECLAMATION, MOVE DATA, or MOVE NODEDATA	Use this trace class to diagnose problems where the server or storage agent report no space available but there should be space available in the storage hierarchy.
ASDEALLOC	Displays information about releasing and de-allocating space on sequential media volumes for the purpose of 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 14. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
ASMOUNT	Displays information about drive (mount point) selection and assignment for sequential media devices.	Use this trace class to 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 preempted.
ASRTRV	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	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	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 preempted.
ASSD	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.

Table 14. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
BF	Information about user data (files) stored in the storage hierarchy. This is an aggregate trace class that enables BFCREATE, BFRTRV, BFSALVAGE, BFLOCK, BFAGGR, BFREMOTE, and BFTRG.	Use this trace class to diagnose general data read or write problems for client operations and server processes.
BFAGGR	Displays information about server aggregation of user data. The server will aggregate 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	Displays information about client operations that store data in the storage hierarchy. Typically, these 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 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.
BFRTRV	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 from a backup set.

Table 14. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
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.
COLLOCATE	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. This can cause a large amount of output trace statements.	Use this trace class to diagnose problems with collocation processing.
CRC	Displays information about generating and managing CRCs on the server or storage agent. This is an aggregate trace class that enables CRCDATA, CRCPROTO, and CRCVAL.	Use this trace class to diagnose data corruption issues where CRC processing did not report data corruption.
CRCDATA	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	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	Displays information about generating and comparing CRC values.	Informational for displaying CRC values during processing.
CRYPTO	Displays information about AES encryption 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 14. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
DBCONN	Traces connection activities.	Use this trace class to trace 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. Use this trace class first to consider 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 concerns 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	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, this still displays information but was superseded by the LANFREE trace class. Typically it is best to enable 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 this to trace general logic paths for data deduplication processing.
DEDUP1	Traces error paths for data deduplication processing.	Use this to trace error paths for data deduplication processing.
DEDUP2	Traces the fingerprinting and digital signatures path.	Use this to trace fingerprinting and digital signature paths.

Table 14. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
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, including client backup and restore, expiration, deletion (DELETE FILESPACE, DELETE VOLUME), export/import, backup set generation and restore, no-query restore, db audit, etc.
DF	Displays information about user data stored on disk volumes. This is an aggregate trace class that enables DFCREATE, DFRTRV, DFMOVE, DFLOCK, DFTXN, and DFCOPY. It is recommended to 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	Displays information about storing user data on disk volumes.	Use this trace class to diagnose writing user data on disk volumes.
DFMOVE	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	Displays information about reading user data from disk volumes.	Use this trace class to diagnose reading user data from disk volumes.

Table 14. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
DS	Displays information about volume selection, space reservation, assignment, and management of data placement on disk volumes. This is an aggregate trace class that enables DSALLOC DSRTRV DSDEALLOC DSVOL. It is recommended to 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	Displays information about reserving and allocating space on disk volumes for the purpose of storing data. This might be storing data on behalf of a client session or for server data movement operations such as MIGRATION, MOVE DATA, or MOVE NODEDATA.	Use this trace class to diagnose problems where the server or storage agent report that no space is available, but there should be space available in the storage hierarchy.
DSDEALLOC	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	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	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 14. 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 filespace functions.	Use this trace class to debug problems related to inventory file spaces (e.g. during DELETE FILESPACE).
LANFREE	Displays general information about LAN-free operations on either the server or storage agent. Also displays error information for LAN-free-related operations. This is an aggregate trace class that enables LNFVERB, LNFMEM, LNFENTRY, and LNFDATA.	Any LAN-free failure.
MMS	Displays information about tape libraries and the server or storage agent use of these. This is an aggregate trace class that enables MMSBASE, MMSTXN, MMSLIB, MMSDRIVE, MMSOP, MMSMAN, MMSSCSI, MMSFLAG, MMSACSL, and MMSHARE. 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 14. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
NA	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. This is an aggregate trace class that enables NALOCK, NAPATH, NAMOVER, NADISK, and NACONFIG. It might be best to include MMS and PVR trace classes when tracing NA.	Use this trace class to diagnose problems with paths to devices.
PROXYNODE	Displays information about proxynode sessions and the commands related to proxynode associations (GRANT, REVOKE, QUERY PROXYNODE).	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.
PVR	Displays information about sequential media devices and the server or storage agent use of these devices. This is an aggregate trace class that enables PVRVOL, PVRCLASS, PVRMP, and PVRREMOTE. Additionally, this aggregate enables the following classes for open platforms only: PVR8MMBASE, PVR4MMBASE, PVRQICBASE, PVRDLTBASE, PVRECARTBASE, PVRDTFBASE, PVRTAPIBASE, PVRGTS, PVRFILE, PVRTAPE, PVRRODSK, PVRNTP, PVRSHARE, and PVRNAS.	Use this trace class to diagnose problems with tape drives, failures reading or writing tape volumes, or other tape-volume-related issues.

Table 14. 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	Displays information about sessions connected to the server, including all verbs sent and received by the server.	This trace class is useful in many cases. It 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 14. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
SHRED	Displays information related to data-shredding operations on the server. The SHRED command is available in Tivoli Storage Manager, Version 5.4.	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's 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. Available in Tivoli Storage Manager Version 5.5 and later.	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. Available in Tivoli Storage Manager Version 5.5 and later.	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.

Table 14. 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. This 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. This 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 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 14. Server or storage agent trace classes (continued)

Trace Classes	Description	Uses
UNICODE	Displays information about code page conversions and Unicode filesystem operations.	Use this trace class to debug problems related to code page conversion problems or unicode filesystem problems.
XI	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 not all of the available SHOW commands.

Table 15. 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 <i>aggrID-high</i> <i>aggrID-low</i> . <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 15. 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 ASVol.	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 15. 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's 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.
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 a 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.
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.

Table 15. 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's inventory. The syntax is SHOW GROUPLeaders <i>objID-high</i> <i>objID-low</i> . <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's inventory. The syntax is SHOW GROUPMembers <i>objID-high</i> <i>objID-low</i> . <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 15. 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</i> . <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 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's inventory. You might issue this command if you are having trouble restoring, retrieving, expiring, or auditing the object. This command reports the following items: <ul style="list-style-type: none"> • New information for archive retention protected objects. • If the archive object is in deletion hold. • If 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 will return the inventory information for all libraries.	Issue this command if there is a problem with the library inventory information. The command will display 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 will return 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 15. 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, it reports this information.
LOGVARS	Displays recovery log global attributes. The syntax is SHOW LOGVars.	Issue this command to determine the state of the recovery log.

Table 15. Server or storage agent SHOW commands (continued)

SHOW Command	Description	Recommendation
MEMTREND	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.	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.
MP	Displays mount points. The syntax is SHOW MP.	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 complimentary information with this command to display the current state of drives and current devclass mount point counts.
NASDEV	Displays the SCSI devices attached to a NAS file server associated with a NAS datamover definition. The syntax is SHOW NASDev.	Create an 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.

Table 15. Server or storage agent SHOW commands (continued)

SHOW Command	Description	Recommendation
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.

Table 15. 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 will take action and cancel or terminate 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 15. 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 SSPool.	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 15. 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=<i>setNum</i> TOUCH=<i>setNum</i>. 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 Table Of Contents (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 15. 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 that must be 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 DEDUPOBJECT 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 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, Sun Solaris, and Windows” on page 158

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. It provides more detail, but because the memory buffer size is fixed, fewer events are traced. It also 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 circumstances where 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, Sun Solaris, and Windows

AIX

Solaris

Windows

The stand-alone utility, `ddtrace`, exactly mimics the DDTRACE server commands.

A stand-alone utility, `ddtrace`, 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 Tivoli Storage Manager server initialization. It 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 data for the client

You can enable tracing on the client or client 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 IBM 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 classes are a shortcut for enabling many related trace classes by 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 16 on page 161 are those that are most typically requested or used for diagnosing problems. The trace class name should be used with the TRACEFLAG options in DSM.OPT.

Table 16. 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 SERVICEtrace 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 of 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.
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 SERVICEtrace class, except for 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.

Table 16. Trace classes (continued)

Trace Class	Description	Recommendation
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 directory, read or write.
DOMAIN	Displays incremental domain processing information. Part of the SERVICE trace aggregate.	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.

Table 16. Trace classes (continued)

Trace Class	Description	Recommendation
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 problems in accepting processing options from server client-options sets.
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.

Table 16. Trace classes (continued)

Trace Class	Description	Recommendation
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 will generally not be 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.
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 such as transaction aborts and retries.

Table 16. Trace classes (continued)

Trace Class	Description	Recommendation
VERBDETAIL	Displays detailed verb information pertinent to client-server sessions. Part of the ALL_SESS trace aggregates.	Use this trace class to determine 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, more recent method, 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 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 `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:

Windows `tracefile c:\service\trace.out`

AIX **HP-UX** **Linux** **Solaris** `tracefile /home/spike/trace.out`
`tracefile trace.txt (Macintosh)`

5. You can also 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).

6. 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: `tracesegsize 200`

Trace file segments are named using the name 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 will be named using the tracefile option with an additional extension using the segment number. For example, `trace.out.1`

7. 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 -tracesegsize=200
```

Related reference

"Client trace classes" on page 160

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.
 1. **AIX** **HP-UX** **Linux** **Solaris** On AIX, HP-UX, Linux, or Sun 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`.
 2. **Windows** On 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 `dsmtrace` is used to configure tracing, it contacts the client through the named pipe on which the client is listening and passes to it the desired 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 client.
2. Configure the client options file with the desired trace options.
3. Restart the client and reproduce the problem.
4. Stop the client.
5. Remove the trace options from the 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 client operation, or when you need to start tracing after the client was running for some time.

The `dsmtrace` utility includes the following features:

- Identify running processes and their process IDs (PIDs)
- Enable client tracing
- Disable client tracing
- Query client trace status

The following table summarizes the availability of this feature:

Table 17. Availability of the `dsmtrace` utility

Client Component	AIX, HP-UX, Linux, or Sun Solaris Program Name	Windows Program Name
Backup-Archive Client (command line)	<code>dsmc</code>	<code>dsmc.exe</code>
Backup-Archive Client (GUI)	N/A	<code>dsm.exe</code>

Table 17. Availability of the dsmtrace utility (continued)

Client Component	AIX, HP-UX, Linux, or Sun Solaris Program Name	Windows Program Name
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 17 on page 167 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 ID (PID) of the client 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>
```

Linux Important note for Linux users: 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. For example:

```
fvttlinuxppc:/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

```
fvttlinuxppc:/opt/tivoli/tsm/client/ba/bin #
```

In such a situation, use 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.

2. 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

We have gathered the known problems and limitations of trace processes to help you resolve problems 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 dsmttrace to enable tracing for such commands, the result can be a stoppage of the dsmttrace process and (AIX, HP-UX, Linux, and Sun Solaris only) a residual named pipe in the /tmp directory.

- **Windows** You must be logged in with a local administrative account in order to use dsmttrace.
- You must be logged in as root to use dsmttrace. 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 dsmt parent process. For an example, refer to the previous example section.
- When multiple instances of the same program are running, you must identify the PID of the instance 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 desired PID. For example, if you want to trace dsmt that is being run by user 'andy' and there are two instances of dsmt (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 dsmterror.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. No is the default. |
| 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 dsmt.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 will close and the trace flags, maximum trace size, maximum trace segment size, and trace file name will all be cleared.

<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 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. Multiple trace flags are separated 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 will grow indefinitely). When the maximum length is reached, then the trace will wrap back to the beginning of the file. Specify a value in MB between 1 and 4095. If tracing is already running, then this option has no effect.

-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-manageable segments. Each segment will be no larger than the specified size. When this option is used, a segment number will be appended to the trace file name for each

segment. Specify a value in MB between 1 and 1000. If tracing is already running, then this option has no effect.

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 the Tivoli Storage Manager reporting and monitoring agent for AIX and Linux

AIX

Linux

You can create a trace for the Tivoli Storage Manager reporting and monitoring agent on AIX and Linux.

After completing the trace steps, a log file named `xxxyyyzzz.log` is created, where:

`xxx` is the agent instance name

`yyy` is the Tivoli Storage Manager port number

`zzz` is the time stamp

The log file is located in the `/install_dir/itm/bin/` directory:

Complete the following steps to create the trace log file:

1. Stop the Tivoli Storage Manager monitoring agent instance.
 - a. Open the CandleManage program.

- b. Select the Tivoli Storage Manager monitoring agent instance that you want to trace.
- c. Right-click the instance and select **Stop**.

Tip: To use the command-line interface, issue the following command:

```
Itmcmd agent -o instance_name stop sk
```

where *instance_name* is the Tivoli Storage Manager monitoring agent instance that you want to trace.

2. Activate tracing. Edit the instance configuration file by changing KSK_TRACE=0 to KSK_TRACE=1. The configuration file is in the following directory:

```
/opt/Tivoli/TSM/Reporting/itm/config/sk_XXX.config
```

where *xxx* is the name of the agent instance.

3. Start the Tivoli Storage Manager monitoring agent instance.
 - a. Open the CandleManage program.
 - b. Select the Tivoli Storage Manager monitoring agent instance that you want to trace.
 - c. Right-click the instance and select **Start**.

Tip: To use the command-line interface, issue the following command:

```
Itmcmd agent -o instance_name start sk
```

where *instance_name* is the Tivoli Storage Manager monitoring agent instance that you want to trace.

Tracing the Tivoli Storage Manager reporting and monitoring agent for Windows

Windows

You can create a trace for the Tivoli Storage Manager reporting and monitoring agent on a Microsoft Windows server.

After completing the trace steps, a log file named *xxxyyyzzz.log* is created, where:

xxx is the agent instance name

yyy is the Tivoli Storage Manager port number

zzz is the time stamp

The log file is located in the `\install_dir\itm\bin\` directory:

Complete the following steps to create the trace log file:

1. Stop the Tivoli Storage Manager monitoring agent instance.
 - a. Open the Tivoli Enterprise Monitoring Services program.
 - b. Select the Tivoli Storage Manager monitoring agent instance that you want to trace.
 - c. At the top of the window, click **Stop**.
2. Activate Tracing. Edit the instance configuration file by changing KSK_TRACE=0 to KSK_TRACE=1. The configuration file is in the following directory:

| \ibm\itm\tmaitm6\kskenv_xxx.config

| where *xxx* is the name of the agent instance.

- | 3. Start the Tivoli Storage Manager monitoring agent instance.
 - | a. Open the Tivoli Enterprise Monitoring Services program.
 - | b. Select the Tivoli Storage Manager monitoring agent instance that you want to trace.
 - | c. At the top of the window, click **Start**.

Tracing data for the API

You can enable tracing for the application programming interface (API).

To enable tracing for the 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` will fail with return code `DSM_RC_CANNOT_OPEN_TRACEFILE` (426).

To enable tracing for the multithreaded API after an application has been started, use the `dsmtrace` utility. The `dsmtrace` utility allows you to turn on tracing while the problem is occurring, without having trace constantly enabled. Refer to the *dsmtrace* section.

Chapter 9. Resolving problems with the Tivoli Storage Manager reporting and monitoring feature

If problems with the reporting and monitoring feature cannot be resolved using the procedures documented in the *IBM Tivoli Storage Manager Installation Guide*, the following topics are available.

Resolving Warehouse Proxy workspace reporting problems

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.

Tivoli reports are displayed in several languages, but problems have arisen 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:

- Modify the UTF8TEST table in the WAREHOUS database.
- Add the environment variable DB2CODEPAGE=1208.
- Recycle all service applications in IBM Tivoli Monitoring by stopping and starting them.

If you are still experiencing problems reporting 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 WAREHOUSE database, select and open the UTF8TEST table and change the “Results” column on the third row to a 0 (zero).
- Stop and restart the Warehouse Proxy Agent within the Tivoli Enterprise Monitoring Server.

Resolving historical data reporting problems in the Warehouse Proxy workspace

If you are experiencing errors in the Warehouse Proxy workspace in the Tivoli Enterprise Portal browser and you are using a non-English system, you might need to reconfigure 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.

Resolving a reporting and monitoring agent installation hang

Windows

A halted installation of the reporting and monitoring agent could be attributed to the release level of IBM Tivoli Monitoring. A stoppage that persists for more than 30 minutes while “TSM Agent” is displayed in the installation status menu indicates a failed installation.

The reporting and monitoring agent software code is designed for IBM Tivoli Monitoring release 6.2, fix level 1. The IBM Tivoli Monitoring software code must be at the 6.2.1 base code level for Windows 2008 server. Complete the following steps to resolve the installation stoppage:

1. Verify that the halted installation is not attributed to a shortage of memory or network interference.
 - a. Change directory to the *installation_dir*\itm\InstallITM directory, where *installation_dir* represents the base directory where IBM Tivoli Storage Manager is installed.
 - b. Examine the IBM Tivoli Monitoring for Storagexxxxxxxxxxxxx.log log file, where xxxxxxxxxxxx represents the date today. For example, IBM Tivoli Monitoring for Storage20100319 1755.log
 - c. Review the installation log file to determine if the Tivoli Storage Manager Agent installation continues processing. If the following entry occurs in the installation log file, continue to step 2. Otherwise, wait for the installation to complete.

```
Updated BuildPresentation.bat: call InstallPresentation.bat
WIN-ZFDTID07JM0 KHD KIT KIW KSK KSY KTM KUM >> BuildPresentation.txt
2>> BuildPresentationErr.txt
```

The entry occurs at the end of the file.

2. Open the Task Manager and stop the KfwServices process. The installation continues after the KfwServices process terminates.
3. After the installation completes, issue the following command:

```
installation_dir\itm\CNPS\BuildPresentation.bat
```

Chapter 10. Help facilities

IBM Tivoli Storage Manager has several places where you can find solutions to any problems you might have with the server or 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) to be 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 have been displayed to fill the display space, or when the end of the requested help text is displayed, you see a prompt along with instructions for what can be entered at that prompt. To continue displaying text for your current selection, press enter or type 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 the prompt 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 will display more comprehensive information within the Windows help utility.

Server or storage agent help

The server and storage agent both provide a help facility, which provides descriptions and syntax for server commands, as well as a full description of server messages which includes **Explanation**, **System Action**, and **User Response**.

Accessing server or storage agent help for commands

You must issue the help command to access help for the server or storage agent.

Issue the following command for help on a server command: `HELP commandName` where *commandName* is the server command for which you want information. For example, the `HELP REGISTER NODE` command displays the description of the `REGISTER NODE` command as well as the syntax and information about the command parameters.

Accessing help for messages

You must 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 will report 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 as well as a full description of client messages, including **Explanation**, **System Action**, and **User Response**.

Help information for the GUI and Web GUI clients is available through the **Help** menu item.

Reporting a problem with a help topic

You must collect certain information before you report a problem with the help system.

1. Record what you clicked on 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 will show you a menu with a **View Source** option. Select **View Source** and a window is displayed with the HTML source code for that window. Write down the title of that window, which will be the URL or the name of the file that the help system is attempting to display.

Chapter 11. Determining data storage problems

Data storage issues can be resolved through several methods.

Using data storage diagnostic tips

If you are experiencing a problem in storing or retrieving data, review the diagnostic tips to try to isolate or resolve the problem.

Checking the server activity log to resolve data storage issues

Check the server activity log for other messages 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 beyond just the message itself. The **Explanation**, **System Action**, or **User Response** sections of the message might provide additional information about the problem. 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 relating 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 restore of another client node's data. 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 the problem is an error reading or writing data from a device, many systems and devices record information in a system error log. Examples of the system error log the are errpt for AIX and the Event Log for Windows..

If a device or volume used by IBM Tivoli Storage Manager is reporting an 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.

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. These storage pool definitions are also used by the storage agent. If attributes of a storage pool were changed, this change might affect data store and retrieve operations.

Review any changes to the storage hierarchy and storage pool definitions. Use QUERY ACTLOG 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 of these 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 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 using 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 will fail.

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's 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 and retrying 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 Tivoli Storage Manager problem.

Device driver hints and tips

There are many possible causes for device driver problems. The problem might be with 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.

Note: 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 will use a host bus adapter (HBA) to communicate. If the device is SCSI attached, the device driver will use 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 that resolves the issue.
- 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 been 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 examples are 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 which have the same bit mode with running kernel modules. In other words, the 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 HBA 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 32-bit HBA API on 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 multiple LUN enabled 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)

Restriction: DDTRACE GET and DDTRACE END are not required.

The 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 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 about how to load drivers to the kernel.

Related reference

“Device driver hints and tips” on page 187

“Adjusting to changes in the HBA or SCSI adapter connecting to the device” on page 188

“Resolving a loose cable connection” on page 189

Updating device information of host systems on a dynamic SAN without restarting

When devices in a storage area network (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 On AIX, issue the `cfgmgr` command to force the operating system to re-configure itself. Then run SMIT to re-configure your Tivoli Storage Manager devices.

HP-UX On HP, 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 On 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 18. HBA adapters and corresponding drivers for all architectures of Linux

HBA Adapters	HBA Driver Name	Available Architectures
Adaptec 7892	<code>aix7xxx</code>	IA32, AMD64
Qlogic 22xx	<code>qla2200</code>	IA32, AMD64
Qlogic 23xx	<code>qla2300</code>	IA32, AMD64
Qlogic 2362	<code>qla2362</code>	EM64T
Emulex	<code>lpfcdd</code>	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 execution 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 that connects to a computer through a SAN 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 due to a difference between Tivoli Storage Manager's 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. Tivoli Storage Manager requires the following conditions for hard disk drives, disk subsystems, vendor-acquired file systems, and remote file systems:

Bypassing cache during write operations

When IBM Tivoli Storage Manager opens database, recovery log, and storage pool volumes, they 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, Tivoli Storage Manager can maintain the integrity of client attributes and data. Bypassing the cache is required because 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, or data that was supposed to be written to the storage pool volumes might be missing.

Hard disk drives installed in 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, or else they 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, depending 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, Solaris, or HP-UX, unless the DSMSEV_DIR environment variable is set. For more information about setting the environment variable, refer to the *IBM Tivoli Storage Manager Installation Guide*.

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 store 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 19 for the file name extension created by the server for scratch volumes, depending on the type of data that is stored.

Table 19. 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 will be sharing 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

There are several possible causes for problems with tape drives and libraries. The problem might be with software on the computer trying to use the device, the connections to the device, or the device itself.

Whenever a device problem is encountered, ask, "Has anything been changed?" Suspect anything 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.

- 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, this situation often accounts for intermittent or persistent failures. Check any cabling changes to verify that they are correct.
- If the device firmware has changed, this situation 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.

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 this fix level, release, or version of the operating system.

Adjusting to device driver changes

A device driver upgrade can 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 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 SCSI connection to the device uses a SCSI adapter, while a fibre-channel (optical) connection to the device uses a host bus adapter (HBA). In either case, the cause of the problem might be if an actual adapter was changed or the computer was opened and other hardware changed or fixed.

Remember: The connecting point for the device to the computer is usually referred to 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 back to the previous adapter to see if that resolves the issue.
- 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 been loosened, which 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 difficult 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.

Storage area network hints and tips

There are many possible causes or problems with a storage area network (SAN). SAN problems might be with software on the computer trying to use the device, connections to the device, or the device itself.

Whenever a SAN problem is encountered, ask “Has anything been changed?” Any kind of changes, from the computer trying to use the device to the device itself might be suspect, especially if the device worked prior to a given change, then stopped working after that change.

To better understand the following discussion on diagnosing problems with a SAN, review the following terminology and typical abbreviations that are used:

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 storage area network. An HBA is similar in function to a network adapter in how it provides access for a computer to a local area network or wide area network.

Storage area network

A storage area network (SAN) is a network of shared devices that can typically be accessed using fibre. Often, a storage area network is used to share devices between many different computers.

Knowing your SAN configuration

Understanding the storage area network (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 storage area network (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 storage area network (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 an 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 will typically provide a device driver for use with their HBA. Similarly, the vendor will provide 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, should be updated to pick up the latest fixes and support.

Configuration

HBAs typically have a number of configurable settings. The settings typically affect how Tivoli Storage Manager functions with a SAN device.

Related reference

“HBA configuration issues”

HBA configuration issues

Host bus adapters (HBAs) typically have many different configuration settings and options.

The HBA vendor usually provides information about the settings for your HBA and the appropriate values for these settings. Similarly, the HBA vendor should provide a utility and other instructions on how to configure your HBA. The following settings are those that typically affect using 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 typically supports many different configurations. The ports on the switch need to be configured appropriately for the type of storage area network (SAN) that is set up, as well as 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 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 storage area network (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 storage area network (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 should provide 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 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 storage area network (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 will typically report errors. The following errors are experienced often:

CRC error, 8b/10b code error, and other similar symptoms

These are recoverable errors 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 types 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 that a faulty gigabit interface connector (GBIC), media interface adapter (MIA), or cable. The recovery for this type of error is disruptive. This error will appear 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.

Note: Often these errors are 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

You might be experiencing problems with your storage agent SAN devices.

See Table 20 for errors that are generated for SAN devices.

Table 20. Common SAN device errors

Error	Explanation
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>For SAN device errors, this message is often issued. 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>
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>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>
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>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 20. 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 QUERY PATH F=D 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"> - UPDate DRive <<i>libraryName</i>> <<i>driveName</i>> SERIAL=AUTODetect <p>This command will force-record the new serial number into the server database. Because the drive was replaced, the element number stays the same.</p> - UPDate PATH <<i>sourceName</i>> <<i>driveName</i>> SRCT=SERVER DESTT=DRIVE LIBRARY=<<i>libraryName</i>> DEVIce=xxxxx AUTODetect=Yes <p>This command will force-record the new serial number into the database. Because the drive was replaced, the element number stays the same.</p> – If the drive or drive path was 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 can be 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: UPDATE PATH AUTODetect=Yes</p>
<p>ANR8972E: Unable to find element number for drive <i>driveName</i> in library <i>libraryName</i></p>	<p>If the ELEMent parameter was set to AUTODetect when defining the drive, Tivoli Storage Manager tries to get the drive's element number automatically. 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 205

SAN device mapping hints and tips

Storage area network (SAN) device discovery and device mapping are supported on Windows 2000, Windows 2003 (32 bits), AIX, Solaris, and Linux (except Linux zSeries®).

The following items illustrate the advantages of 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 will help 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 via HBA. Instead, the server uses the SNIA 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 performed to obtain 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 Solaris: libhbaapi32.so
- 64-bit Solaris: libhbaapi64.so

If any of these files are missing, message "ANR1791W HBAAPI wrapper library xxxxxxxx failed to load or is missing." is displayed.

Related concepts

"SAN device mapping errors" on page 205

Disabling SAN device mapping:

Occasionally you must disable SAN device mapping to circumvent a problem or to 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, Sun Solaris, and Windows platforms. SANDISCOVERY OFF is the default for all other platforms.

Platform-specific information:

When working on your SAN device mapping, it is important that you know your platform-specific information.

AIX 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 has provided 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.

Note: For a list of supported HBAs and required driver levels by operating system, go to IBM Support and search on keyword TSMHBA.

SAN device mapping errors

The errors that are most often generated during storage area network (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 HBAAPIIn 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 (e.g. 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. For IBM tape devices, make sure the latest firmware is installed. Firmware prior to 4772 for IBM 3580 tape devices caused problems with Qlogic HBAAPI.

ANR1790W SAN discovery failed

Error message ANR1790W is a general message indicating that the HBAAPI function failed and we cannot perform SAN discovery.

Verify that all SAN devices are working properly (e.g. 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.

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 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** **Solaris** AIX, HP-UX, Linux, or Sun Solaris systems (except on Linux zSeries) store their HBAAPI libraries in the location specified by the /etc/hba.conf file, and Windows file 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 Solaris file)
- /opt/JNIsnia/Solaris/Jni/64bit/JniHbaLib.so (JNI's 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 once.

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 above setopt SANDISCOVERY command can be issued as many times as necessary.

Another way to disable/enable SAN discovery is to put the following option in the `dmserv.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, Sun Solaris, and Windows platforms. SANDISCOVERY is defaulted to OFF for all other 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 list displays a number of 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 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 was encountered 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.

ANR8227E Fileset `devices.common.IBM.fc.hba-api` is not at the required level

Error message ANR8227E is only displayed for AIX.

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 204

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 storage area network (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 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`

Note: Alternatively, reinstalling 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 having 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 `tsmscsi` 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`

Note: Alternatively, reinstalling 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/tsmscsi/lbinfo` and `/dev/tsmscsi/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 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 OS level.

Perform the following steps to resolve your configuration issues:

1. 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 to make sure the HBA driver level and OS 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

During filer-to-server operations, IBM Tivoli Storage Manager uses up to two additional TCP/IP ports; a control port used internally by Tivoli Storage Manager during both backup and restore operations, and a data port used during network data management protocol (NDMP) backup operations to a Tivoli Storage Manager native storage pool.

Tivoli Storage Manager defaults to the standard 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.

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 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 were issued, the data that Tivoli Storage Manager reports from these devices might provide sufficient detail 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. Accessibility features for Tivoli Storage Manager

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Accessibility features

The following list includes the major accessibility features in Tivoli Storage Manager:

- 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
- User documentation provided in HTML and PDF format. Descriptive text is provided for all documentation images.

The Tivoli Storage Manager Information Center, and its related publications, are accessibility-enabled.

Keyboard navigation

Windows The Tivoli Storage Manager for Windows Console follows Microsoft conventions for all keyboard navigation and access. Drag and Drop support is managed using the Microsoft Windows Accessibility option known as MouseKeys. For more information about MouseKeys and other Windows accessibility options, please refer to the Windows Online Help (keyword: MouseKeys).

AIX Tivoli Storage Manager follows AIX operating system conventions for keyboard navigation and access.

HP-UX Tivoli Storage Manager follows HP-UX operating-system conventions for keyboard navigation and access.

Linux Tivoli Storage Manager follows Linux operating-system conventions for keyboard navigation and access.

Tivoli Storage Manager follows Macintosh operating-system conventions for keyboard navigation and access.

Solaris Tivoli Storage Manager follows Sun Solaris operating-system conventions for keyboard navigation and access.

Vendor software

Tivoli Storage Manager 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.

Related accessibility information

You can view the publications for Tivoli Storage Manager in Adobe® Portable Document Format (PDF) using the Adobe Acrobat Reader. You can access these or any of the other documentation PDFs at the IBM Publications Center at <http://www.ibm.com/shop/publications/order/>.

IBM and accessibility

| See the IBM Human Ability and Accessibility Center for more information about
| the commitment that IBM has to accessibility: <http://www.ibm.com/able>.

Appendix B. Using gt script

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 C. 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 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 21. Supported operating system environments and components

Operating System Environment	Tivoli Storage Manager Component
AIX AIX	Server Storage Agent Device Driver
HP-UX HP-UX	Server Storage Agent Device Driver
Linux Linux	Server Storage Agent Device Driver
Solaris Sun Solaris	Server Storage Agent Device Driver
Windows Windows	Server Storage Agent Device Driver Management Console

Note: 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 Sun 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 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, SUN, 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 Sun Solaris only)
- MGMTCONSOLE (Windows only)
- DEVDRIVER
- SYSTEM
- CLIENAPI (AIX, HP-UX, Linux, or Sun 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 D. IBM Global Security Kit return codes

Version 6.2 of the IBM Tivoli Storage Manager server uses the IBM Global Security Kit (GSKit) Version 7.0.4.27 for Secure Sockets Layer (SSL). The V6.2 backup-archive client uses GSKit Version 8.0.13.3 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\gsk7\  
notepad gskitinfo.txt
```

CAUTION:

You can damage the system registry if you use `regedit` incorrectly.

- For the 64-bit AIX server, issue the following from the command line:
`gsk7ver_64`

See Table 22 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 23 on page 227 for the key management return codes.

Table 22. 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 Secure Sockets Layer (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 22. 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, so 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 22. 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 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.
0x000001a2	418	GSK_ERROR_NO_READ_FUNCTION	The <code>read()</code> failed. Report this error to IBM Software Support.

Table 22. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
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	When validating a certificate, unable to access the specified user registry.
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 SSLv2 formatted 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's 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.
0x0000025a	602	GSK_MISC_INVALID_ID	A valid ID was not specified for the gsk_secure_soc_misc() function call.

Table 22. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x000002bd	701	GSK_ATTRIBUTE_INVALID_ID	The function call does not have a valid ID. 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.

Table 23. IBM Global Security Kit key management 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 Secure Sockets Layer (SSL) handle is not valid. The specified handle was not the result of a successful open() 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.

Table 23. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
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 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 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.
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.

Table 23. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
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 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 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 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	An SSL protocol that is not valid received from partner.

Table 23. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
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	When validating a certificate, unable to access the specified user registry.
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 SSLv2 formatted 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's 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.

Table 23. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
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.

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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



Java[™]

COMPATIBLE Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Glossary

A glossary is available with terms and definitions for the IBM Tivoli Storage Manager server and related products.

The glossary is located in the Tivoli Storage Manager information center: <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r2>

Index

A

- accessibility features 215
- administration center
 - configuring the IP address 80
 - create a user 76
 - create a user ID 76
 - establishing a connection with the IBM Tivoli Storage Manager server 73
 - messages versus IBM Tivoli Storage Manager messages 89
 - support utility 85
 - task fails
 - check the server activity log 87
 - internal errors 88
 - portlets 87
 - wizard 87
 - task fails with message 86
 - trace 127
 - unable to access the server from a Web browser 81
- AIX JFS2
 - image backup 20
 - snapshot-based backup-archive 20
- allocating additional memory 42
- ANR1221E
 - error message 60
- ANR2317W
 - error message 60
- API
 - option file 23
- application programming interface (API) instrumentation 21
 - tracing 177
- auto-update 93, 94

B

- BACKINT
 - troubleshooting 106
- backup application
 - files automatically excluded 10
 - files excluded by EXCLUDE DIR 12
 - files excluded by include/exclude statements 9
 - files excluded due to incremental copy frequency 43
 - include/exclude due to compression, encryption, and subfile backup statements 13
 - include/exclude statements coded wrong 14
 - platform-specific include/exclude statements 13
- BACKUP DB
 - ANR2971E with SQL code 52
 - common errors 53
 - incorrect environment variables 51
- backup-archive client
 - automatic deployment 95

- backup-archive client (*continued*)
 - help 181
 - restarting 95
 - SHOW commands 34

C

- cache
 - bypass during write operations 194
- certificate authority 70
- client
 - authentication failure 6
 - can problem be reproduced 2
 - error messages
 - examining 1
 - generating errors
 - connected to the server 69
 - identifying when and where problems occur 1
 - image backup 17
 - resolving problems 1
 - scheduler 7
 - server activity log
 - examining 1
 - trace classes 160, 165
- client option sets
 - resolving problems 4
 - using 5
- client schedule log 7
- collector tool 77
- communication errors
 - resolving 69
- compressed data during backup-archive 174
- copy frequency 43
- CTGTRV009E
 - error message 64
- CTGTRV011E
 - error message 64
- customer support
 - contact xi

D

- Daemon traceflags
 - client and journal 160
- data
 - sent to the IBM Tivoli Storage Manager storage agent or server 24
- data protection 97
 - Exchange with VSS backup-restore support
 - gathering information before calling IBM 116
 - Exchange with VSS backup/restore support
 - determining the issue 114
 - gathering files before calling IBM 117
 - general help 113

- data protection (*continued*)
 - Exchange with VSS backup/restore support (*continued*)
 - tracing when using VSS 116
 - troubleshooting 119
- data storage
 - diagnostic tips 185
- data storage hints and tips
 - backup or copy problem with specific node 187
 - change the server policies 186
 - change the storage hierarchy 186
 - HELP 185
 - reading or writing to a device 186
 - recreate the problem 185
 - server activity log 185
 - specific volume 187
- database error messages 50
- database ID file missing or incorrect 49
- database restore errors 49
- DB2 password
 - expired 54
- debug print output
 - obtaining 29
- delete
 - troubleshooting 111
- deployment 93, 94
- device driver
 - 32-bit Linux kernel modules 189
 - 64-bit Linux kernel modules 189
 - Adaptec SCSI requirements 192
 - error messages in the system error log 189
 - HBA changes 188
 - HBA drivers on the Linux 2.6.x kernels 190
 - Linux server running on x86_64 architecture 189
 - loose cable connection 189
 - multiple LUN support on Linux kernels 190
 - operating system changes 188
 - performing ddtrace from version 5.3.2 on Linux 191
 - Qlogic fibre-channel HBA BIOS requirements 192
 - SCSI adapter changes 188
 - updating device information 192
- device driver trace
 - from a command shell - AIX, Sun Solaris, Windows 158
 - from the server console/admin client 157
- diagnostic tips
 - client 1
 - data storage 185
 - storage agent 121
- documentation
 - to resolve client problems 2
- dsmc/dsmadm/dsmj/
 - no start 3

E

- education
 - see Tivoli technical training ix
- encrypted data during backup-archive 174
- encrypted file system 17
- error message
 - CTGTRV009E 64
 - CTGTRV011E 64
- error messages
 - ANR1330E 61
 - ANR1331E 61
- establishing a connection with the IBM Tivoli Storage Manager server 73

F

- FILE directory mapping 194
- File Manager
 - troubleshooting 106
- fixes, obtaining xi

G

- GSKit
 - return codes 223
- gt script 217

H

- halt
 - installation 180
- health monitor 82
 - how it works 83
 - re-synchronizing the ADMIN_CENTER administrator ID password 85
 - warning or critical database status 84
 - warning or critical storage status 83
- help
 - server or storage agent 182
- help facilities 181
- help system
 - CLI for server or storage agent 183
 - dsmcutil 182
 - GUI and Web GUI clients 183
 - reporting a problem 183
 - server or storage agent
 - commands 182
 - messages 183
 - Windows 182
- hints and tips
 - device driver 187
 - disk subsystems 193
 - hard disk drives 193
 - NDMP filer-to-Tivoli Storage Manager server operations 213
 - SAN 197
 - SAN configuration 198
 - SAN device mapping 204
 - tape drives and libraries
 - adapter firmware changes 195
 - cabling between the computer and device changes 195
 - device driver changes 196

- hints and tips (*continued*)
 - tape drives and libraries (*continued*)
 - device firmware changes 195
 - error messages in system error log 197
 - loose cable connections 197
 - operating system changes 196
 - other hardware changed or fixed 196
 - replaced adapter 196
- historical reporting
 - Warehouse Proxy workspace 179

I

- IBM Global Security Kit
 - key management return codes 227
 - return codes 223
- IBM Software Support
 - submitting a problem xiii
- IBM Support Assistant x
- IDLETIMEOUT option 48
- image backup
 - client 17
 - error 17, 19
- INCLEXCL option 9
- InitSID.sap 99
- InitSID.utl 99
- inquire function
 - BR*Tools 112
 - BRRESTORE 112
- InstallAnywhere
 - exit codes 37
- Internet, searching for problem resolution x, xi
- IP address
 - configuring 80
- IPv6 98

J

- Japanese language
 - reporting problems 179
- Java archive
 - collector tool 77
- journal
 - restarting 27
- journal-based backup (JBB)
 - database viewing utility 28
 - determining 27
 - running in foreground 28

K

- key database file
 - out-of-synch 72
 - password recovery 72
- knowledge bases, searching x
- Korean language
 - reporting problems 179

L

- LAN-free setup
 - storage agent 123

- Linux image backup error 17
- Linux installation
 - SELinux AVC denial 37
- Linux Snapshot image backup error error message ANS1258E 19
- log files
 - DB2 upgrade 49
 - installation 37
 - location 104, 105
- LVSA
 - examining Windows system event log 29
 - forcing a memory dump 30
 - full memory dump 30
 - problem determination 29

M

- maintenance
 - client deployment 93, 94
- message definitions
 - IBM Tivoli Storage Manager 90
- messages
 - when using BR*Tools 105
- Microsoft diagnostic information
 - VSS 33
- Microsoft tuning
 - VSS 32
- monitoring 179
- moving data to other volumes 194

N

- NIS server
 - stoppage 45
- non-root user
 - running applications using the API 25
- ntbackup.exe 34

O

- open file support
 - best practices 31
 - problem determination 29
- option
 - IDLETIMEOUT 48

P

- performing ddtrace on HP-UX 191
- portlets
 - errors caused by starting or stopping 87
- problem determination
 - describing problem for IBM Software Support xii
 - determining business impact for IBM Software Support xii
 - submitting a problem to IBM Software xiii
- process ended 58
- process started 57
- process symptoms
 - files not expired 65

- process symptoms (*continued*)
 - migration does not run 66
 - migration only uses one process 66
- publications
 - download vii
 - order vii
 - search vii
 - Tivoli Storage Manager viii

R

- reporting 179
- reporting and monitoring agent
 - installation hang 180
 - tracing 175
- restore
 - BRRESTORE 113
 - troubleshooting 113
- RESTORE DB
 - ANR2971E with SQL code 52
 - common errors 53
 - incorrect environment variables 51
- RMAN
 - troubleshooting 108

S

- SAN
 - configuration 210
 - configuration between devices 200
 - configuration problems 212
 - devices supported 198
 - fibre channel switch configuration 199
 - fibre-channel link error report 201
 - gateway port settings 200
 - host bus adapter configuration 199
 - host bus adapters 198
 - vendor support 212
- SAN device mapping
 - disabling 205
 - errors 205
 - missing from the display of QUERY SAN 209
- SAN devices
 - storage agent 202
- scheduled event
 - status 7
- scheduler
 - client service restart 8
- SCSI devices 213
- secure sockets layer
 - determining errors 70
 - general return codes 223
- sequential media volume
 - tape 214
- server
 - database 48
 - diagnostic tips
 - change server options or the settings create errors 41
 - checking the server activity log 41
 - code page conversion failure 43
 - failing a scheduled client operation 42

- server (*continued*)
 - diagnostic tips (*continued*)
 - recreating the problem 41
 - resolving errors from reading or writing to a device 41
 - resolving failed connections by client or administrators 69
 - resolving server space issues 42
 - process 54
 - process messages 54
 - stoppage or loop errors 44
 - storage pool
 - ANR0522W error message 67
 - Collocate?=Yes 67
 - COPY ACTIVATEDATA command 68
 - high volume usage 67
 - resolving problems 66
 - simultaneous write 68
 - unable to store data 68
- server activity log
 - checking 87
 - checking for errors 7
- server command definition file 92
- server or storage agent
 - trace class 125
 - trace classes 131
- server stoppage
 - activity log 48
 - library files 46
 - resolving general problems 43
 - server error file (dsmserv.err) 45
 - system image 46
 - system logs 48
- SHOW commands
 - server or storage agent 144
- SID 103, 104, 105, 107, 108, 110, 111, 113
- Snapshot Difference
 - resolving problems 14
- snapshot directory 16
- software distribution 93, 94
- Software Support
 - contact xi
 - describing problem for IBM Software Support xii
 - determining business impact for IBM Software Support xii
- SSL 70
- stack trace
 - server or storage agent 130
- status
 - scheduled event 7
- 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 123
 - storage pool configured for simultaneous write 124
 - testing LAN-free configuration 124

- storage agent (*continued*)
 - SAN devices 202
- support for API
 - before calling IBM
 - files to gather 22
 - information to gather 21
- support information ix

T

- test flags
 - VSS 32
- Tivoli Integrated Portal
 - excessive memory consumption 80
- server stoppage
 - log analyzer tool 79
 - user authority problems 75
- Tivoli technical training ix
- trace
 - administration center 127
 - client
 - backup-archive client 165
 - device driver 157
 - enable client trace on command line 165
 - enable client trace while client is running 167
 - known problems and limitations 170
 - options 171
 - server or storage agent 129
 - trace classes
 - administration center 125
 - client 160
 - server or storage agent 131
 - trace data
 - is it compressed during backup-archive 174
 - is it encrypted during backup-archive 174
 - tracing
 - application programming interface (API) 177
 - client 159
 - tracing the reporting and monitoring agent
 - Windows 176
 - Traditional Chinese language
 - reporting problems 179
 - training, Tivoli technical ix
 - transient errors
 - VSS 31
 - troubleshooting
 - BACKINT 106
 - backup function 111
 - delete function 111
 - File Manager 106
 - IBM support 101
 - inquire function 112
 - random problems 97
 - reproducible problems 97
 - restore function 113
 - RMAN 108
 - tsmdiag utility 219

U

uninstall stoppage 54
util_par_file 99

V

vendor environment file 100
Volume Shadow Copy Services
 Windows 31
vsreq.exe sample program 33
VSS
 Microsoft diagnostic information 33
 Microsoft tuning 32
 ntbackup.exe 34
 test flags 32
 trace 33
 transient errors 31
 vsreq.exe sample program 33
 Windows 31
 Windows fixes 32

W

Windows
 VSS 31
Windows fixes
 VSS 32
wizards
 errors caused by starting or
 stopping 87



Program Number: 5608-ARM, 5608-ISM, 5608-ISX, 5608-SAN

Printed in USA

GC23-9789-02

