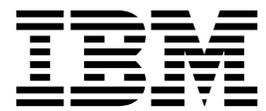


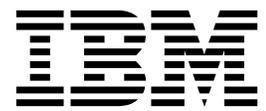
IBM Tivoli Storage Manager
Version 7.1.6

Problem Determination Guide



IBM Tivoli Storage Manager
Version 7.1.6

Problem Determination Guide



Note:

Before you use this information and the product it supports, read the information in “Notices” on page 211.

Third edition (June 2016)

This edition applies to version 7, release 1, modification 6 of IBM Tivoli Storage Manager (product number 5608-E01, 5608-E02, 5608-E03), and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 1993, 2016.

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

Contents

About this publication vii

Who should read this guide. vii

Publications vii

Chapter 1. Help facilities 1

Backup-archive client help 1

 Access help for the Client Service Configuration

 Utility (**dsmcuti1**). 2

Server or storage agent help 2

 Accessing server or storage agent help for
 commands 2

 Accessing help for messages 3

 Command-line interface help for the client 3

Reporting a problem with a help topic. 3

Chapter 2. Resolving client problems 5

Examining error messages 5

Examining the server activity log messages 5

Identifying when and where the problem can occur 5

Reproducing the problem 6

Collecting documentation to resolve problems with
the client application 6

Determining why the **dsmc**, **dsmadmc**, **dsm**, or **dsmj**
programs do not start 7

Resolving problems with client option sets 9

 Scenarios for resolving problems with client
 option sets 9

Resolving password expiration problems 10

Resolving LDAP-authenticated password problems 10

 Verifying your password authentication
 configuration 11

 The Tivoli Storage Manager server does not
 accept the LDAPPASSWORD 12

 Resolving problems with the LDAP directory
 server 12

 Auditing the LDAP directory server to clean up
 the server 14

 Error messages for LDAP authenticated
 passwords. 14

Resolving client scheduling problems. 17

 Determining the status of a scheduled event 18

 Checking for errors in the server activity log 18

 Starting and stopping the client service 19

Resolving errors when including or excluding client
files during backup processing 20

 Identifying files that are included or excluded by
 the server client option set 21

 Excluding files automatically from backup
 processing. 21

 Excluding files with the EXCLUDE.DIR
 statement 23

 Determining whether compression, encryption,
 and subfile backup statements include or exclude 24

 Using delimiters to include or exclude files. 25

 Resolving errors due to the incorrectly coded
 include or exclude list. 25

Resolving Snapshot Difference problems. 26

 Resolving snapshot directory problems for
 NetApp or N-Series file system volumes. 28

 Resolving login problems when using the
 encrypted file system on AIX operating systems 28

Resolving image backup errors 29

 Resolving Linux image backup errors. 29

 Resolving backup failures when you use Linux
 snapshot image backup 30

 Resolving errors during AIX JFS2 snapshot-based
 backup-archive and image backup. 31

Support solutions for the Tivoli Storage Manager
API 33

 Gathering API information before calling IBM
 support. 33

 Gathering API files before calling IBM Support 34

Determining whether data is sent to the Tivoli
Storage Manager storage agent rather than the
server 36

Running applications that use the API as a non-root
user ID 37

Journal Based Backup problem determination 39

 Determining if a backup will be journal-based. 39

 Running the journal daemon in the foreground 40

 The Journal Database Viewing utility. 40

Using Windows Volume Shadow Copy Services 42

 Defining VSS transient errors 42

 Defining Windows VSS test flags 42

 Volume Shadow Copy Services tuning 43

 Gathering VSS diagnostic information for
 Microsoft assistance 43

 Troubleshooting errors using a VSS trace 44

 Running VSS API calls with the vsreq.exe sample
 program 44

 Comparing Tivoli Storage Manager and
 Ntbackup.exe interaction with VSS 45

SHOW commands for the backup-archive client 45

Resolving problems for recovery of individual
Microsoft SQL databases from a virtual machine
backup 47

 Resolving database access problems 47

 Viewing active copies of Microsoft SQL databases 49

 Microsoft SQL databases with DBCS names 49

 Responses to messages for virtual machine
 backups with application protection 49

 Saving VSS XML manifest files 50

 Determining whether a virtual machine backup
 might fail 51

Chapter 3. Resolving Tivoli Storage Manager server problems 53

Recreating the problem 53

Checking the server activity log file and other log files	53
Checking system error log files for device errors	54
Reverting server options or settings	54
Restarting the scheduling service	54
Resolving server space issues	55
Allocating additional server memory	55
Configuring a server instance to use shared memory	56
Changing the copy frequency	56
Resolving RELABEL operation errors	57
Avoiding communication errors during import processing	57
Adding a self-signed certificate to the keystore	58
Determining why summary records for a client backup event are missing	59
Resolving installation and upgrade problems	60
Installation log files	60
Installation wizard fails to start	60
Resolving GSKit installation problems	60
Server instances are not created during an upgrade	61
Resolving a stopped uninstallation process	62
Client automatic deployment did not upgrade the client software	62
Resolving server stoppages	63
Resolving a stoppage or loop	64
Resolving wait state problems with external user repository servers	65
Finding the server error file (dsmserv.err)	66
Finding the system image (core file)	66
Retrieving library files for core analysis	67
Retrieving system log files	68
Retrieving the activity log	68
Detecting errors after a server service starts and stops	69
sqlib/db2dump directory causes shutdown	69
Resolving issues with database page verification	70
Resolving database errors	71
Resolving database manager starting problems	71
Tracing the user ID and Password plug-in	72
Limiting DB2 memory allocation	73
Retrieving DB2 version information	74
Locating DB2 diagnostic log files	74
DB2 upgrade log files	75
Resolving a missing or incorrect database ID file problem	76
Resolving problems with the BACKUP DB and the RESTORE DB commands	77
Characteristics of the \$\$_TSMDBMGR_\$\$ user ID	81
Resolving database reorganization problems	82
Analyzing the process symptoms to resolve problems	82
Reviewing process messages to determine the state of server operations	82
Analyzing the ANR1221E error message	88
Analyzing the ANR2317W error message	89
Analyzing error messages ANR1330E and ANR1331E	90
Files are not expired after reducing versions	93
Process symptoms indicate migration errors	93
Resolving storage pool issues	94

“ANR0522W Transaction failed...” message received	94
Storage pool experiences high volume usage after increasing MAXSCRATCH value	95
Storage pool is set to use collocation, but volumes contain data that is not collocated	95
Resolving storage problems for active data pools	96
Resolving issues with cloud-container storage pools	96

Chapter 4. Resolving Operations Center problems	99
Log files overview	99
Viewing the Operations Center log from within the Operations Center	100
Alerts are not updated immediately	100
Active tasks are not canceled immediately	101
Further known issues with the Operations Center	102

Chapter 5. Resolving communication problems	103
Resolving errors created when connecting to the server	103
Resolving failed connections by clients or administrators	103
Resolving Secure Sockets Layer errors	104
Recovering the key database file password	106
Troubleshooting the certificate key database	107

Chapter 6. Resolving storage agent problems	109
Checking the server activity log for storage agent information	109
Resolving an error caused by reading or writing to a device	109
Resolving problems caused by changing storage agent options	110
Resolving problems caused by changing server options or settings	110
Storage agent LAN-free setup	110
Resolving the issue of data being sent directly to the server	110
Resolving a disqualified LAN-free-enabled storage pool	111
Ensuring that data is transferred using a LAN-free environment	112

Chapter 7. Using trace to resolve problems	113
Starting an extended trace of the Operations Center	113
Tracing the Operations Center by enabling logging functions from within the Operations Center	113
Tracing the Operations Center by enabling functions in the logging configuration file	114
Enabling a trace for the server or storage agent	115
Enabling a stack trace for messages for the server or storage agent	116
Trace classes for a server or storage agent	117
Show commands for the server or storage agent	131

Enabling a trace for the Tivoli Storage Manager device driver	140
Tracing from the server console	141
Tracing data from a command shell for AIX, Oracle Solaris, and Windows	142
Tracing to detect a code page conversion failure	142
Tracing data for the client	142
Client and Journal Daemon traceflags	143
Client trace classes	144
Enabling a backup-archive client trace	149
Determining if data is encrypted or compressed during backup-archive by using a trace.	158
Tracing data for the API.	159
Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on an AIX or Linux system	160
Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on a Windows operating system.	162
Chapter 8. Resolving data storage problems	165
Resolving unreadable data problems	165
Checking the server activity log to resolve data storage issues	165
Checking HELP for messages issued for a data storage problem	165
Recreating the data storage problem.	165
Resolving data storage errors related to reading or writing to a device	166
Changing the storage hierarchy to resolve data storage problems	166
Changing the server policies to resolve data storage problems	166
Resolving a data storage backup or copy problem that occurs only with a specific node	167
Resolving a data storage problem that occurs only for a specific volume	167

Hints and tips for storage	167
Device driver hints and tips	167
Hard disk drives and disk subsystems hints and tips.	173
Tape drives and libraries hints and tips.	175
SAN hints and tips	177
NDMP filer-to-Tivoli Storage Manager server operation hints and tips	192
Resolving SCSI device problems	192
Resolving sequential media volume (tape) errors through messages ANR0542W or ANR8778W	193

Appendix A. Getting call stack information from a core file 195

Appendix B. Running the tsmdiag utility 197
tsmdiag utility options 198

Appendix C. IBM Global Security Kit return codes. 199

Appendix D. Accessibility features for the Tivoli Storage Manager product family. 209

Notices 211

Glossary 215

Index 217

About this publication

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

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

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

Who should read this guide

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

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

Publications

The Tivoli Storage Manager product family includes IBM Tivoli Storage FlashCopy® Manager, IBM Tivoli Storage Manager for Space Management, IBM Tivoli Storage Manager for Databases, and several other storage management products from IBM.

To view IBM product documentation, see IBM Knowledge Center.

Chapter 1. Help facilities

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

Backup-archive client help

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

Syntax

```
▶▶—dsmc help—▶▶
```

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

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

Parameters

command-name [*subcommand-name*]

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

option-name

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

TOC-section-number

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

[*ANS*]*message-number*

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

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

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

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

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

Access help for the Client Service Configuration Utility (dsmcutil)

Windows

To obtain help information for the Tivoli Storage Manager Client Service Configuration Utility, you must issue the **DSMCUTIL HELP** command.

When you issue the **DSMCUTIL HELP** command, the help information is displayed within the Windows help utility.

Server or storage agent help

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

Accessing server or storage agent help for commands

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

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

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

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

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

```
help 3.46.3
```

Use topic numbers to display command-line help for subcommands. **DEFINE DEVCLASS** is an example of a command that has subcommands. For example, you can specify the **DEFINE DEVCLASS** command for 3590 device classes and for 3592 device classes:

3.0 Administrative commands

```
...
3.13.10 DEFINE DEVCLASS (Define a device class)
  3.13.10.1 DEFINE DEVCLASS (Define a 3590 device class)
  3.13.10.2 DEFINE DEVCLASS (Define a 3592 device class)
  ...
```

To display help for the **DEFINE DEVCLASS** command for 3590 device classes, type:
help 3.13.10.1

Accessing help for messages

Issue the help command to access help for messages.

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

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

Command-line interface help for the client

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

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

Reporting a problem with a help topic

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

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

Chapter 2. Resolving client problems

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

Examining error messages

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

If it is set, the Tivoli Storage Manager client QUIET option suppresses the display of all messages in screen output. However, all messages are still logged in the log files. Turning off the QUIET option might facilitate troubleshooting operations because you can see the messages on screen, as they occur.

Look for any ANSnnnnx messages that are issued to the console. Messages are also logged. Scheduler messages are logged in the `dsmsched.log` file. Client messages are logged in the `dsmerror.log` file. Descriptions of the messages and API return codes are provided in Messages, return codes, and error codes. Online help is also available for system messages. To get online help for a message when you are using the command-line client, type **HELP ANS_***nnnnx*, where *nnnn* is the message number and *x* is the message type.

Examining the server activity log messages

Use the **QUERY ACTLOG** command to view the server activity log file and the messages issued for this 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 client processing often occur only when you are performing specific operations, at certain times, or only on certain client computers.

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

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

Reproducing the problem

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

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

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

Collecting documentation to resolve problems with the client application

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

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.

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 `dsm Sched.log`.
- Web client log. The error log for the web client is `dsmwebcl.log`.
- Options files. Information about the options that you set for the clients can facilitate troubleshooting and problem resolution. Much of this information is contained in the following files:
 - The client options file (`dsm.opt`). This file exists for all clients on all operating systems.
 - The client system-options file (`dsm.sys`). This file is only used on AIX®, HP-UX, Linux, Mac OS X, and Oracle Solaris clients.
 - The include-exclude file. This file contains the objects to include or exclude from client operations. Its location is set by the client `incl excl` option.
- Trace data. If the tracing facility was active, the file that contains the trace data can be provided to support.
- Application dump. When the backup-archive client stops running unexpectedly, many platforms generate an application dump. The operating system provides the application dump.
- Memory dump. If the backup-archive client stops, a memory dump can be generated that can then be used to help with diagnosis. The type of system determines how the memory dump occurs, and the operating system provides the memory dump.

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

- A list of all the software that is 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 that are defined on the server that apply to this client node. Issue the **QUERY CLOPTSET** command to search for the client option sets.
- Server options. A number of server options are used to manage the interaction between the backup-archive client and server. An example of one such server option is `TXNGROUPMAX`.
- Information about this node as it is defined to the server. To collect this information, issue the **QUERY NODE *nodeName* F=D** command by using an administrative client that is connected to the server.
- Schedule definitions for the schedules that apply to this node. The schedule definitions can be queried from the server when you issue the **QUERY SCHEDULE** command.
- The policy information that is configured for this node on the server. The policy information can be queried from the server when you issue the **QUERY DOMAIN**, **QUERY POLICYSET**, **QUERY MANAGEMENTCLASS**, or **QUERY COPYGROUP** commands.

Determining why the `dsmc`, `dsmadmc`, `dsm`, or `dsmj` programs do not start

The backup-archive client uses the `dsmc`, `dsmadmc`, `dsm`, or `dsmj` programs in its startup procedure. When one of these programs does not start, the backup-archive client does not start.

The `dsmc`, `dsmadmc`, `dsm`, or `dsmj` programs have the following definitions:

dsmc The backup-archive command-line client.

dsmadmc

The administrative command-line client.

Windows **dsm**

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

The backup-archive client graphical user interface (GUI). The Oracle Java™ runtime version is checked when you first start the Java GUI. In some cases, this check is not completed properly and the `dsm` or `dsmj` startup might fail with a “bad number” message.

Processing stops and the following message is displayed if the `dsmc`, `dsmadmc`, `dsm`, or `dsmj` program does not start:

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

Remember: The `dsmerror.log` file is used only as an example file in the message.

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

1. The name and the directory that is specified by the `ERRORLOGNAME` option are used.

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

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

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

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

The client does not try to create log files in the root directory. Message ANS1398E is displayed when the method in the first rule 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 is preserved by Tivoli Storage Manager code.

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

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

Windows 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.
```

AIX **HP-UX** **Linux** **Mac OS X** **Solaris** Extra 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 receive the ANS1398E error if you try to run Tivoli Storage Manager applications by using an error log file that was already generated by root, that is left with default permissions. For data protection clients, you might receive only 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/messages/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/messages/tsm/dsmerror.log`

Resolving problems with client option sets

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

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

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

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

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

Scenarios for resolving problems with client option sets

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

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

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

Scenario 1: Having an environment where restoring is a high priority.

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

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

Scenario 2: Using workstations that are on a slow network with limited space for data on the server.

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

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

Scenario 3: Using a database that does not stop.

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

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

Scenario 4: Finishing backups using a fast network and wanting to make the best possible use of client resources.

Set the RESOURCEUTILIZATION option to the maximum amount. Issue the following server command:

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

Resolving password expiration problems

If you receive a client authentication error, it might be as a result of an expired password. Password expiration does not apply to node or administrator passwords that authenticate with an LDAP directory server.

Procedure

Complete the following steps to change the expired password period:

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

If a Windows client node cannot connect to the server after it is 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 and uses the **DSMC SCHED** command, Tivoli Storage Manager uses the node name in the client options file to contact the server. However, when the scheduler runs as a Windows service, Tivoli Storage Manager uses the node name in the Windows registry.

2. For the Windows client, issue the **DSMCUTIL UPDATE SCHEDULE** command to achieve the following results:
 - With the *node* parameter, address how to change the node name that is 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)

Resolving LDAP-authenticated password problems

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

Before you can use the LDAP-authenticated password, you must configure the LDAP directory server to communicate with the Tivoli Storage Manager server. Ensure that the access control list on the LDAP directory server grants full authority to a user (LDAPUSER) over the base distinguished name (Base DN).

Verifying your password authentication configuration

If you configured the server to authenticate passwords with an LDAP directory server and you are receiving errors, review the configuration steps. You must ensure that the Tivoli Storage Manager server and the LDAP directory server are properly configured.

Procedure

1. Open the `dsmserv.opt` option file and find the **LDAPURL** option, which contains the server and the base distinguished name (Base DN). You can add more values to the **LDAPURL** option, with each URL value being up to 1024 characters. The port number is optional. The default port number is 389. Each URL configuration must contain the following values:

- An LDAP directory server name
- The Base DN of the namespace or suffix that the LDAP directory server maintains. The DN format must comply with the directory server that you choose.

The **LDAPURL** option must conform to the following specifications:

- If you specify multiple URLs, follow these guidelines:
 - Each URL must be on a separate line
 - Each URL must point to a different external directory and all external directories must contain the same data
- Each URL must begin with the following value: `ldap://`

For example:

```
LDAPURL ldap://zapp.storage.dallas.gov/ou=tsmdata,dc=storage,dc=dallas,dc=com
```

The URL that you specify cannot be a secure URL, meaning that it cannot begin with `ldaps://`.

2. View the **LDAPUSER** or **LDAPPASSWORD** settings by issuing the **QUERY STATUS** command. Define the **LDAPUSER** who can then add or remove entries, and change or reset passwords. If the **LDAPUSER** is not defined, issue the **SET LDAPUSER** command to define the LDAP directory server administrator.

Important: If the value for the **LDAPUSER** parameter includes special characters, enclose the value with quotation marks. For example:

```
set ldapuser "cn=bill cook,cn=users,dc=storage,dc=dallas,dc=gov"
```

3. View the **LDAPUSER** or **LDAPPASSWORD** settings by issuing the **QUERY STATUS** command. If a password is not defined, designate one for the **LDAPUSER** by issuing the **SET LDAPPASSWORD** command.

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

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

Requirement: If you use any special characters when you issue the **SET LDAPPASSWORD** command, surround them with quotation marks. For example:

```
set ldappassword "Pa$$=w0rd"
```

The Tivoli Storage Manager server does not accept the LDAPPASSWORD

If you receive a warning that the LDAPPASSWORD is not valid, the problem might not be with the password.

If you issue a **SET LDAPPASSWORD** command and receive error messages ANR3114E or ANR3116E, Tivoli Storage Manager might not be configured correctly. Examine any server messages that occurred around the time that ANR3114E or ANR3116E were issued to determine the cause of the errors. A common problem that you might see is that an incorrect value is set for the **SET LDAPUSER** command. The user must be entered in distinguished name (DN) format. For example:

```
ou=armonk,cn=tsmdata,uid=9A73819745
```

If the value does not conform to the DN, the **LDAPUSER** is not defined and you are unable to set the LDAPPASSWORD. A DN typically consists of a comma-separated list of naming attributes and value pairs. The following list shows the more commonly used naming attributes:

- The common name (cn)
- The user ID (uid)
- The organizational unit (ou)
- The domain component (dc)
- The organization (o)
- The country (c)

For example:

```
cn=Jack Spratt,ou=marketing,dc=tucson,dc=storage,dc=com  
uid=abbynornal,ou=sales,dc=tucson,dc=storage,dc=com  
uid=cbukowski,ou=manufacturing,o=storage,c=us
```

Resolving problems with the LDAP directory server

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

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

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

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

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

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

Without SSL/TLS

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

With SSL/TLS

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

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

Without SSL/TLS

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

With SSL/TLS

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

For the previous commands, the following parameters apply:

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

Resolving problems with locked nodes and administrators

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

Procedure

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

1. Return the password to the server by issuing the following example command:
`update node node_x new_pw authentication=local`
2. Clean out the LDAP directory server by issuing the following example command:
`audit ldapdirectory fix=yes wait=no`
This command removes nodes or administrator IDs that are stored on the LDAP directory server that do not authenticate passwords with an LDAP directory server.
3. Log out of the node.
4. Issue the following command:
`update node node_x newest_pw authentication=ldap`
5. Log on to the node with the new password.

Auditing the LDAP directory server to clean up the server

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

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

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

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

Example: Auditing the LDAP directory server

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

```
AUDIT LDAPDIRECTORY FIX=YES
```

The command produces a list of all nodes and administrators that are removed from the LDAP directory server. Also produced is a list of all nodes and administrators that are missing from the LDAP directory server. If you want to see what is out-of-sync, use the default **FIX=NO** setting to report on the discrepancies between the servers.

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

Error messages for LDAP authenticated passwords

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

These error messages are the result of communicating with an LDAP directory server:

ANR3114E

Message ANR3114E is issued whenever an unexpected error is encountered during an LDAP operation. The message gives you more information to assist you in resolving the error. For example,

```
ANR3114E LDAP error
```

LDAP error code (error description) occurred during operation.

LDAP error code

The error number that is returned by either the LDAP client interface or the LDAP Directory server.

error description

A description of the *LDAP error code*, indicating the cause of the error.

operation

The LDAP client operation that is running when the error occurred.

In the following example, error code 53 is returned by the LDAP client interface or the LDAP directory server. The operation that was in progress at the time of the error is also flagged. In this example, `ldap_search_s`.

ANR3114E

LDAP error 53 (*DSA is unwilling to perform*) occurred during `ldap_search_s`.

ANR3115E

Message ANR3115E is issued when there is an error with the LDAP directory server. For example,

ANR3115E The LDAP directory server returned the following error message (*LDAP server message*) with the LDAP error.

LDAP server message

This message text is returned by the LDAP directory server and gives more information about the error that just occurred.

ANR3116E

Error message ANR3116E is issued when the Global Security ToolKit (GSKit) component encounters an error during an LDAP operation. GSKit provides Secure Sockets Layer/Transport Layer Security (SSL/TLS) for LDAP operations. This error message is usually related to SSL/TLS, certificates, cryptography, or network operations. For example:

ANR3116E LDAP SSL/TLS error *GSKIT error code* (*error description*) occurred during *operation*.

GSKit error code

The error number that is returned by the GSKit component.

error description

A text description that is associated with the *error code* indicating the cause of the error.

operation

The LDAP client operation that is running when the error occurred.

If you cannot determine the cause of the errors, work through the following steps:

1. Examine the server messages that were issued around the same time as the error message to determine the cause and the impact of the error. Issue the **QUERY ACTLOG** command to view the activity log file and to search for error messages.
2. Look for network problems.
3. Check the status of the LDAP directory server.
4. For error message ANR3116E, look for problems with the certificates that the LDAP directory server uses or the Tivoli Storage Manager server key database (`cert.kdb`).
5. Examine the LDAP directory server log files.

6. Use LDAP utilities such as “ldapsearch” or “ldp” to isolate the problem.

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

Table 1. Errors that might occur when you authenticate passwords with an LDAP directory server

Error messages	Resolution
ANR3114E LDAP error 118 (The SSL library cannot be loaded) ANR3116E LDAP SSL/TLS error 118 (Unknown SSL error) ANR3103E Failure occurred while initializing LDAP directory services	The library path might not be set properly. Make sure that you are using the correct version of the GSKit.
ANR3114E LDAP error 116 (Failed to connect to the SSL server) ANR3116E LDAP SSL/TLS error 406 (I/O error) ANR3103E Failure occurred while initializing LDAP directory services ANR2732E Unable to communicate with the LDAP directory server	The level of GSKit might be incorrect on the Tivoli Directory Server. Upgrade GSKit to the correct level. See technote 1469388. For Active Directory, disable automatic root certificates updates with Windows Update if an internet connection is not available.
ANR3114E LDAP error 52 (DSA is unavailable) ANR3103E Failure occurred while initializing LDAP directory services ANR2732E Unable to communicate with the LDAP directory server	The Active Directory server does not have a certificate available for TLS/SSL. Create a signed certificate that can be used by Microsoft Active Directory.
ANR3114E LDAP error 116 (Failed to connect to SSL server) ANR3116E LDAP SSL/TLS error 414 (Bad certificate) ANR3103E Failure occurred while initializing LDAP directory services ANR2732E Unable to communicate with the LDAP directory server	The LDAP directory server certificate is not trusted. Add the root certificate authority (CA) certificate to the Tivoli Storage Manager server key database file (cert.kdb) and verify that the certificates are not expired.
ANR3094E The distinguished name (DN) that is specified in the LDAPURL option does not exist on the LDAP directory server ANR3103E Failure occurred while initializing LDAP directory services	If the DN exists, the LDAPUSER might not have full access control rights to the Base DN that is specified in the LDAPURL option.

Table 1. Errors that might occur when you authenticate passwords with an LDAP directory server (continued)

Error messages	Resolution
<p>ANR3114E LDAP error 50 (Insufficient access)</p> <p>ANR1885E LDAP directory service initialization: Permission was denied when the LDAP directory entry was accessed as LDAPUSER</p> <p>ANR3103E Failure occurred while initializing LDAP directory services</p> <p>ANR1885E SET LDAPPASSWORD: Permission was denied when the LDAPUSER entry was accessed</p>	<p>The LDAPUSER does not have full access control rights to the base DN that is specified in the LDAPURL option.</p>
<p>ANR3114E LDAP error 116 (Failed to connect to SSL server)</p> <p>ANR3116E LDAP SSL/TLS error 420 (Socket closed)</p>	<p>For Tivoli Directory Server, the SSL_TIMEOUT_MILLISEC is not set high enough. See technote 1233758.</p>
<p>ANR3114E LDAP error 4 (Size limit exceeded)</p>	<p>Increase the LDAP server search size limit to accommodate the total number of LDAP-authenticated nodes and administrators.</p>
<p>ANR3114E LDAP error 91 (Connection error) occurred during ldap_sasl_bind.</p> <p>ANR3103E Failure occurred while initializing LDAP directory services.</p>	<p>The LDAP server is not active or is offline.</p>

Resolving client scheduling problems

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

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

- Additions and changes to the client options are not recognized by the client scheduler until the next scheduled start. Deletions made to the client option set do not take effect until you restart the scheduler.
- Additions, deletions, and changes made to the client-acceptor-managed schedules are recognized at the next scheduled start.
- Use the **SHOW PENDING** diagnostic tool to show schedules, nodes, and the next scheduled run time.
- From the client options file, view the `dsm.sys` stanza for the node and the `MANAGEDSERVICES`, `PRESCHEDCMD`, and `POSTSCHEDCMD` option values for information after a node misses a scheduled event.

Determining the status of a scheduled event

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

About this task

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

Procedure

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
```

What to do next

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 client. See the **QUERY EVENT** command for scheduled and completed events.

Checking for errors in the server activity log

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

When you are checking the server activity log, narrow the query results down to the time frame around 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
```

Afterward you can issue one of the following queries:

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

The 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 to inspect the schedule log file. The schedule log is typically saved to the `dsmsched.log` file, and is typically stored in the same directory as the `dsmerror.log` file. The location of the log file can be specified by using client options, so you might refer to the options file to see whether the `SCHEDLOGNAME` option was used to relocate the log file. On Windows, the schedule log can also be relocated by an option setting that is part of the schedule service definition. You can issue the **DSMCUTIL QUERY** command to check whether this option was set. When you locate the schedule log, search through the

file to find the time period corresponding with the start date and time of the scheduled event in question. The following list shows common search parameters:

- If you are investigating a missed event, check the details of the previous event, including the time at which the previous event finished.
- 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 (typically named `dsmerror.log`).

Starting and stopping the client service

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

Tip: When you manage many clients that run 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 operating systems, standard operating system utilities are required.

Windows To remotely manage the client scheduler service by using the **DSMCUTIL** command 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
```

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

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

AIX **HP-UX** **Linux** **Solaris** If you use the traditional method to manage the scheduler, 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. The following example shell script shows you how to recycle the Tivoli Storage Manager scheduler process:

```
#!/bin/ksh
# Use the following script to kill the currently running instance
# of the TSM scheduler, and restart the scheduler in nohup mode.
#
# This script will not work properly if more than one scheduler
# process is running.
# If necessary, the following variables can be customized to allow an
# alternate options file to be used.
# export DSM_DIR=
# export DSM_CONFIG=
# export PATH=$PATH:$DSM_DIR
# Extract the PID for the running TSM Scheduler
```

```

PID=$(ps -ef | grep "dsmc sched" | grep -v "grep" | awk {'print $2'});
print "Original TSM scheduler process using PID=$PID"
# Kill the scheduler
kill -9 $PID
# Restart the scheduler with nohup, redirecting all output to NULL
# Output will still be logged in the dsmsched.log
nohup dsmc sched 2>&1 > /dev/null &
# Extract the PID for the running TSM Scheduler
PID=$(ps -ef | grep "dsmc sched" | grep -v "grep" | awk {'print $2'});
print "New TSM scheduler process using PID=$PID"

```

AIX **HP-UX** **Linux** **Solaris** **Mac OS X** If you want to use the CAD managed method to manage the client scheduler, set the `managementservices` option to **schedule** or **schedule webclient** in the `dsm.sys` file. For Mac OS X, if you do not specify the `managementservices` option, the CAD manages both the scheduler and the web client, by default.

AIX **HP-UX** **Solaris** Add the following entry into the system startup file (`/etc/inittab` for most platforms):

```

tsm::once:/usr/bin/dsmcad > /dev/null 2>&1 # TSM Client
Accepter Daemon

```

Linux The backup-archive client installation program creates a startup script for the CAD (**dsmcad**) in the `/etc/init.d` directory. You can start, stop, restart, and query the CAD by using the standard **service** command on Linux. For example:

```

# service dsmcad start
# service dsmcad stop
# service dsmcad restart
# service dsmcad status

```

To enable the CAD to start automatically after a system restart, add the service as follows, at a shell prompt:

```
# chkconfig --add dsmcad
```

Mac OS X You can start or stop the CAD with the **launchd** utility. To start the CAD, issue the following command in the Terminal window:

```
/bin/launchctl load -w com.ibm.tivoli.dsmcad
```

To stop the CAD, issue the following command in the Terminal window:

```
/bin/launchctl unload -w com.ibm.tivoli.dsmcad
```

You can also control the CAD with the **TSM Tools for Administrators** application.

Resolving errors when including or excluding client files during backup processing

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

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

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

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

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

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

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

Excluding files automatically from backup processing

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

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

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

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

Windows For example, the following output is displayed when you issued the **DSMC QUERY INCLEXCL** command:

```
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 2 for the files that are automatically excluded.

Table 2. Files automatically excluded during backup processing

Platform	Files Excluded
<p>Windows</p>	<ul style="list-style-type: none"> • Files that are enumerated in the HKLM\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup registry key • The client staging directory C:\ADSM.SYS • Internet Information Server (IIS) metafiles (these files are processed in the system object or system state backup) • Registry files (these files are processed in the system object or system state backup) • Client trace files • System files <p>Windows system files are silently excluded from the system drive backup processing and cannot be included.</p> <p>To process these Windows system files, you must issue a DSMC BACKUP SYSTEMSTATE command.</p> <p>The Windows system files are excluded from the system drive backup processing because they are sent during the system object or system state backups. System files are boot files, catalog files, performance counters, and files that are protected by the Windows system file protection (sfp). These files are not processed during backup of the system drive. However, the files are excluded internally from the system drive processing 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 dsmfinfo.txt file) ===== SFP c:\windows\system32\ahui.exe (protected) c:\windows\system32\apphelp.dll (protected) c:\windows\appatch\apphelp.sdb (protected) c:\windows\system32\asycfilt.dll (protected)</pre>
<p>AIX HP-UX Linux Solaris</p>	<p>Client trace file</p>
<p>Mac OS X</p>	<ul style="list-style-type: none"> • Volatile, temporary, and device files that are used by the operating system • Client trace files

Excluding files with the EXCLUDE.DIR statement

The EXCLUDE.DIR statement excludes all directories and files under the parent directory.

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

For example, consider this set of include-exclude statements:

```
Linux Mac OS X Solaris
exclude.dir /usr
include ../../*.o
```

```
Windows
exclude.dir C:\Users
include C:\...\*.o
```

The INCLUDE statement in this example indicates that all files with a .o extension are included, but the preceding EXCLUDE.DIR statement excludes all files in the /usr or C:\Users 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 that end with .o, use the following syntax:

```
AIX HP-UX Linux Mac OS X Solaris
exclude /usr/../../*
include ../../*.o
```

```
Windows
exclude C:\Users\...\*
include C:\...\*.o
```

When you use wildcards in include-exclude statements, use * if you want to include or exclude all the files, rather than the *.*.*.* pattern. The *.*.*.* pattern means to include or exclude all files that contain at least one dot (.) character, while * means to include or exclude all files. If you use *.* files that contain no dot characters (such as C:\MYDIR\MYFILE on Windows) are not filtered.

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

If you use the command-line client to start a selective backup, or a partial incremental backup, of a single file, the file is processed, even if there is an EXCLUDE.DIR statement that excludes one of the parent directories in the file path.

For example, consider the following include-exclude statement that is used in subsequent command-line actions:

```
AIX HP-UX Linux Mac OS X Solaris
exclude.dir /home/spike
```

```
Windows
exclude.dir C:\Users\spike
```

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

```
HP-UX Linux Mac OS X Solaris
dsmc selective /home/spike/my.file
```

Windows

```
dsmc selective C:\Users\spike\my.file
```

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

AIX

HP-UX

Linux

Mac OS X

Solaris

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

Windows

```
dsmc selective "C:\Users\spike\my.*"
```

Important: A subsequent incremental backup of the /home file system renders inactive the /home/spike/my.file file. Likewise, on Windows, a subsequent incremental backup of the C:\Users directory renders inactive the C:\Users\spike\my.file file.

Do not end EXCLUDE.DIR statements with a directory delimiter.

The following examples show incorrect EXCLUDE.DIR statements, due to a directory delimiter at the end of the directory path:

AIX

HP-UX

Linux

Solaris

```
exclude.dir /usr/
```

Mac OS X

```
exclude.dir /Users/
```

Windows

```
exclude.dir c:\directory\
```

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

AIX

HP-UX

Linux

Solaris

```
exclude.dir /usr
```

Mac OS X

```
exclude.dir /Users
```

Windows

```
exclude.dir c:\directory
```

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 is included for backup processing.

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

Consider the following example:

AIX

HP-UX

Linux

Mac OS X

Solaris

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

Windows

```
exclude c:\Users\file.o
include.compress c:\Users\*.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.” Do not interpret the `INCLUDE.COMPRESS` statement as “back up all files that match the pattern `/usr/*.o` and compress them.” If you want to back up the `/usr/file.o` file in this example, you must remove the exclude statement.

Using delimiters to include or exclude files

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

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

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

Windows

Using the backwards slash “\” and the volume delimiter “:”

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

AIX

HP-UX

Linux

Mac OS X

Solaris

Using the forward slash “/”

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

Resolving errors due to the incorrectly coded include or exclude list

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

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

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

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

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

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

Resolving Snapshot Difference problems

AIX

Linux

Windows

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 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 on AIX and Linux, or one with administrative authority on Windows. Set the authority level to be the same as the authority level used when you map or mount the filer volume. Ensure that you use the fully qualified host name or the dotted IP address format for the filer name. Issue the backup-archive client **SET PASSWORD** command to save this user ID/password information.

Remember: The **DSMC SET PASSWORD** command is extended to save “filer” type passwords.

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

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

- NetApp/N-Series filers that are 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** **Linux** Network file system (NFS) attached volumes

The Snapshot Difference feature does not support the following features:

- SAN-attached NetApp/N-Series volumes
- QTrees or subdirectories
- Vfiler volumes with a filer that is running ONTAP V8.1.0 or earlier are not supported. Vfiler volumes with a filer that is running ONTAP V8.1.1 or later are supported.

Windows

Verifying the filer volume type

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

Snapshot Difference restrictions

The lack of Unicode support from NetApp prevents Tivoli Storage Manager from processing any files that use characters that are not in the 7-bit ASCII range. Tivoli

Storage Manager can back up only names that contain ASCII characters. Two Snapshot Difference behaviors were noted when testing with Unicode characters:

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

The following are situations under which files and 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 do remove the rule that excluded the file. A snapshot-assisted incremental backup command with the `snaptiff` option does not detect this `incl/excl` change because it detects file changes only between two snapshots. The files themselves must be changed in order for the Snapshot Difference API to detect the change and for Tivoli Storage Manager to back up the file.
- You added an include statement to the options file. This include statement takes effect only if the file is detected to have changed by the Snapshot Difference API. The files might not get backed up because Tivoli Storage Manager is not inspecting each file on the volume during the backup operation.
- You explicitly delete a file from Tivoli Storage Manager inventory by issuing the **DSMC DELETE BACKUP** command. The Snapshot Difference API does not detect that a file was manually deleted from Tivoli Storage Manager by you. Therefore, the file remains unprotected in 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 back up the file again.
- Policy changes such as changing the policy from `Mode=modified` to `mode=absolute` are not detected. The entire file space is deleted from 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 `snaptiff` 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.

You can use the following trace flags for Snapshot Difference processing:

- `enter`
- `exit`
- `general`
- `snapshot`
- `hci`
- `hci_detail`

- diskmap
- diskmap_detail
- hdw
- hdw_detail
- bacache
- snapdiffdb

AIX

Linux

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.
```

AIX

Linux

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

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

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

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

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

Important: Linux If you run a full NetApp SnapDiff backup and then use the NFS4 method to mount the NetApp volume onto the server, another full NFS backup occurs. To avoid a full backup, use the undocumented **SNAPDIFFINCR** test flag to force incremental processing on entries that have already been processed. For example, `-test=snapdiffincr`.

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

AIX

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

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

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

Start the client in one of the following ways:

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

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

```
efskeymgr -n
```

Resolving image backup errors

AIX

Linux

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

Resolving Linux image backup errors

Linux

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

About this task

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:

Procedure

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

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

or

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

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

The output is similar to the following output:

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

The library version must be Version 1.01 or later.

2. Verify the installation after the upgrade.

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

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

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

Resolving backup failures when you use Linux snapshot image backup

Linux

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

Before you begin

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

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

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

About this task

The following example shows the output that is 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.
```

Procedure

Complete the following steps to load the modules:

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

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

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

```
# modprobe dm_snapshot
```

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

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

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

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

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

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

Results

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

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

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

AIX

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

Before you begin

The following situations illustrate when 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 started 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 started 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.

About this task

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 accept 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 that prompts the user to delete the snapshots manually.

Procedure

To manually delete a snapshot, issue the following commands in the specified order:

1. `snapshot -q -c ' ' <SRCFS>`
2. `df -k`
3. `umount -f /tsm*`
4. `rmdir /tsm*`
5. `snapshot -d /dev/tsm*`

If the snapshot delete process fails with “Device Busy” or some other error message, issue the `umount -f <srcfs>` command to unmount the source file system. Then, try to delete the snapshot again.

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.

Results

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 file system. The following cases indicate where Tivoli Storage Manager does not try to delete older snapshots:

- If the snapshot was not created by Tivoli Storage Manager. Tivoli Storage Manager names its snapshots with a “tsm” prefix to distinguish them from other snapshots that are created for the same file system. If the snapshot was not

created by Tivoli Storage Manager, an error message is generated that asks the user to delete the older snapshot and the operation again.

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

- Snapshot monitoring
- Snapshot deletion

To use these features, ensure that your operating system level is at AIX 6.1 or later.

Support solutions for the Tivoli Storage Manager API

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

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

See the *Using the Application Programming Interface* or IBM Support Assistant for more information.

Gathering API information before calling IBM support

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

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

- On what operating system is the problem being experienced?
- What is the exact level of the operating system, including all service packs and hot fixes that were applied?
- What is the exact level of the IBM Tivoli Storage Manager API?
- What is the exact level of the Tivoli Storage Manager server?
- What is the Tivoli Storage Manager server platform and operating system level?
- What is the exact level of the Tivoli Storage Manager storage agent (if LAN-free environment)?
- What is the Tivoli Storage Manager storage agent platform and operating system level (if LAN-free environment)?
- What applications are running on the system?
- What steps are required to recreate the problem? If you cannot recreate the problem, what steps caused the problem?

Gathering API files before calling IBM Support

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

Gather as many of the following files before you contact IBM Support:

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

The following two options files are used on Linux and UNIX 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 that is referenced by the **DSMI_CONFIG** environment variable. For Linux and UNIX, the default options file is `dsm.sys` and is found in the directory that is referenced by the **DSMI_DIR** environment variable.

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

DSMI_CONFIG

The fully qualified name for the client options file.

DSMI_DIR

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

DSMI_LOG

The `DSMI_LOG` variable points to the path for the `dsierror.log` file. If the client `errorlogname` option is set, the location that is specified by that option overrides the directory that is specified by `DSMI_LOG`.

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

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

Verifying that the API uses the correct options file

AIX

HP-UX

Linux

Mac OS X

Solaris

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

Procedure

Complete the following steps to verify that the API uses the correct options file or server stanza:

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

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

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

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

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

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

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

3. Check the environment variable values that were mentioned in “Gathering API files before calling IBM Support” on page 34 or repeat steps 1 and 2 with a different options file or server stanza.
4. Remove the option that is inserted in step 1.

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

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

Procedure

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

1. Add the following trace options to the client options file before you back up or archive objects:

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

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

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

This statement is followed by the following trace statement:

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

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

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

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

For example:

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

What to do next

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

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

AIX

HP-UX

Linux

Mac OS X

Solaris

You must perform specific steps if you are logged on as a non-root user ID who is attempting to run an application that uses the application programming interface (API).

Procedure

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

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

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

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

```
$ 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. Optional: Check the ownership and permissions of the Trusted Communication Agent (TCA) only if the **PASSWORDDIR** option is not used or points to a directory that the user does not have read/write access to. This file is in the directory indicated by the **DSMI_DIR** environment variable. For example, the following TCA has the correct ownership and permissions:

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

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

```

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

```

- The root or authorized user must generate the TSM.PWD password file by using either the backup-archive client or the **dapism** sample API application. An authorized user is any non-root user ID who has read and write access to the stored password (TSM.PWD file). The 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*:

```

# dapism
*****
* 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.

```

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

Journal Based Backup problem determination

Windows

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

Determining if a backup will be journal-based

Windows

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

About this task

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

Procedure

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

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

2. Perform a full incremental backup on the corresponding file system while the file system is actively being journalized. This full incremental backup must set the “Last Backup Completed” date on the 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 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

AIX

Linux

Windows

You can increase performance by restarting a valid journal.

The reasons for restarting a valid journal:

- Error conditions in the journal daemon
 - Buffer overflow errors caused by excessive change activity on the journal file system being monitored for changes
 - Journal database access errors (disk full errors, etc.)
- Request by a backup client
- Clients will issue a journal restart request when it is determined that a journal file system lacks integrity for one of the following reasons:
 - Server filespace no longer exists
 - 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

Windows

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

Start the journal daemon from a Windows command prompt as follows:

```
tsmjbbd.exe i
```

The Journal Database Viewing utility

Windows

The Journal Database Viewing utility provides valuable information to help in problem determination of journal-based backups.

The Journal Database Viewing utility provides the following information:

- The current state of the journal
- The file system that is 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: For backup-archive clients that are older than Version 6.3.1, you cannot view the contents of open journals with the viewing utility. An open journal is one that is opened by another process, such as the journal daemon. You can, however, view the contents of an open journal control record. The viewing utility is available with Tivoli Storage Manager V6.3.1 and newer backup-archive clients. For more information about the viewing utility, see the following technote: Run the dbviewb.exe utility in batch mode.

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

The syntax of this utility is as follows:

```
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 Windows Volume Shadow Copy Services

Windows

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

Defining VSS transient errors

Windows

The 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 by using two test flags (**TESTFLAG SETVSSMAXRETRY** and **TESTFLAG SETVSSDELAY**). The client considers the following VSS errors to be transient:

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

Defining Windows VSS test flags

Windows

The 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
```

Volume Shadow Copy Services tuning

Windows

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

Controlling the VSS diff area size

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

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

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

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

Event log maximum size

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

Gathering VSS diagnostic information for Microsoft assistance

Windows

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

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

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

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

Troubleshooting errors using a VSS trace

Windows

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

About this task

Perform the following steps to complete a VSS trace:

Procedure

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

Results

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

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

Running VSS API calls with the vsreq.exe sample program

Windows

The Volume Shadow Copy Services (VSS) Software Developers Kit (SDK) contains the **vsreq** (VSS requester) sample program. The VSS requester program performs a sequence of VSS API calls like the calls that are performed by the 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 Microsoft support to help in diagnosing the VSS problem.

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

Comparing Tivoli Storage Manager and Ntbackup.exe interaction with VSS

Windows

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

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

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

SHOW commands for the backup-archive client

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

Depending upon the information that a **SHOW** command displays, there might be instances where the information is changing or cases where it might cause the application (client, server, or storage agent) to stop running. The **SHOW** commands must be used only when development or service suggests it. The **SHOW** commands in table Table 3 are not all of the available **SHOW** commands.

Table 3. *SHOW* commands for the backup-archive client

SHOW Command	Description	Information
CACHE	Displays information about the subfile cache.	For Microsoft Windows clients that were configured to use subfile backup, this command is useful to display information about the configured subfile cache.
CLUSTER	Displays information about the disk mappings in a Microsoft Cluster.	Useful to display 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 to display information and summarize the DOMAIN, DOMAIN.IMAGE, and DOMAIN.NAS client options.
OPTIONS	Displays the client options.	Useful to determine the settings of client options.

Table 3. *SHOW* commands for the backup-archive client (continued)

SHOW Command	Description	Information
OPTTABLE	Displays information about options that are administered by the server versus the options that are managed by the client option file.	The client might receive its option settings from either the client option file or from the server. To receive the option from the server, a client option set must be defined with the DEFINECLOPTSET command. This command helps you to determine whether the client is using an option that is configured from the option file or an option that is 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 extra capabilities, such as image backup. This SHOW command displays the plug-ins that are installed for this client and it also displays 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 between a client and a server. This SHOW command reports the capabilities available by this server and client.
SYSTEMSTATE	For Windows clients, displays the SYSTEM STATE data that is available on this client.	The SYSTEMSTATE SHOW command is helpful in determining which SYSTEM STATE files are installed on this Windows system and the files that might be backed up.
TRACEFLAGS	Displays information about trace classes and aggregate trace classes for this client.	The TRACEFLAGS SHOW command is helpful in determining which trace classes and aggregate trace classes might be used for this client.
VERSION	Displays the version and build date for this client.	The VERSION SHOW command is helpful in determining which client is running and when it was built.

Resolving problems for recovery of individual Microsoft SQL databases from a virtual machine backup

Windows

You can use IBM Tivoli Storage Manager for Virtual Environments Data Protection for VMware to recover individual Microsoft SQL databases from a virtual machine backup. When you recover a database, you might have to troubleshoot common problems that occur with individual SQL databases.

When you use self-contained application protection for Microsoft SQL on Data Protection for VMware, you can back up a guest virtual machine that hosts a Microsoft SQL Server application. If you want to restore an individual Microsoft SQL database from a virtual machine backup, you must use Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server.

The following table contains solutions to common problems that you might encounter when you try to recover individual Microsoft SQL databases from a virtual machine backup.

Table 4. Troubleshooting information for recovery of individual Microsoft SQL databases from a virtual machine backup

Problem	Solution or explanation
You cannot access the database backups with Data Protection for SQL.	"Resolving database access problems"
You can see only inactive copies of SQL databases when you use the Data Protection for SQL GUI or the <code>tdpsqlc</code> command.	"Viewing active copies of Microsoft SQL databases" on page 49
You cannot view SQL database names that contain characters from the double-byte character set (DBCS) with Data Protection for SQL.	"Microsoft SQL databases with DBCS names" on page 49
You used application protection during a virtual machine backup, and you received warnings and error messages.	"Responses to messages for virtual machine backups with application protection" on page 49
You want to determine which SQL databases were on the guest virtual machine at the time of a virtual machine backup.	"Saving VSS XML manifest files" on page 50
You want to view the status of VSS writers within the guest virtual machine.	"Determining whether a virtual machine backup might fail" on page 51

Resolving database access problems

Windows

If you backed up a guest virtual machine that hosts a Microsoft SQL Server application, you might not be able to access the databases with Data Protection for SQL.

Procedure

To resolve the database access problems, complete the following steps:

1. Verify that application protection was used when you created the virtual machine backup:

- a. From the Command Prompt window, issue the following backup-archive client command to display the list of successful virtual machine backups on the server:

```
dsmc -node=datacenter_node query vm vm_name -detail
```

Where *datacenter_node* is the name of the virtual node that holds the data in the data center, and *vm_name* is the name of the virtual machine that you backed up.

- b. Verify that the output of this command contains the following output fields:
 application protection type: 'TSM VSS'
 application(s) protected: 'MS SQL 2008 – database-level recovery'

If the command output does not contain these output fields, or the second field does not include the database-level recovery text, complete the following steps:

- 1) Ensure that the Tivoli Storage Manager V7.1 or higher backup-archive client is installed on the data mover node and the client options file contains the `include.vmtsmvss vm_name` option.
 - 2) Back up the guest virtual machine again.
2. Verify that the computer name of the guest virtual machine did not change after the backup of the virtual machine was created.
 3. Verify that the in-guest DSMAGENT node has access to the data center node virtual machine backups.

- a. Issue the following command to verify that the client node has access to backup versions of the virtual machine on the server:

```
dsmc -node=datacenter_node query access
```

Where *datacenter_node* is the name of the virtual node that holds the data in the data center.

- b. Verify that the command output contains the following fields:

Type	Node	User	Path
Backup	<i>dsmagent_node</i>	*	\VMFULL- <i>vm_name</i> **

If the output does not contain this information, rerun the **set access** command on the data mover node to give the DSMAGENT node access to the guest virtual machine backups. For example, issue the following command:

```
dsmc set access backup -type=vm dsmagent_node vm_name
```

Where *dsmagent_node* is the backup-archive client node name in the virtual machine guest, and *vm_name* is the name of the virtual machine that you backed up.

What to do next

Access the individual databases with Data Protection for SQL again.

Viewing active copies of Microsoft SQL databases

Windows

To be able to view active copies of Microsoft SQL databases with Data Protection for SQL, you must run all primary and subsequent incremental backups of the SQL databases by using Data Protection for VMware with application protection.

If you used application protection for the primary backup of Microsoft SQL databases, but did not use application protection for the subsequent incremental backups, no valid active backups can be used for individual SQL database restore operations. Data Protection for SQL examines the virtual machine backup and can display only the SQL databases from the virtual machine backups that were successfully backed up with application protection.

Ensure that you enable application protection when you run the primary backup and any subsequent incremental backups of the virtual machine that hosts the Microsoft SQL application. This method ensures that the active copies of SQL databases that you backed up from a virtual machine can be displayed by Data Protection for SQL.

Microsoft SQL databases with DBCS names

Windows

While Data Protection for VMware is enabled for Unicode and can back up Microsoft SQL databases with DBCS names, Data Protection for SQL is not enabled for Unicode. Therefore, you cannot use Data Protection for SQL to restore databases with DBCS names from a virtual machine that was backed up with application protection.

To restore a virtual machine backup that contains SQL databases with DBCS names, you must restore the entire virtual machine backup with Data Protection for VMware.

Responses to messages for virtual machine backups with application protection

Windows

You might get some warnings or error messages during virtual machine backup operations when you use application protection.

The following messages might be displayed. If so, take the following actions:

ANS2196W An incompatible disk configuration is detected. Individual SQL Database Restore of database '<database_name>' is not supported.

You can use only Microsoft SQL databases that are on basic disks with Master Boot Record (MBR) partitioning for individual SQL database recovery. This warning message identifies one or more SQL databases that have an unsupported disk configuration.

ANS2330E Failed to unfreeze the VSS writers because the snapshot time exceeded the 10 second timeout limitation.

Determine whether there is an error by taking the following actions:

1. Use the vSphere client to create a quiesced virtual machine snapshot. If this action is successful, proceed to the next step.

If this action is not successful, the problem is likely to be related to VMware. You might need to contact VMware Support regarding the problem.

2. Back up the virtual machine without application protection:
 - a. Disable application protection by removing the `INCLUDE.VMTSMVSS vmname` option from the client options file.
 - b. Back up the virtual machine by running the following command from the Command Prompt window:

```
dsmc backup vm vmname -vmbackuptype=fullvm
```

Where *vmname* is the name of the virtual machine that you want to back up.

The steps that you followed so far can help you further diagnose and solve the problem. However, if this step is not successful, there is a problem with either the Windows guest virtual machine or the backup-archive client on the data mover node. You might need to find support information for Tivoli Storage Manager at the Support Portal.

Saving VSS XML manifest files

Windows

Saving VSS XML manifest files can help you determine which Microsoft SQL databases were found on the guest virtual machine at the time of backup.

About this task

VSS XML manifest files contain VSS writer information that is generated during a virtual machine backup operation. The VSS XML manifest files are required for VSS restore operations of selected Microsoft SQL databases.

Procedure

To save the VSS XML manifest files on the data mover node, complete the following steps:

1. Add the following statement to the client options file:

```
testflag VMBACKup_SAVE_LOCAL
```
2. Start a virtual machine backup with application protection of the guest virtual machine that hosts the SQL Server application.

After the virtual machine backup operation is completed, the VSS XML manifest files are saved to the following location on the data mover node:

```
C:\mnt\tsmvbackup\fullvm\vm_tsmvss\vm_name
```

Where *vm_name* is the name of the virtual machine that is backed up.

3. View the list of SQL databases that are found on the guest virtual machine at the time of backup by opening the `sqldbinfo.xml` file with a text editor. Ensure that the `sqldbinfo.xml` file contains complete information about the SQL databases that were backed up.

Determining whether a virtual machine backup might fail

Windows

Check the status of the VSS writers in a guest virtual machine to determine whether a virtual machine backup with application protection might fail.

About this task

Use the **vssadmin list writers** command to display the status of the VSS writers. This command lists all the writers that are available on the guest virtual machine, including the status of the writers. If one or more of the VSS writers are not in a stable state, then the virtual machine backup with application protection will fail.

Procedure

From the Command Prompt window, issue the following command:

```
vssadmin list writers
```

The following sample shows the command output:

```
Writer name: 'SqlServerWriter'  
  Writer Id: {a65faa63-5ea8-4ebc-9dbd-a0c4db26912a}  
  Writer Instance Id: {debc861a-7709-48b4-86a5-0a62457dc4a0}  
  State: [1] Stable  
  Last error: No error
```

The State field indicates the status of the VSS writer.

Chapter 3. Resolving Tivoli Storage Manager server problems

When working with IBM Tivoli Storage Manager, you might experience problems specific to the server. The 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:

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

Recreating the problem

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

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

Checking the server activity log file and other log files

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

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

List of additional log files

IBM Software Support might request that you send them the following log files:

- Web-server log files:
 - console.log
 - messages.log
- First-failure-data-capture (FFDC) log files:
 - exception_summary_date_time.log
 - ffdc_date_time.log

Location of log files

- The web-server log files are in the following directory:

AIX **Linux** `installation_dir/ui/Liberty/usr/servers/
guiServer/logs`

Windows `installation_dir\ui\Liberty\usr\servers\guiServer\logs`

where *installation_dir* represents the directory in which Tivoli Storage Manager is installed. For example:

AIX **Linux** `/opt/tivoli/tsm`

Windows `c:\Program Files\Tivoli\TSM`

- The FFDC log files are in the same location, but in the `ffdc` subdirectory.

Checking system error log files for device errors

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

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

The following are some examples of system error logs:

- `errpt` for AIX
- Event Log for Windows

Reverting server options or settings

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

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

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

Restarting the scheduling service

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

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

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

Resolving server space issues

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

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

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

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

Allocating additional server memory

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

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

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

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

Configuring a server instance to use shared memory

Configure a server instance to use shared memory to help resolve slow database backups that can occur because of Transmission Control Protocol (TCP) loopback problems.

About this task

In the following procedure, you must update the database backup node configuration for your server to enable shared memory.

- `AIX` `HP-UX` `Linux` `Solaris` `server_bin_directory/dbbkapi/dsm.sys`
- `Windows` `server_instance_directory\tsmdbmgr.opt`

Procedure

1. Ensure that the server options file, `dmserv.opt`, contains the following lines:

```
COMMMethod SHAREdmem
SHMPort 1510
```

2. `AIX` `HP-UX` `Linux` `Solaris` Modify the stanza for the database backup node in the client API system options file, `dsm.sys`.

- Remove the following lines from the stanza:

```
COMMMethod TCPip
TCPServeraddress 127.0.0.1
TCPport 1500
```

- Add the following lines to the stanza:

```
COMMMethod SHAREdmem
SHMPort 1510
```

3. `Windows` Modify the stanza for the database backup node in the client API system options file, `tsmdbmgr.opt`.

- Remove the following lines from the `tsmdbmgr.opt` file:

```
COMMMethod TCPip
TCPServeraddress 127.0.0.1
TCPport 1500
```

- Add the following lines to the `tsmdbmgr.opt` file:

```
COMMMethod SHAREdmem
SHMPort 1510
```

Changing the copy frequency

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

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

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

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

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

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

Resolving RELABEL operation errors

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

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

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

To resolve a RELABEL error, complete the following steps:

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

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

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

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

Avoiding communication errors during import processing

Communication errors are displayed in the activity log of the target server if you cancel the import process from the target server.

If you cancel the import process from the target server, communication error messages in the activity log show the node name that started the export operation from the source server. For example, the following messages might be displayed in the server activity log:

```
ANR0440W Protocol error on session 2 for node ADMIN
(Solaris SPARC) - invalid verb header received.
ANR3174E Communication error with managed server ADMIN.
ANR0484W Session 2 for node ADMIN (Solaris SPARC) terminated -
protocol violation detected.
```

You can ignore the communication error messages in the activity log of the target server that pertain to import processing. Alternatively, if you cancel import

processing from the source server, no communication error messages occur on either the source or target server.

Adding a self-signed certificate to the keystore

You can set up secure communications by using a self-signed certificate with your object storage system. In this situation, Tivoli Storage Manager uses HTTPS instead of HTTP when it communicates with the object storage system. The following steps provide a method for importing certificates.

About this task

Use a web browser to get a copy of the certificate used by the object storage system. The following steps are specific to Firefox, but other browsers provide similar functions. Refer to your preferred browser's instructions on exporting certificates.

Procedure

1. Get the certificate that is used by OpenStack Swift server.
 - a. Type the URL for your object storage system in the browser Address bar and press **Enter**. Use the keystone server URL for OpenStack, or the accesser node URL for Cleversafe.
 - b. Accept any warnings displayed by the browser.
 - c. Click the lock icon in the browser Address bar.
 - d. Select **More Information** in the pop-up window.
 - e. Select **View Certificate** in the Page Info window.
 - f. Click the **Details** tab in the Certificate Viewer page, and then select **Export**.
 - g. Save the exported file to the location that you want.
2. Add the certificate to the Java default keystore.

The following steps assume your client nodes are on Linux, and your server is running on Linux. Because each Cleversafe accesser has its own certificate by default, add the certificate for each accesser to the keystore, and use a different alias for each certificate.

- a. Open a terminal and change directory to the `jre/bin` directory.
The default installation location is `/opt/tivoli/tsm/jre/bin`.
- b. Make a backup copy of the Java cacerts file by running the following command: `cp ../lib/security/cacerts ../lib/security/cacerts.original`.
On a Windows system, the location of the Java cacerts keystore is: `install_dir\jre\lib\security\`, and the location of the keytool is `install_dir\jre\bin\`.
- c. Import the saved certificate from the previous procedure by running the following command: `./keytool -import -keystore ../lib/security/cacerts -alias somealias -file yourfile`
where *somealias* is a unique alias for this certificate in the keystore, which is important if you have more than one certificate, and *yourfile* is the path and file name of the certificate from the first step of these instructions.
- d. When you are asked for the password, type *changeit*. If you changed your password from the default password, type your current password.
- e. When you are asked to trust this certificate, type *yes*.

The following message is shown when the certificate is added successfully: Certificate was added to keystore. The default certificates have a short expiration. When they expire, you might lose access to the object storage until you update the certificates. You can create your own certificates and use them, but creating and installing these certificates on object storage systems is outside the scope of this document.

- f. Restart the IBM Tivoli Storage Manager server.

Determining why summary records for a client backup event are missing

If a client/server communications session ends abnormally, you might experience a delay before the summary records for a client backup event are added to the server database.

Symptom

After a client backup process is completed, the record is not added to the database immediately. It can take several hours for the summary record to be added to the database.

Causes

It can take several hours for the summary records to be added to the server database because a server session must wait for abend processing to be completed. A session can end abnormally for the following reasons:

- Network failures
- Session timeouts

Session timeouts can occur when backup processes take longer than expected.

Resolving the problem

1. To determine why client/server communication sessions ended abnormally, take the following actions:
 - a. Review the activity log by issuing the **QUERY ACTLOG** command.
 - b. Review the client error log, `dsmerror.log`, in the client installation directory.
 - c. If you cannot determine the cause of the problem by reviewing the log activity, enable tracing for the backup-archive client.
2. Resolve any communication issues. You can work with your network team to collect and analyze network data.

Related reference:

“Enabling a backup-archive client trace” on page 149

Resolving installation and upgrade problems

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

Installation log files

If errors occur during the installation process, these errors are recorded in log files.

To view installation log files, click **File > View Log** from the Installation Manager tool. To collect these log files, click **Help > Export Data for Problem Analysis** from the Installation Manager tool.

The log files are stored in the IBM Installation Manager logs directory:

AIX **HP-UX** **Linux** **Solaris** /var/ibm/InstallationManager/logs
Windows C:\ProgramData\IBM\Installation Manager\logs

Installation wizard fails to start

The IBM Installation Manager requires gtk libraries to support the graphical user interface (GUI) on AIX systems. If these libraries are not installed before you install the Tivoli Storage Manager server, the installation might not start. An error about missing gtk libraries is issued.

Related information:

 [Installing Tivoli Storage Manager by using the installation wizard](#)

Resolving GSKit installation problems

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

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

Windows Issue the following command:
`set PATH=X:\Program Files\IBM\gsk8\bin;X:\Program Files\IBM\gsk8\lib64;%PATH%`

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

HP-UX **Linux** **Solaris** Update the LD_LIBRARY_PATH or shell by issuing the following command:
`export LD_LIBRARY_PATH=platform-specific-gskit-library-directory:$LD_LIBRARY_PATH`

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

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

AIX For AIX, issue the following command:
`export LIBPATH=/usr/opt/ibm/gsk8_64/lib64:$LIBPATH`

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

- *instance_directory/sqllib/usercshrc*
- *instance_directory/sqllib/userprofile*

For the *instance_directory/sqllib/usercshrc* file, add the following lines:

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

For the *instance_directory/sqllib/userprofile* file, add the following lines:

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

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

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

If your GSKit version is not 8.0.14.28 or later, you must reinstall the server. The reinstallation ensures that the correct GSKit version is available.

Server instances are not created during an upgrade

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

About this task

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

- **AIX** **HP-UX** **Linux** **Solaris** Secure shell (SSH)
- **Windows** Windows server message block (SMB)

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

upgrade the server instance.

AIX

HP-UX

Linux

Solaris

Procedure

1. Close the installation wizard.
2. Complete the following steps for each server instance:
 - a. After the upgrade finishes, issue the following command to recreate the instance:

```
/opt/tivoli/tsm/db2/instance/db2icrt -u instance_user instance_name
```
 - b. Recreate the variables in the instance file. Issue the **db2set -i** command for each variable in the instance file. For example, set the variable *DB2COMM* to be only TCPIP for instance MYINST:

```
/opt/tivoli/tsm/db2/instance/db2set -i MYINST DB2COMM=TCPIP
```

To display a list of all defined variables specify the **-all** parameter, for example **db2set -all**.
 - c. Issue the **db2stop** command to stop the database instance.
 - d. Use the user ID that owns the server instance to issue the **db2start** command to start the database instance.
 - e. Catalog and upgrade each database by issuing the following commands:

```
db2 catalog db TSMDB1 on "database_path"
db2 upgrade db TSMDB1
```
 - f. Issue the **db2stop** command.
 - g. Start the server.

Resolving a stopped uninstallation process

An expired DB2 instance user password might cause the IBM Tivoli Storage Manager uninstallation process to stop before it completes.

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

Client automatic deployment did not upgrade the client software

If the deployment schedule completes but the client software is not upgraded to the target level, review the log files on the client system.

Symptom

The client software was not upgraded to the target level after the automatic deployment schedule completed.

Cause

The following examples are some of the reasons why the client software does not upgrade to the target level.

- There is not enough space on the client file system to complete the upgrade.
- Network issues prevented the data from the server to be transferred to the client system.

Resolving the problem

You can resolve the client upgrade failure by reviewing the log and trace files on the client system. The deployment manager writes log and trace data for a deployment operation to the client file system. The location for the log files is specified in the deployment schedule definition on the server.

Complete the following steps to resolve the client upgrade failure:

1. Change directories to the location of the log files.

- **AIX** The default location is `/usr/tivoli/client/IBM_ANR_UNX/Vxxxx/log`.
- **HP-UX** **Linux** **Solaris** The default location is `/opt/tivoli/tsm/client/IBM_ANR_UNX/Vxxxx/log`.
- **Mac OS X** The default location is `/Library/Application Support/tivoli/tsm/client/ba/bin/IBM_ANR_MAC/Vxxxx/log`.
- **Windows** The default location is `C:\Program Files\Tivoli\TSM\IBM_ANR_WIN\Vxxxx\log`.

Where the `Vxxxx` directory in the path represents the target level of the deployed backup-archive client.

2. Review the log and trace files from the deployment manager to determine the root cause of the client upgrade failure. Table 5 shows a list of log files that you can review.

Table 5. Description of log files

Log name	Description
setup.log	Error log that shows error, warning, and informational messages.
trace.txt	Client trace that shows detailed information about the client upgrade process.
updatemgr.log	Deployment manager log that shows information about the deployment process.

Resolving server stoppages

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

The server might stop for one of the following reasons:

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

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

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

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

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

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

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

Related tasks:

Appendix B, “Running the `tsmdiag` utility,” on page 197

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 can be the entire 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 can affect the entire server.

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

- For a stoppage or a loop where the server can respond to commands, issue the following commands to help determine the cause of the stoppage:
 - **QUERY SESSION f=d**
 - **QUERY PROCESS**
 - **SHOW RESQ**
 - **SHOW THREADS**
 - **SHOW DEADLOCK**
 - **SHOW TXNT**
 - **SHOW DBTXNT**
 - **SHOW LOCKS**
 - **SHOW LIBR**
 - **SHOW MP**
 - **SHOW SESS**
 - **SHOW ASQ**
 - **SHOW ASVOL**
 - **SHOW DBV**
 - **SHOW SSS**
 - **SHOW CSV** (Issue this command only when the problem is related to scheduling.)

- When a server hangs or loops, issue the following commands to provide a detailed diagnostic snapshot of the Tivoli Storage Manager environment:

```
db2fodc -hang -all dbs
db2support . -d database -s
```

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

- In addition to the output from the listed commands, or in the cases of a server that cannot respond to commands, collect a dump. The way that you collect a dump depends on the operating system.

- **AIX** **HP-UX** **Linux** **Solaris** Issue the **KILL -11** command on the dsmserv process to create a core file. To run the “kill” command, obtain the process ID by issuing the **PS** command.

- **Windows** Search for collecting user mode dumps at the Microsoft website at <http://support.microsoft.com/>.

Resolving wait state problems with external user repository servers

AIX **Linux** **Solaris**

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

Before you begin

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

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

About this task

Complete the following steps to resolve a wait state problem that occurs with the following servers when you are using an LDAP server:

Procedure

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

Solaris Issue the following commands:

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

Linux Issue the following commands:

- a. db2set DB2_ALTERNATE_GROUP_LOOKUP=GETGROUPLIST
 - b. db2stop force
 - c. db2start
3. Restart the 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.

Before you begin

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

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

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

About this task

Perform the following steps to capture the server error file:

Procedure

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

Finding the system image (core file)

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

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

- **AIX** **HP-UX** **Linux** **Solaris** A file is typically created called core. Be sure that there is enough space in the server directory to accommodate a dump operation. It is common to have a dump file as large as 2 GB for the 32-bit Tivoli Storage Manager server. Additionally, make sure the ulimit for core files is set to unlimited to prevent the dump file from being truncated.

- **Windows** The contents of the system are dumped automatically through a system application programming interface (API) call. If the server is running as a service, the dump file is called `dsmsvc.dmp`. Otherwise, the dump file is called `dsmserv.dmp`

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

Retrieving library files for core analysis

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

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

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

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

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

- **AIX** Collect the following files:
 - `/usr/ccs/lib/libpthread.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 dsmserv core_file`** command. Then, from the `dbx` prompt, issue the **`map`** command to show all of the libraries that are loaded and needed for core analysis.
- **HP-UX** Issue the **`chatr dsmserv`** 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** Issue the **`ldd dsmserv`** command and send in all the dynamic shared libraries. For example:
 - `libm.so.6 =>/lib64/libm.so.6`
 - `libnsl.so.1 =>/lib64/libnsl.so.1`
 - `libpthread.so.0 =>/lib64/libpthread.so.0`
 - `libdl.so.2 =>/lib64/libdl.so.2`

- libc.so.6 =>/lib64/libc.so.6
- /lib64/ld64.so.1 =>/lib64/ld64.so.1
- **Solaris** Issue the following commands to collect the needed libraries:
 - **sh**
 - **cd /usr**
 - **find . -name "ld.so" -print ; **
 - **find . -name "ld.so.?" -print ; **
 - **find . -name "libm.so.?" -print ; **
 - **find . -name "libsocket.so.?" -print ; **
 - **find . -name "libnsl.so.?" -print ; **
 - **find . -name "libthread.so.?" -print ; **
 - **find . -name "libthread_db.so.?" -print ; **
 - **find . -name "libdl.so.?" -print ; **
 - **find . -name "libw.so.?" -print ; **
 - **find . -name "libgen.so.?" -print ; **
 - **find . -name "libCrun.so.?" -print ; **
 - **find . -name "libc.so.?" -print ; **
 - **find . -name "libmp.so.?" -print ; **
 - **find . -name "libc_psr.so.?" -print ; **
 - **find . -name "librtld_db.so.?" -print) > runliblist**
 - **tar cfh runliblist.tar -I runliblist**

Retrieving system log files

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

Retrieve the following log files to give to IBM service:

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

Retrieving the activity log

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

View and collect the activity log entries that start at least two hours prior to the stoppage and 30 minutes after the stoppage by issuing the **QUERY ACTLOG** command.

Detecting errors after a server service starts and stops

Windows

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

About this task

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

Procedure

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

sqllib/db2dump directory causes shutdown

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

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

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

- Add space to the file system.
- Move the files to another file system. See Table 6 on page 70.
- Use the server to archive the files, and then delete them by using the following steps:
 1. Run the `db2support` utility to collect the DB2 system diagnostic information.
 2. Archive the `db2support.zip` file and diagnostic files that are listed in Table 6 on page 70 to server with the client.
 3. Delete the files that are listed in Table 6 on page 70.

Table 6. Files that can be deleted after these files are archived

File name	Description
instance_name.nfy instance_name.n.nfy (where <i>n</i> is a number)	Administration notification logs
db2dasdiag.log	DB2 administration server (DAS) diagnostic log
db2eventlog.xxx (where <i>xxx</i> is the database partition number)	DB2 event log
nnnnnnn.nnnnnn.nnn.dump.bin (where <i>n</i> is a number)	Binary dump files of key in-memory structures
nnnnnnn.n.nnn.trap.txt (where <i>n</i> is a number)	Trap files
nnnnnnn.nnnnnn.nnn.apm.bin (where <i>n</i> is a number)	Access the plan-manager binary dump files
nnnnnnn.nnnnnn.nnn.stack.txt (where <i>n</i> is a number)	Stack traces
F0DC_xxxx/core<pid>	Core files These F0DC_xxx directories contain the time stamp in the directory name. Keep the most recent directories and their files. The history can be useful for diagnosing possible future problems that are related to the database. A guideline is to keep at least 1 weeks worth.
events/db2optstats.n.log (where <i>n</i> is a number)	Statistics log file

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

Related reference:

“Locating DB2 diagnostic log files” on page 74

Resolving issues with database page verification

Page validation failure during database backup processing can indicate corruption in the database, requiring a repair action to correct the problem. If page validation fails, the database backup also fails.

Procedure

- Contact IBM support for assistance in diagnosing and repairing any database corruption.
- If a full database backup was in progress to free space in the archive log directory, take one of the following actions:
 - Increase the amount of space in the archive log directory.
 - Issue the ARCHFAILOVERLOGDIRECTORY option to specify a failover archive log directory in which the server can store log files that cannot be stored in the archive log directory.

Ensuring adequate space in the archive log directory allows the server to continue to run while the database is repaired.

Resolving database errors

Database errors might be caused by issues such running out of space and by errors that are caused by insert, update, or delete operations.

Users who are experienced DB2 administrators can run advanced SQL queries and use DB2 tools to monitor the database, the space that is being used, and any errors. When you are running these queries, do not use DB2 tools to change DB2 configuration settings from those settings that are preset by Tivoli Storage Manager or, do not use any other software to modify these settings. The server must be used with the data definition language and database configuration that Tivoli Storage Manager deploys.

HP-UX **Solaris** Tivoli Storage Manager installs the DB2 database. After you install and configure Tivoli Storage Manager, check that the system kernel parameters are set properly. To check these parameters, run the **db2osconf** utility.

For more information, see the DB2 product information.

Related information:

 [DB2 V10.5 product information](#)

Resolving database manager starting problems

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

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

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

You might also receive this error:

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

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

An example from the db2diag.log:

```
2011-04-26-16.04.11.820963-420 I2345542E1168 LEVEL: Error
PID : 25178 TID : 47207843621184PROC : db2sysc 0
INSTANCE: hannigan NODE : 000
EDUID : 1 EDUNAME: db2sysc 0
FUNCTION: DB2 Common, OSSe, OSSHLlibrary::load, probe:80
MESSAGE : ECF=0x90000076=-1879048074=ECF_LIB_CANNOT_LOAD
          Cannot load the specified library
DATA #1 : Hex integer, 4 bytes
0x00000002
DATA #2 : String, 58 bytes
/home/hannigan/sqllib/security64/plugin/server/dsmdb2pw.so
CALLSTCK:
[0] 0x00002AEF63DD267E pdOSSeLoggingCallback + 0x20C
[1] 0x00002AEF68486A42 /home/hannigan/sqllib/lib64/libdb2osse.so.1 + 0x1C4A42
[2] 0x00002AEF6848825E ossLog + 0xA6
[3] 0x00002AEF684928E9 _ZN11OSSHLlibrary4loadEPKcm + 0x1D3
[4] 0x00002AEF63F63BDC _Z20secLoadPluginGenericP19SEC_PLUGIN_HANDLE_TPc + 0x68
```

```

[5] 0x00002AEF63F62FBB _Z23secLoadServerAuthPluginP19SEC_PLUGIN_HANDLE + 0x57
[6] 0x00002AEF63F6C833 _Z25sqllexLoadAllPluginsServerP5sqlca + 0x3B5
[7] 0x00002AEF6431737C /home/hannigan/sqlllib/lib64/libdb2e.so.1 + 0x123637C
[8] 0x00002AEF643164C5 sqloRunInstance + 0x191
[9] 0x00000000040D31D DB2main + 0xD41

2011-04-26-16.04.11.825930-420 I2346711E1178 LEVEL: Error
PID : 25178 TID : 47207843621184PROC : db2sysc 0
INSTANCE: hannigan NODE : 000
EDUID : 1 EDUNAME: db2sysc 0
FUNCTION: DB2 Common, OSSe, OSSHLibrary::load, probe:90
MESSAGE : ECF=0x90000076=-1879048074=ECF_LIB_CANNOT_LOAD
          Cannot load the specified library
DATA #1 : String, 109 bytes
../shared/gskit8/lib/linux64_x86/libgsk8iccs_64.so: cannot open shared object
file: No such file or directory
CALLSTCK:
[0] 0x00002AEF63DD267E pdOSSeLoggingCallback + 0x20C
[1] 0x00002AEF68486A42 /home/hannigan/sqlllib/lib64/libdb2osse.so.1 + 0x1C4A42
[2] 0x00002AEF6848825E ossLog + 0xA6
[3] 0x00002AEF6849294D _ZN11OSSHLibrary4loadEPKcm + 0x237
[4] 0x00002AEF63F63BDC _Z20secLoadPluginGenericP19SEC_PLUGIN_HANDLE_Tpc + 0x68
[5] 0x00002AEF63F62FBB _Z23secLoadServerAuthPluginP19SEC_PLUGIN_HANDLE + 0x57
[6] 0x00002AEF63F6C833 _Z25sqllexLoadAllPluginsServerP5sqlca + 0x3B5
[7] 0x00002AEF6431737C /home/hannigan/sqlllib/lib64/libdb2e.so.1 + 0x123637C
[8] 0x00002AEF643164C5 sqloRunInstance + 0x191
[9] 0x00000000040D31D DB2main + 0xD41

```

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

```

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

```

Tracing the user ID and Password plug-in

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

To set up automatic tracing for the user ID and password plug-in, complete the following steps:

AIX HP-UX Linux Solaris

1. Ensure that the server has write authority to the `~/sqlllib/db2dump/` directory.
2. Add the following text to the `~instance/sqlllib/userprofile` file:
`export DB2_DS MDB2PW_TRACEFILE=filename`

where *filename* is a fully qualified path and file name of the trace file, for example `~/sqlllib/db2dump/dsmdb2pw.trc`.

3. Restart DB2

After DB2 restarts, trace output is stored in the specified file and directory.

Windows

1. To verify that the `DB2_VENDOR_INI` `db2set` is set, run the `db2set` command.
2. If the `DB2_VENDOR_INI` variable is not set, create a configuration file, for example:
`c:\Program Files\Tivoli\TSM\s1\tsmdbmgr.env`

3. Update the configuration file that is listed in the *DB2_VENDOR_INI* with the location of the trace file:

```
DB2_DS MDB2PW_TRACEFILE=c:\Program Files\Tivoli\TSM\s1\sql1lib\dsmdb2pw.trc
```

4. Set up the trace file by issuing the following command:

```
db2set -i server_instance DB2_VENDOR_INI=configuration_file_location
```

for example:

```
db2set -i s1 DB2_VENDOR_INI=c:\Program Files\Tivoli\TSM\s1\tsmdbmgr.env
```

5. Stop the Tivoli Storage Manager server and restart it by issuing the following commands:

```
halt
```

```
dsmserv -k server_instance
```

After the server restarts, trace output is stored in the specified file and directory.

Tip: You can use the file name and directory of your choice for the name and location of the trace file.

Limiting DB2 memory allocation

When DB2 uses a large amount of memory, you can limit the amount of memory that DB2 uses by issuing the **db2 update** command.

About this task

By default, DB2 is installed and configured to use automatic memory management, which causes DB2 to use a large percentage of the physical memory. To restrict the amount of memory, use the **db2 update** command to specify the memory limit.

Procedure

Issue the **db2 update** command:

```
db2 update dbm cfg using instance_memory memory_value
```

where *memory_value* is specified in 4 KB blocks.

Example

To limit the DB2 memory allocation to use 3,200,000 KB of memory, divide 3,200,000 KB by 4 KB blocks, which gives a result of 800000. Then, issue the following command:

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

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

Retrieving DB2 version information

The version of DB2 that is installed with the IBM Tivoli Storage Manager server is updated periodically. If database problems occur, you must know the version of DB2 and its location so that you can provide this information to IBM Software Support.

Procedure

Issue the **db2level** command to show where DB2 products are installed on your server, and to list the DB2 product level.

The following output shows sample results of the **db2level** command.

```
AIX HP-UX Linux Solaris
> db2level
DB21085I This instance or install (instance name, where applicable:
"cetinst1") uses "64" bits and DB2 code release "SQL10051" with level
identifier "0602010E".
Informational tokens are "DB2 v10.5.0.1", "special_31150",
"IP23526_31150", and Fix Pack "1".
Product is installed at "/opt/tivoli/tsm/db2".
```

```
Windows
C:\>db2level
DB21085I This instance or install (instance name, where applicable: "SERVER1")
uses "64" bits and DB2 code release "SQL10051" with level identifier
"0602010E".
Informational tokens are "DB2 v10.5.100.64", "special_31150",
"IP23521_31155", and Fix Pack "1".
Product is installed at "C:\PROGRA~1\Tivoli\TSM\db2" with DB2 Copy Name
"DB2TSM1".
```

Locating DB2 diagnostic log files

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

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

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

1. Log in as the server user instance.
2. Issue the following command:

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

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

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

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

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

3. If no path is specified in the **DIAGPATH** configuration parameter, the DB2INSTPROF directory path is used. Find the path that was set in the *DB2INSTPROF* environment variable. Issue the following command from the DB2 command-line prompt:

```
db2set db2instprof
```

The output from this command shows the location of DB2 data files. The diagnostic log file is in the instance sub directory of the directory that is specified by the *DB2INSTPROF* registry variable. For example, for the server instance TSMSEVER1, the **db2set db2instprof** command shows this path:

```
C:\ProgramData\IBM\DB2\DB2TSM1
```

The diagnostic log file is in the TSMSEVER1 sub directory:

```
C:\ProgramData\IBM\DB2\DB2TSM1\TSMSEVER1
```

4. If the *DB2INSTPROF* environment variable is not set, then *x:\SQLLIB\DB2INSTANCE* is used. *x:\SQLLIB* is the drive reference and it is also the directory that is specified in the *DB2PATH* registry variable. The value of *DB2INSTANCE* is the name of the instance. You do not need to call the directory *SQLLIB*. The first part of the output from the **db2set db2path** command is the diagnostic data directory path with the instance name added. The output shows the following directory path:

```
C:\Program Files\Tivoli\TSM\db2\TSMINST1
```

where *DB2PATH* is *C:\Program Files\Tivoli\TSM\db2* and the instance name is *TSMINST1*.

Related reference:

“Installation log files” on page 60

DB2 upgrade log files

When you upgrade the Tivoli Storage Manager server, the DB2 **db2ckupgrade** script runs and creates log files for the server databases.

During the upgrade, the wizard automatically fixes some errors in the database. You must fix other errors manually. Check the log files for the errors that you must fix. The log files contain the results from the **db2ckupgrade** command for each database.

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

If you receive a database error message when the script is running that the wizard does not fix, you must cancel or close the wizard, fix the error, and start the upgrade again. If it is a silent installation that is being completed, you must check the *log.text* file for errors, fix any errors in this file, and start the upgrade again. For details about error messages that are listed in the log files, see the DB2 product information.

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 V6.1 databases. You must start the Tivoli Storage Manager V6.2 server to generate a new backup image in DB2. After Tivoli Storage Manager V6.2 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 V6.2 server for incremental database backups or database restores to be successful.

If you started the upgraded Tivoli Storage Manager V6.2 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 V6.2. If you drop the database before a backup image is generated, you must reinstall Tivoli Storage Manager V6.1 server and then restore the database.

If you must restore a Tivoli Storage Manager V6.1 database and the database does not exist, you must restore it through Tivoli Storage Manager V6.1. You can then upgrade to Tivoli Storage Manager V6.2.

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

When a database is restored, the database ID file must stay in sync with the database. With Tivoli Storage Manager V6.2, if you format the database before you restore it, the database ID file changes. This change causes a mismatch of the date and time in the database and keeps the server from starting.

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

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

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

Resolving problems with the BACKUP DB and the RESTORE DB commands

The 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 server through the client application programming interface (API).

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

1. Open a DB2 command-line interface:

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

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

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

```
? sql2033
```

You can use the details of the error condition to debug the problem with the **BACKUP DB** or **RESTORE DB** command. If the SQLERRMC code is also displayed, it is explained in the SQLCODE description that you are provided. You can find more information about the API return codes through the following files:

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

Resolving incorrect environment variables for BACKUP DB and RESTORE DB

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

About this task

Requirement:

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

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 c:\tsminst1\tsmdbmgr.env. You can verify the file name and location by issuing the db2set -i tsminst1 DB2_VENDOR_INI command, where tsminst1 is the DB2 instance.

The DSMI_* variables are initially set up automatically by the Tivoli Storage Manager instance configuration wizard.

Procedure

Open the `/home/tsminst1/sql1lib/userprofile` file and review the statements. If you change this file, stop and restart the DB2 instance so that the changes are included. For example, consider the following scenario. The `userprofile` file has statements like the following example text:

```
export DSMI_CONFIG=server_instance_directory/tsmdbmgr.opt
export DSMI_DIR=server_bin_directory/dbbkapi
export DSMI_LOG=server_instance_directory
```

The `tsmdbmgr.opt` file has the following text:

```
SERVERNAME TSMDBMGR_TSMINST1
```

The `server_bin_directory/dbbkapi/dsm.sys` file has the following text:

```
SERVERNAME TSMDBMGR_TSMINST1
commethod tcpip
tcpserveraddr localhost
errorlogname /tsminst1/tsmdbmgr.log
```

Verify that the `SERVERNAME` entry in the `tsmdbmgr.opt` file matches the `SERVERNAME` entry in the `dsm.sys` file.

Linux Do not add the `PASSWORDACCESS generate` option to the `dsm.sys` configuration file. This option can cause the database backup to fail.

Resolving error message ANR2968E

Error message ANR2968E is surfaced during the **BACKUP DB** command.

About this task

There are two causes for this error message:

- If the Tivoli Storage Manager error log file is owned by the root user ID rather than the server instance user ID.
- **Windows** If you use quotation marks to surround the paths that are in the `tsmdbmgr.env` file. Use a path that does not contain spaces or use the Windows short name for the path.

To correct the error caused by the root user ID, complete the following steps:

Procedure

1. Log on using a Tivoli Storage Manager server instance ID and verify the name of the error log file. For example:

```
$ grep -i "ERRORLOGNAME" $DSMI_DIR/dsm.sys
ERRORLOGNAME /home/db2inst1/tsminst1/tsmdbmgr.log
```

where `db2inst1` is the server instance user ID and `/home/db2inst1/tsminst1/` is the server instance directory.

2. Issue the following example command to verify the current owner of the error log file:

```
$ ls -la /home/db2inst1/tsminst1/tsmdbmgr.log
-rw-r--r-- 1 root system 834 May 05 09:43 /home/db2inst1/tsminst1/tsmdbmgr.log
```

3. If the error log file is not owned by the Tivoli Storage Manager instance user ID, remove it. You must have root authority to remove the file. Issue the following example command to remove the log file:

```
$ su root password
# rm /home/db2inst1/tsminst1/tsmdbmgr.log
# exit
```

4. Issue the **BACKUP DB** command and verify that the command completed successfully. Verify that the log file is owned by the server instance ID. For example:

```
$ ls -la /home/db2inst1/tsminst1/tsmdbmgr.log
-rw-r--r-- 1 db2inst1 db2iadm1 834 May 05 09:50
/home/db2inst1/tsminst1/tsmdbmgr.log
```

Troubleshooting error message ANR2971E using the SQL code

Error message ANR2971E might show when you are restoring or backing up a database operation, and the process stops. Use the SQL code attached to the error to help you resolve this problem.

Before you begin

If you are restoring a database because the server stopped during normal operation, review the db2diag.log file *before* backing up or restoring the database.

The following message can be issued when you are restoring or backing up data:
ANR2971E Database backup/restore/rollforward terminated - DB2 sqlcode -2581 error

In the following scenario, the **DSMSERV RESTORE DB** process failed with a DB2 SQL 2581 code. This following scenario does not pertain to problems with the DSMI environment variables.

Procedure

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   : B1JRPM01
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=SQLQ_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 you issue the **BACKUP DB** or **RESTORE DB** commands.

ANR2968E - Database backup terminated DB2 SQLCODE -2033 SQLERRMC 406

To resolve the SQL 406 error message, ensure 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 ~/sqllib/userprofile and ~/sqllib/usercshrc.

DB2 SQLCODE: -2033, DB2 SQLERRMC: 106

If you receive the SQL 106 error message, it can mean that there is a permissions problem with the log file that the application programming interface (API) for the client writes.

To resolve the error, find the log file with the permissions problem, log in using the root user ID, and delete the file.

DB2 SQLCODE: -2033, DB2 SQLERRMC: 168

Verify that the DSMI_DIR environment variable points to the client API executable directory that contains the trusted communication agent dsmtca.

ANR2971E - Database backup/restore/rollforward terminated DB2 SQLCODE - 2071 error

The library cannot be loaded because the library or a library that is required by this library does not exist or does not have a valid format. If a library cannot be loaded, it means that a 32-bit library is being loaded in a 64-bit instance, or a 64-bit

library is being loaded in 32-bit instance. If a library cannot be loaded, it also indicates that the DSMI_DIR environment variable points to the wrong Tivoli Storage Manager client API executable files. To obtain information about the error, open a DB2 command line processor window and issue the following command:

```
db2 => ? sq12071
```

Verify that if any changes were made to the `tsmdbmgr.opt`, `dsm.sys`, or `sql1lib/userprofile` files, the DB2 instance is recycled so that it picks up the new values. To recycle the DB2 instance, stop and start the Tivoli Storage Manager server. Also, verify that the **EXPORT** command precedes the `DSMI_*` entries in the `sql1lib/userprofile` file.

Error message indicates that the node is locked

You might receive an error when DB2 contacts the server and a particular node, and receives an error message that states that the node is locked.

To correct the error, use the localhost address instead of an explicit loopback address, for example 127.0.0.1. Specify the `tcpserveraddress localhost` option in the `SERVERNAME TSMDBMGR_TSMINST1` stanza of the `dsm.sys` file.

Problems with database backup performance

In some cases, particularly on AIX systems, if the server is configured to use TCP/IP for database backup and restore operations, database backups might be slow. To resolve the problem, configure the server instance to use shared memory.

Related tasks:

“Configuring a server instance to use shared memory” on page 56

Characteristics of the `$$_TSMDBMGR_$$` user ID

The Tivoli Storage Manager server generates the `$$_TSMDBMGR_$$` user ID at startup.

You can view the `$$_TSMDBMGR_$$` user ID in the results of a **QUERY SESSION** command. This ID is also present in the activity log file and other server log files.

The server uses the `$$_TSMDBMGR_$$` user ID to back up the server database. By using the `$$_TSMDBMGR_$$` user ID, you can make the database accessible for processing if the server is unavailable. Changing this ID jeopardizes the ability to recover or restore a server if a disaster occurs.

Restriction: You cannot change the `dsm.sys` or `dsm.opt` file to set up or use a different client node name. The local Tivoli Storage Manager server database uses the `dsm.sys` or `dsm.opt` file to back up its own database.

Resolving database reorganization problems

Database table reorganization and index reorganization require a significant amount of system resources. To avoid occupying system resources that can be used elsewhere, run your reorganization routines on off times.

Unexpected database growth and unexpected active and archive log space requirements can occur if tables or the indexes associated with tables are not reorganized. Tivoli Storage Manager reorganizes tables by default. If automatic reorganization is affecting server performance, you can schedule reorganization manually.

The following suggestions might help when you configure your reorganization:

- Turn on index reorganization if you are running deduplication on your server. See the server option `ALLOWREORGINDEX`.
- By default, table reorganization is turned on 24-hours-a-day. Run reorganization during a relatively idle time during the day. See the following server options for defining an idle time when reorganization can run:
 - `REORGBEGINTIME`
 - `REORGDURATION`

Analyzing the process symptoms to resolve problems

You can sometimes determine the cause of errors by observing the process symptoms.

You might encounter one of the following process symptoms:

- Insufficient space in a target copy storage pool
- Damaged file found on volume
- Files are not expired after reducing the number of versions that need to be kept
- Migration does not run for sequential media storage pool
- Migration only uses one process
- Process running slow

Reviewing process messages to determine the state of server operations

Server processes, whether run in the foreground or background, will always issue a “process started” message and a “process ended” message in addition to the general process messages. You can use these messages to determine the status of your server operation.

Processes that run on the server

A server process is a task that is performed on the server. You can assign the task to perform a specific operation, such as migrating data from a storage pool to the next storage pool in the hierarchy. Issue the server processes to resolve problems that you are having with your server.

Server processes are typically initiated as an automated process on the server. The process might or might not be influenced by a server option or other setting. The server process can also be started by a command.

Many server processes can be run in the `BACKGROUND` or asynchronously. Processes that run in the `BACKGROUND` 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 7 for the descriptions of the server processes.

Table 7. 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 7. 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:</p> <ul style="list-style-type: none"> • IMPORT ADMIN • IMPORT NODE • IMPORT POLICY • IMPORT SERVER 	
LABEL LIBVOLUME	Label one or more library volumes in a library.	
Migration	<p>Migrate data from one storage pool to the next in the storage hierarchy.</p> <p>Migration starts and stops, based on the HighMig and LowMig thresholds defined for the storage pool. Whenever UPDATE STGPOOL is issued, these values are reexamined and, if appropriate, MIGRATION is started. Otherwise, the server monitors the percentage utilization for non-migrated data in a storage pool. As the server needs, it starts migration processing for that storage pool when the HighMig threshold is exceeded. You can also issue the MIGRATE STGPOOL command to manually start migration processing.</p>	Migration might be configured to run multiple concurrent processes. The multiple processes are controlled by the 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 7. Server processes (continued)

Process or command	Description	Runs in the foreground or as a multiple process
Reclamation	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. 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.	
RESTORE STGPOOL	Restore all files for a given storage pool from a copy storage pool.	The RESTORE STGPOOL can run as a synchronous process by specifying WAIT=YES . RESTORE STGPOOL might be run using multiple concurrent processes, which is controlled by the MAXPROCESS parameter specified on the RESTORE STGPOOL command.
RESTORE VOLUME	Restore all files for a given volume from a copy storage pool.	The RESTORE VOLUME command can run as a synchronous process by specifying WAIT=YES . RESTORE VOLUME might be run using multiple concurrent processes, which is controlled by the MAXPROCESS parameter specified on the RESTORE VOLUME command.

Messages issued when processes start

When the server runs tasks as processes, the processes are assigned an identification message and report that they have started.

The reported start is issued in the following message:

ANR0984I Process *process_id* for *process_name* started in the *process_state* at *time*

The following list defines the variables from this message:

process_id

Numeric process identifier.

process_name

The name of the process.

process_state

BACKGROUND or **BACKGROUND**. If the process is running in the foreground, the command was issued with the **WAIT=YES** parameter. Foreground processing causes the administrative session that issued the command to wait until the processing completes. A process running in the background returns

immediately to the administrative session that issued the command, indicating that a process was started while the process still runs. Processes running in the background might be monitored with the **QUERY PROCESS** command.

time The time that the process was started.

Messages issued when processes end

When the server runs tasks as processes, the processes report when they end. The “process ended” messages that are issued vary from process to process. The message depends on whether the process must report about items and bytes processed, no items or bytes processed, items processed, or just bytes processed.

Process ended

When a process completes and it does not have bytes or number of files to report, the following message is issued:

```
ANR0985I Process process_id for process_name running in the process_state completed with the completion_state at time
```

The following list defines the variables from this message:

process_id

Numeric process identifier.

process_name

The name of the process.

process_state

FOREGROUND or **BACKGROUND**. If the process is running in the foreground, the command was issued with the **WAIT=YES** parameter. Foreground processing causes the administrative session that issued the command to wait until the processing completes. A process running in the background returns immediately to the administrative session that issued the command, indicating that a process was started while the process still runs. Processes running in the background can be monitored with the **QUERY PROCESS** command.

completion_state

SUCCESS or **FAILURE**.

time The time that the process was started.

Process ended with items and bytes

When a process completes and has bytes and items processed to report, the following message is issued:

```
ANR0986I Process process_id for process_name running in the process_state processed number_of_items items for a total of bytes_processed bytes with a completion state completion_state at time
```

The following list defines the variables from this message:

process_id

Numeric process identifier.

process_name

The name of the process.

process_state

FOREGROUND or **BACKGROUND**. If the process is running in the foreground, the

command was issued with the **WAIT=YES** parameter. Foreground processing causes the administrative session that issued the command to wait until the processing completes. A process running in the background returns immediately to the administrative session that issued the command, indicating that a process was started while the process still runs. Processes running in the background can be monitored with the **QUERY PROCESS** command.

number_of_items

The number of items processed.

bytes_processed

The number of bytes processed.

completion_state

SUCCESS or FAILURE.

time The time that the process was started.

Process ended with items

When a process completes and has items processed to report, the following message is issued:

```
ANR0987I Process process_id for process_name running in the process_state
processed number_of_items items with a completion state completion_state at time
```

The following list defines the variables from this message:

process_id

Numeric process identifier.

process_name

The name of the process.

process_state

FOREGROUND or **BACKGROUND**. If the process is running in the foreground, the command was issued with the **WAIT=YES** parameter. Foreground processing causes the administrative session that issued the command to wait until the processing completes. A process running in the background returns immediately to the administrative session that issued the command, indicating that a process was started while the process still runs. Processes running in the background can be monitored with the **QUERY PROCESS** command.

completion_state

SUCCESS or FAILURE.

time The time that the process was started.

Process ended with bytes

When a process completes and has bytes processed to report, the following message is issued:

```
ANR0988I Process process_id for process_name running in the process_state
processed bytes_processed bytes with a completion state completion_state at time
```

The following list defines the variables from this message:

process_id

Numeric process identifier.

process_name

The name of the process.

process_state

BACKGROUND or **BACKGROUND**. If the process is running in the foreground, the command was issued with the **WAIT=YES** parameter. Foreground processing causes the administrative session that issued the command to wait until the processing completes. A process running in the background returns immediately to the administrative session that issued the command, indicating that a process was started while the process still runs. Processes running in the background can be monitored with the **QUERY PROCESS** command.

bytes_processed

The number of bytes processed.

completion_state

SUCCESS or **FAILURE**.

time

The time that the process was started.

Analyzing the ANR1221E error message

When you receive error message ANR1221E, the cause is typically due to insufficient space in the target copy storage pool.

About this task

Perform the following steps to resolve error message ANR1221E:

Procedure

1. Issue the **QUERY STGPOOL** *stgpoolName* **F=D** command.
2. Issue the following SQL select statement from an administrative client to this server: "select *stgpool_name,devclass_name,count(*)* as 'VOLUMES' from volumes group by *stgpool_name,devclass_name*."
3. Compare the number of volumes reported by the select statement to the maximum scratch volumes allowed (as reported by the **QUERY STGPOOL** command). If the number of volumes reported by the **select** is equal to or exceeds the "Maximum Scratch Volumes Allowed," update the storage pool and allow more scratch volumes. If scratch volumes are not used in the storage pool (*scratch=0*), then ensure that you add more private volumes. Issue the **UPDATE STGPOOL** *stgpoolName* **MAXSCR=nn** command, where *stgpoolName* is the name of the storage pool to update and *nn* is the increased number of scratch volumes to make available to this copy storage pool.

Important: The tape library should have this additional number of scratch volumes available, or you need to add scratch volumes to the library prior to issuing this command and retrying the **BACKUP STGPOOL** operation.

Analyzing the ANR2317W error message

The ANR2317W error message is issued when a process determines that there is a damaged file.

About this task

The message is shown with the following information:

```
ANR2317W Audit Volume found damaged file on volume volumeName: Node nodeName,  
Type fileType, File space fileSpaceName, fsId fileSpaceID,  
File name fileName is number version of totalVersions versions.
```

Perform the following steps to resolve error message ANR2317W:

Procedure

1. Issue the **QUERY VOLUME** *volumeName* **F=D** command.
2. Issue the following SQL select statement from an administrative client to this server: “select* from VOLHISTORY where VOLUME_NAME='volume_name' AND TYPE='STGNEW.” The results of the **QUERY VOLUME** command indicate when this volume was last written. The information from the **SELECT** operation reports when this volume was added to the storage pool. Often, **AUDIT VOLUME** might report files as damaged because at the time that the data was written, the hardware was malfunctioning and did not write the data correctly, even though it reported to the IBM Tivoli Storage Manager server that the operation was successful. As a result of this device malfunction, many files on many different volumes might be affected. Perform the following steps to resolve this issue:
 - a. Evaluate the system error logs or other information about this drive to determine if it still reports an error. If errors are still reported, they must first be resolved. To resolve a hardware issue, work with the hardware vendor to correct the problem.
 - b. If this storage pool is a copy of a storage pool volume, simply delete this volume using the **DELETE VOLUME** *volumeName* **DISCARDATA=YES** command. The next time a storage pool backup is run for the primary storage pool or storage pools where this damaged data resides, it will be backed up again to this copy storage pool and no further action is necessary.
 - If this storage pool is a primary storage pool volume and the data was written directly to this volume when the client stored the data, then it is likely that there are no undamaged copies of the data on the server. If possible, back up the files again from the client.
 - If this storage pool is a primary storage pool volume but the data was put on this volume by **MIGRATION**, **MOVE DATA**, or **MOVE NODEDATA** commands, there might be an undamaged copy of the file on the server. If the primary storage pool that contained this file was backed up to a copy storage pool prior to the **MIGRATION**, **MOVE DATA**, or **MOVE NODEDATA** commands, then an undamaged file might exist. If an undamaged file exists, issue the **UPDATE VOLUME** *volumeName* **ACCESS=DESTROYED** command and then issue the **RESTORE VOLUME** *volumeName* command to recover the damaged files for this volume from the copy storage pool.

Analyzing error messages ANR1330E and ANR1331E

You might receive error message ANR1330E or ANR1331E while data is being read from a Tivoli Storage Manager storage pool volume.

When the 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, and tape drive) encountered a problem while the data was read.
- An error occurred while the data is being written and the data is damaged.
- A database restore operation was performed and a volume was not appropriately audited so that it is in sync with the point-in-time (PIT) restore time.

You must first determine whether the 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 caused the damage.

You can ignore the error, but only if the error occurs infrequently. Hardware occasionally encounters an error while reading data. In most cases, the hardware recognizes that an error occurred and recovers without having to report it. But there are times when the data is read in an altered (damaged) state because of a temporary hardware error. The following list defines the results of reading data and receiving an error:

Audit OK, error reading intact data on media

Tivoli Storage Manager checks the self-describing information and reports the data as damaged if it does not match what is expected. In messages ANR1330E and ANR1331E, the data is reported as damaged.

If after you audit the volume, messages ANR1330E and ANR1331E are displayed frequently, determine which hardware device is causing the data to be read incorrectly. Query the activity log to find the date and time that messages ANR1330E and ANR1331E were issued and provide the information to your hardware support team. With this information, they can examine the hardware error logs for any operations that might have completed abnormally. Also, have your hardware support team ensure that the device drivers and microcode maintenance for the hardware is up to date.

A common place for such errors to occur is on a storage area network (SAN). Typically, these errors occur if many link level interrupt (LLI) errors occur on the switch or the network. LLI errors indicate that the system is performing poorly and are known to cause data to be modified during retransmission. Ask your hardware support team to examine the network

error logs for instances of LLI errors. Look for LLI errors that were logged around the time that the ANR1330E and ANR1331E message were issued.

Audit failed, data damaged on media

If the audit reports the data as damaged, an error might have occurred that caused the data to be written incorrectly onto the media. Also, a database restore operation might have a volume that was not appropriately audited to synch-up with the PIT restore time. Determine, from the audit reports, when the data was written and examine message ANR1331E to find out which hardware device damaged the data. See the following example data:

ANR1330E

The server has detected possible corruption in an object being restored or moved. The actual values for the incorrect frame are: magic C6A2D75D hdr version 35134 hdr length 43170 sequence number 160421181 data length 7E53DCD8 server id 348145193 segment id 327643666840426461 crc 06E04914.

ANR1331E

Invalid frame detected. Expected magic 53454652 sequence number 00000023 server id 00000000 segment id 2062.

The segment ID number in message ANR1331E in this example is 2062. To determine the date that the data was inserted into the server, issue the following command:

```
SHOW INVO 0 2062
```

The following example shows the output from the **SHOW INVO** command:

```
OBJECT: 0.2062 (Backup):  
Node: NODE1 Filespace: \\node1\c$ (Unicode).  
\5400\BF\ BFDEFS.H  
Type: 2 (File) CG: 1 Size: 0.89088 HeaderSize: 364
```

BACKUP OBJECTS ENTRY:

```
State: 1 Type: 2 MC: 1 CG: 1  
\\node1\c$ (Unicode) : \TESTFILES\ FILE1.TXT (MC: DEFAULT)  
Active, Inserted 11/29/2009 13:28:26
```

EXPIRING OBJECTS ENTRY:

Expiring object entry not found.

Find the Inserted field and note the date and time. In this example, the object was inserted on 11/29/2009 at 13:28:26. Provide your hardware support team with the date and time. The support team can examine the hardware error logs for any operations that completed abnormally. Also, ask the support team to ensure that the device drivers and microcode maintenance for the hardware is up to date. Your hardware support team must examine the SAN network error logs. Look for errors around the time that the data was inserted into Tivoli Storage Manager.

If the **SHOW INVO** command returns unhelpful output, issue the following command to determine the date of insertion:

```
SHOW BFO 0 xxx
```

where xxx is the segment group ID. The example shows the output from the **SHOW BFO** command:

```
Bitfile Object: 0.xxx  
**Super-bitfile 0.xxx contains following aggregated bitfiles  
(offset/length)  
0.2063 0.75295 0.3071 Active  
0.2064 0.78366 0.88780 Active  
0.2065 0.167146 0.13831 Active  
0.2066 0.180977 0.21254 Active  
0.2067 0.202231 0.3808 Active
```

0.2068 0.206039 0.11261 Active

```
**Disk Bitfile Entry
Bitfile Type: PRIMARY
Storage Format: 22
Logical Size: 0.217364
Physical Size: 0.221184
Number of Segments: 1,
Deleted: False
Storage Pool ID: 1
Volume ID: 2
Volume name: TapeVol1
```

Get an aggregated bit file number from the first entry on the list of aggregated bit files. In the preceding example, the first aggregated bit file number is 2063. Issue the **SHOW INVO** command using 2063.

No hardware errors at time of insertion

If the hardware support team discovers that no hardware errors occurred at the time the data was inserted into Tivoli Storage Manager, contact the IBM support team. Provide the team with the activity log at the time that messages ANR1330E and ANR1331E were issued. Also, issue the **AUDIT VOLUME FIX=NO** command with the following trace, and provide the Tivoli Storage Manager support team with the trace:

```
TRACE ENABLE BF AF DF SS AS DS SSFRAME
TRACE DISABLE BFLOCK AFLOCK SSLOCK
TRACE BEGIN filename
```

Fixing damaged files on media

If you find that the data is damaged on a volume, issue the **AUDIT VOLUME FIX=YES** command on the volume. If the following conditions are true, the data remains marked as damaged on the primary pool volume:

- The volume is a primary pool volume
- The data is backed up to a copy storage pool
- The data is damaged

After the **AUDIT VOLUME FIX=YES** command completes, issue the **RESTORE VOLUME** command for the primary pool volume. The damaged data is replaced with a new copy of the data. If the **AUDIT VOLUME FIX=YES** command successfully read the data, the data is no longer marked as damaged in the primary storage pool.

If there is no backup copy, the **AUDIT VOLUME FIX=YES** command deletes the data. If the data that was deleted is backup data, it is placed on the server the next time the client backup runs.

If the data that is being deleted by the **AUDIT VOLUME FIX=YES** command is on a copy-storage-pool volume, the data is deleted from the copy pool volume. The next time that the primary storage pool is backed up, a new copy is added to the copy storage pool.

Files are not expired after reducing versions

You can update the server policies to reduce the number of versions of a file that you want to keep, however errors can sometimes be generated as a result of these updates.

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 the process might expire these files. The server applies only the **Versions Data Exists** and **Versions Data Deleted** values to files at the time that they are backed up to the server. When a file is backed up, the server counts the number of versions of that file and if that exceeds the number of versions that must be kept, the server marks the oldest versions that exceed this value to expire.

Process symptoms indicate migration errors

You might be faced with process symptoms that point to migration as the cause for errors.

Migration does not run for sequential media storage pool

If migration does not run for sequential media storage pools, issue the **QUERY STGPOOL** *stgpoolName* **F=D** command.

Migration from sequential media storage pools calculates the “Pct. Util” as the number of volumes in use for the storage pool, relative to the total number of volumes that can be used for that storage pool. Similarly, it calculates the “Pct. Migr” as the number of volumes with data that can be migrated, in use for the storage pool, relative to the total number of volumes that can be used for that storage pool. Because it might be considering unused scratch volumes in this calculation, there might not appear to be sufficient data that can be migrated in the storage pool to require migration processing.

Migration uses only one process

Issue the **QUERY STGPOOL** *stgpoolName* **F=D** and **QUERY OCCUPANCY * * STGPOOL=** *stgpoolName* command.

The following are the reasons why only one migration process is running:

- The Migration Processes setting for the storage pool is set to one or is not defined (blank). If true, issue the **UPDATE STGPOOL** *stgpoolName* **MIGPROCESS=*n*** command, where *n* is the number of processes to use for migrating from this pool. Note that this value must be less than or equal to the number of drives (mount limit) for the NEXT storage pool where migration is storing data.
- If the **QUERY OCCUPANCY** command only reports a single client node and file space in this storage pool, migration can only run a single process even if the Migration Processes setting for the storage pool is greater than one. Migration processing is partitioning data, based on client node and file space. For migration to run with multiple processes, data for more than one client node needs to be available in that storage pool.

Resolving storage pool issues

Storage pools are integral to a successful server operation. The IBM Tivoli Storage Manager database contains information in storage pools about registered client nodes, policies, schedules, and the client data.

This information must be available and valid in order for Tivoli Storage Manager to function correctly. Storage pool errors can be related to the following issues:

- Failed transactions
- A storage pool experiencing a high volume usage after increasing the MAXSCRATCH value
- A storage pool having “Collocate?=Yes” but volumes still containing data for many nodes
- Unable to store data in an active data pool by using the simultaneous-write function or by issuing the **COPY ACTIVATEDATA** 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.

About this task

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:

Procedure

1. Issue **QUERY VOLUME volname F=D** for the volumes in the referenced storage pool. For any volumes reported with access other than Read/Write, check that volume. A volume might be marked Read/Only or Unavailable because of a device error. If the device error is resolved, issue the **UPDATE VOLUME volname ACCESS=READWRITE** command to let the server select and try to write data to that volume.
2. Issue **QUERY VOLUME volname** for the volumes in the referenced storage pool. Volumes that report “pending” for the volume status are volumes that are empty but waiting to be reused again by the server. The wait time is controlled by the REUSEDELAY setting for the storage pool and displayed as “Delay Period for Volume Reuse” on the **QUERY STGPOOL** command. Evaluate the REUSEDELAY setting for this storage pool and, if appropriate (based upon your data management criteria), lower this value by issuing the **UPDATE STGPOOL stgpoolname REUSEDELAY=nn** command, where *stgpoolname* is the name of the storage pool and *nn* is the new reuse delay setting. The key to getting the data collocated is to have sufficient space in the target storage pool for the collocation processing to select an appropriate volume. Having sufficient space in the target storage pool is significantly influenced by the number of scratch volumes in a storage pool.
3. Issue the **QUERY STGPOOL F=D** command to verify whether the ACCESS is Read/Write.

Storage pool experiences high volume usage after increasing MAXSCRATCH value

For collocated sequential storage pools, increasing the **MAXSCRATCH** value might cause the server to use more volumes.

The server uses more storage pool volumes in this case because of the collocation processing. Collocation groups user data for a client node onto the same tape. During a client backup or archive operation, if no tapes currently have data for this client node, the server selects a scratch volume to store the data. Then, for other client nodes storing data, the server again selects a scratch volume. The reason that scratch volumes are not selected prior to changing the **MAXSCRATCH** setting is that if there is no scratch volume available and no preferred volume already assigned for this client node, the volume selection processing on the server ignores the collocation request and stores the data on an available volume.

Storage pool is set to use collocation, but volumes contain data that is not collocated

When a storage pool is collocation enabled (the **COLLOCATION** parameter is set to **GROUP**, **NODE**, or **FILESPEC**), many volumes might contain data that is not collocated.

There are two possibilities for this situation:

- The data was stored on volumes in this storage pool before enabling the storage pool for collocation.
- The storage pool ran out of scratch tapes and stored data on the best possible volume, even though the request to collocate was ignored.

If data for multiple nodes ends up on the same volume for a storage pool that is collocation enabled, use one of the following actions:

- Issue the **MOVE DATA** command for the volume or volumes affected. The process reads the data from the specified volume and moves it to a different volume in the same storage pool if:
 - Scratch volumes are available, or
 - Volumes with sufficient space are assigned to this client node for collocating their data
- Allow migration to move all the data from that storage pool by setting the **HIGHMIG** and **LOWMIG** thresholds. By allowing the migration of all data to the **NEXT** storage pool, the collocation requirements are processed if the following are true:
 - The **NEXT** storage pool is collocation enabled
 - The **NEXT** storage pool has sufficient scratch volumes
 - The **NEXT** storage pool is assigned volumes to satisfy the collocation requirements
- Issue the **MOVE NODEDATA** command for the client nodes whose data is in that storage pool. If scratch volumes are available or volumes with sufficient space are assigned to this client node for collocating their data, the following events happen:
 - The **MOVE NODEDATA** process reads the data from the volumes that this node has data on
 - The **MOVE NODEDATA** process moves the data to a different volume or volumes in the same storage pool

The key to getting the data collocated is to have sufficient space in the target storage pool for the collocation processing to select an appropriate volume. There must be enough empty volumes available in the storage pool to allow collocation to select a new volume. Ensure that enough empty volumes are available instead of using a volume that already has data on it from a different node. Empty volumes can be scratch volumes if the storage pool is defined with enough scratch volumes, or define the empty volumes by issuing the **DEFINE VOLUME** command.

Resolving storage problems for active data pools

You might experience difficulty in storing data in an active data pool by using the simultaneous-write function or by issuing the **COPY 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.

Resolving issues with cloud-container storage pools

With Tivoli Storage Manager, you can back up data to and restore data directly from a cloud-container storage pool. At times, you might encounter performance issues or limitations with cloud-container storage pools.

Use the following guidance to resolve issues and handle limitations:

Problems with deleting objects from the cloud

The performance of a cloud-container storage pool relies on the network connection between the server and the cloud. Depending on the cloud provider, deleting objects from cloud storage can take a significant amount of time. For example, if you are using Swift OpenStack as your cloud provider, you must individually delete cloud objects, and network latency impacts each delete operation. If you need to delete many cloud objects in a short time, make sure that Tivoli Storage Manager can eventually delete all the objects that are saved in the cloud. For example, you might encounter poor performance if you use an off-premises cloud provider and you regularly delete large file spaces from Tivoli Storage Manager.

Performance issues with restoring files

If you encounter poor performance when restoring files, verify that the restore operation is available in your environment. See technote 1659833.

Restrictions for cloud-container storage pools

The following functions are not compatible with cloud-container storage pools:

- Replication of a cloud-container storage pool with the **PROTECT STGPOOL** command
- Migration
- Reclamation
- Aggregation
- Collocation
- Simultaneous-write operations
- Storage pool backup operations
- Use of virtual volumes

In addition, you cannot use the **NEXTSTGPOOL** parameter with the **DEFINE STGPOOL** command on a cloud-container storage pool or a directory-container storage pool because Tivoli Storage Manager cannot determine when the cloud storage provider is full. Use the **NEXTSTGPOOL** parameter to specify only a random-access or primary sequential storage pool. As a result, the overflow capability is not available for container-based storage pools.

No failover to the cloud after the local storage becomes full

If you use storage pool directories with a cloud-container storage pool, and the directories contain no more free space, backup operations stop prematurely. To avoid this situation, allocate more storage pool directories to give the storage pool more local storage space for backup operations. You can also wait for the data to be automatically cleaned up from the local directories after the data moves to the cloud.

Limitations on using node replication with a cloud-container storage pool

You can use a cloud-container storage pool as a destination storage pool on a target replication server. However, you cannot use a cloud-container storage pool as a source storage pool on a source replication server. To provide redundancy, use the replication capabilities that are available from the cloud storage provider.

File types to avoid with cloud-container storage pools

For a cloud-container storage pool, avoid storing client data types that are not optimized for storing data in removable media storage pools. For example, avoid storing IBM Tivoli Storage Manager for Virtual Environments control files and Data Protection for Microsoft SQL Server metadata files. For more information, see the following documents:

- Technote 1659833
- IT11763: METADATA CONSIDERATIONS ARE MISSING IN DATA PROTECTION FOR SQL SERVER DOCUMENTATION.

Chapter 4. Resolving Operations Center problems

AIX

Linux

Windows

If a problem occurs with the Tivoli Storage Manager Operations Center and you cannot solve it, you can consult the descriptions of known problems for a possible solution. You might also need to review log files and enable extended tracing for the Operations Center.

Log files overview

AIX

Linux

Windows

If you contact IBM Software Support about a problem with the Operations Center, they might request that you send them log files.

List of log files

IBM Software Support might request that you send them the following log files:

- Up to eight Operations Center log files:
 - `tsm_opscntr.log`
 - `tsm_opscntr1.log`
 - `tsm_opscntr2.log`
 - `tsm_opscntr3.log`
 - `tsm_opscntr4.log`
 - `tsm_opscntr5.log`
 - `tsm_opscntr6.log`
 - `tsm_opscntr7.log`

More than one Operations Center log file can exist for the following reasons:

- If the Operations Center log is up to 8 MB in size, the current version is `tsm_opscntr.log`, the previous version is `tsm_opscntr1.log`, the version before that is `tsm_opscntr2.log`, and so on.
- If the Operations Center log is greater than 8 MB in size, the log is spread over multiple files, each with a maximum size of 8 MB. For example, if the log is 15 MB in size, it is spread over the files `tsm_opscntr.log` and `tsm_opscntr1.log`.

Tip: If IBM Software Support requests that you perform an extended trace of the Operations Center, you can identify which of the Operations Center log files are created during the trace from the modification times of the files.

- Web-server log files:
 - `console.log`
 - `messages.log`
- First-failure-data-capture (FFDC) log files:
 - `exception_summary_date_time.log`
 - `ffdc_date_time.log`

Location of log files

- The Operations Center and web-server log files are in the following directory:

AIX **Linux** `installation_dir/ui/Liberty/usr/servers/guiServer/logs`

Windows `installation_dir\ui\Liberty\usr\servers\guiServer\logs`

where *installation_dir* represents the directory in which Tivoli Storage Manager is installed. For example:

AIX **Linux** `/opt/tivoli/tsm`

Windows `c:\Program Files\Tivoli\TSM`

Tip: You can also view the Operations Center log from within the Operations Center.

- The FFDC log files are in the same location, but in the `ffdc` subdirectory.

Related tasks:

“Starting an extended trace of the Operations Center” on page 113

Viewing the Operations Center log from within the Operations Center

AIX **Linux** **Windows**

The Operations Center log contains data from a trace of Operations Center events. You can view the log in the Operations Center, or you can go to the directory that contains the log file and open the file.

Procedure

To view the Operations Center log while you are logged in to the Operations Center, complete the following steps:

1. On the Operations Center menu bar, hover over the question-mark icon and select **About Operations Center**.
2. In the window that is displayed, click **Installation Details**.
3. Click the **View Log** tab.
4. Click **Display Log**.

Related tasks:

“Starting an extended trace of the Operations Center” on page 113

Alerts are not updated immediately

AIX **Linux** **Windows**

On the Alerts page of the Operations Center, when you attempt to assign multiple alerts to an administrator or to close multiple alerts, the alerts are not assigned or closed immediately.

Symptom

Table 8 on page 101 shows sample data from one test environment when an administrator updated multiple alerts. These results might differ from the results in your storage environment.

Table 8. Approximate delay times when alerts were updated in a controlled environment

Number of alerts updated	Delay for hub-server alerts	Delay for alerts from spoke servers with Tivoli Storage Manager Version 7.1.0	Delay for alerts from spoke servers with Tivoli Storage Manager V6.3.4
1	6 seconds	7 seconds	7 seconds
10	6 seconds	7 seconds	9 seconds
100	6 seconds	8 seconds	40 seconds
1,000	10 seconds	20 seconds	5.5 minutes
10,000	45 seconds	1.25 minutes	1 hour

For example, when the administrator selected 10,000 hub-server alerts and clicked **Close**, it took approximately 45 seconds for the alerts to close.

Solution

Wait until the alerts are updated, or update fewer alerts at a time. For faster response rates, upgrade spoke servers that are running V6.3.4 to run V7.1.0 or V7.1.1.

Active tasks are not canceled immediately

AIX Linux Windows

When you select multiple tasks on the Active Tasks page of the Operations Center and attempt to cancel them, the tasks are not canceled immediately. There is a longer delay for spoke-server tasks than there is for hub-server tasks.

Symptom

Table 9 shows sample data from one test environment when an administrator canceled multiple tasks. These results might differ from the results in your storage environment.

Table 9. Approximate delay times when tasks were canceled in a controlled environment

Number of tasks canceled	Delay for hub-server tasks	Delay for spoke-server tasks
1	5 seconds	5 seconds
10	5 seconds	7 seconds
100	10 seconds	25 seconds
1000	40 seconds	3.5 minutes

For example, when the administrator selected 1000 hub-server tasks and clicked **Cancel**, it took approximately 40 seconds for the tasks to be canceled.

Solution

Wait until the tasks are canceled, or cancel fewer tasks at a time.

Further known issues with the Operations Center

AIX

Linux

Windows

Known issues are documented in the form of technotes in the Support knowledge base. As problems are discovered and resolved, IBM Software Support updates the knowledge base. By searching the knowledge base, you can quickly find workarounds or solutions to problems.

- For a list of known issues, see the following web page in the Support knowledge base: [Known issues with Tivoli Storage Manager Operations Center V7.1.6](#).
- For more issues that become known after product release, see the following search results: [Search results for known issues with Tivoli Storage Manager Operations Center V7.1.6](#).

Chapter 5. Resolving communication problems

The need for connectivity in IBM Tivoli Storage Manager means that any communication error can render your application useless. Communication errors might be attributed to TCP/IP configuration, client and server connections, and other causes.

Resolving errors created when connecting to the server

Problems that are generated while you are connecting to the server might be related to your communication options.

To correct the error, perform any or all of the following options:

- Review the changes in the client communication options in the client option file (if they exist) and try to revert back to the previous values. Retry the connection.
- If the server communication settings were changed, either update the client communication options to reflect the changed server values or revert the server back to its original values.
- If any network settings were changed, such as the TCP/IP address for the client or server (or a firewall), work with the network administrator to update the client, server, or both for these network changes.

Resolving failed connections by clients or administrators

The two main cases for connection failures are general failure, where no connections at all are allowed, or an isolated failure where some connections are allowed but others fail.

If no connections at all are possible, it might be necessary to run the server in the foreground so that a server console is available, and additional diagnostic steps can be taken. Check the settings to verify the proper configuration for communicating with the server:

- Ensure that the server is able to bind to a port when it is started. If it is unable to bind to a port, then it is likely that some other application is using that port. The server can not bind (use) a given TCP/IP port if another application is already bound to that port. If the server is configured for TCP/IP communications and successfully binds to a port on startup for client sessions, the following message is issued:

```
ANR8200I TCP/IP driver ready for connection with clients on port 1500.
```

If a given communication method is configured in the server options file, but a successful bind message is not issued during server startup, then there is a problem initializing for that communication method.

- Verify that the code **TCPPORT** setting in the server options file is correct. If the code setting is inadvertently changed, the clients fail to connect. That is because the clients are trying to connect to a different TCP/IP port than the one the server is listening on.
- If multiple servers are using the same TCP/IP address, ensure that the **TCPPORT** and **TCPADMINPORT** for each server are unique. For example, there are two servers at the same TCP/IP address. The first server has a **TCPPORT** of 1500 and a **TCPADMINPORT** of 1500. The second server has a **TCPPORT** of 1501 and a

TCPADMINPORT of 1500. The first server to grab port 1500 locks out the other server from port 1500 and clients can no longer access the first server. Administrative clients always connect to the second server. A better choice of ports for each server would be 1500 and 1501 for **TCPSPORT**; 1510 and 1511 for **TCPADMINPORT**.

- Check that the server is enabled for sessions. Issue the **QUERY STATUS** command and verify that “Availability: Enabled” is set. If the result states “Availability: Disabled,” issue the **ENABLE SESSIONS** command.
- If specific clients are unable to connect to the server, check the communication settings for those clients. For TCP/IP, check the **TCPSERVERADDRESS** and **TCPSERVERPORT** options in the client options file.
- If only a specific node is rejected by the server, verify that the node is not locked on the server. Issue the **QUERY NODE** *nodeName* command, where *nodeName* is the name of the node to check. If the result states “Locked?: Yes,” then evaluate why this node is locked. Nodes can only be locked by using the **LOCK NODE** administrative command. If it is appropriate to unlock this node, issue the **UNLOCK NODE** *nodeName* command, where *nodeName* is the name of the node to unlock.
- If the computer on which the server is running is having memory or resource allocation problems, it might not be possible to start new connections to the server. The memory or resource allocation problem might be cleared up temporarily if you either halt and restart the server, or if you halt and restart the computer itself. This action is a temporary solution, and diagnosis should be continued for either the operating system or the server because the memory and resource allocation problem might indicate an error in either.

Resolving Secure Sockets Layer errors

Secure Sockets Layer (SSL) errors can be attributed to an incorrect environment setup, a bad server certificate, connection problems, out-of-sync conditions, or other causes.

Use the following guidance to resolve common SSL client-to-server and server-to-server problems:

Not connecting to the server after using a vendor-acquired certificate authority (CA) certificate

If you are using a vendor-acquired certificate and it was not added to the server, specify the root certificate as trusted in the server key database. To add the root certificate to the database, issue this command:

```
gsk8capicmd -cert -add -db cert.kdb -pw password  
-label name -file .der_file -format ascii
```

The CA root certificate was not added to the client

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

```
gsk8capicmd -cert -add -db dsmcert.kdb -pw password  
-label my CA -file ca.arm -format ascii
```

Unable to run gsk8capicmd.exe (IBM Global Security Kit [GSKit])

In most cases, this Windows error is generated by an incorrect environment setup. Set up the **PATH** variable as directed before you run the **gsk8capicmd** utility.

ANS1595E Bad server certificate

This error is reported when the server certificate is not known to the client or server. The “bad server certificate” error can occur under these conditions:

- The certificate was never imported
- The cert256.arm certificate file was corrupted before the certificate was imported
- The command to import the certificate was entered incorrectly
- The *DSM_DIR* variable points to the wrong directory, which contains an incorrect client key database (dsmcert.kdb)
- The server is set up for Transport Layer Security (TLS) 1.2 but the client is not at a sufficient level (6.3 is required).
- The server is set up for TLS 1.2 but the client imported the cert.arm file instead of the cert256.arm file.
- The server is set up for TLS 1.2 but the client imported the cert256.arm file instead of the cert.arm file.

Repeat all the steps necessary for importing the server certificate and check the *DSM_DIR* variable. For more information about the failure, see the dsmerror.log file. The client error log might also contain information about specific IBM GSKit failure.

ANS1592E Failed to initialize SSL protocol

This error occurs on the client and indicates that the SSL connection was not established. For more information about the failure, see the client error log. The server does not accept SSL sessions on the port to which the client or server is trying to connect. Determine whether the client or server points to the correct server port (TCPPort), which can be a port number that is different from 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 server at a port that is defined by SSLTCPPORT or SSLTCPADMINPORT, the server establishes a session and initiates an SSL “handshake.” If the client is not SSL-enabled, it cannot complete the SSL handshake process. The session then seems to stop, but times out through the server IDLEWAIT option or end when the server administrator issues the **CANCEL SESSION** command to manually cancel it. The example illustrates a session in this state, from the server:

```
TSM:SERVER1>query session
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY SESSION
```

Sess Number	Comm. Method	Sess State	Wait Time	Bytes Sent	Bytes Recvd	Sess Type	Platform	Client Name
1	SSL	IdleW	17 S	0	0	Node		

Important: Because the computing environment might cause a valid handshake process to take some time, do not assume that the result always indicates a non-SSL client.

ANR8583E and GSKit return code 420, and ANR8581E with GSKit return code 406 occur for the same Tivoli Storage Manager client session

When server messages ANR8583E and ANR8581E occur for the same client session, it is likely that the client generated an ANS1595E message. Message ANS1595E typically occurs 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.

ANR3338E TLS is at an earlier level than 1.2

This error is reported when the server and the storage agent attempt to connect with an SSL protocol earlier than TLS 1.2. For server and storage agent communication, if the `SSLDISABLELEGACYTLS` option is specified, TLS sessions must connect at a minimum level of TLS 1.2 or the session is rejected.

Cross-defining servers without SSL=YES causes a server hang

If you plan to use SSL communication, the SSL infrastructure must be in place on the source and target replication servers. Required SSL certificates must be in the key database file that belongs to each server. The SSL function is active if the server options file contains the `SSLTCPPORT` or `SSLTCPADMINPORT` option or if a server is defined with `SSL=YES` at startup.

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 that is being used can be uniquely identified.

Related reference:

Appendix C, "IBM Global Security Kit return codes," on page 199

Recovering the key database file password

If you forgot the current key database file password, IBM Tivoli Storage Manager can help you to recover it.

Before you begin

You must have system privileges to administer the key database file password recovery.

About this task

Complete the following steps to recover and update the key database file password:

Procedure

1. Issue the `QUERY SSLKEYRINGPW` command to display the current key database password.
2. Issue the following command to use the server record of the key database password to update the password:

```
SET SSLKEYRINGPW password UPDATE=Y
```

where *password* is the password retrieved by the `QUERY SSLKEYRINGPW` command.

What to do next

Tip: If the `cert.kdb` file does not exist, you can create a new file by restarting the server. The 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.

Troubleshooting the certificate key database

Backup copies of the `cert.kdb` file ensure that Transport Layer Security (TLS) starts when you restore the Tivoli Storage Manager server. If you have a backup copy, you can restore the file and restart the server.

Procedure

To create a backup copy of the certificate key database, `cert.kdb`, complete the following steps:

1. Issue the **DELETE KEYRING** server command to delete the password information in the Tivoli Storage Manager key database.
2. Delete all remaining `cert.*` files.
3. Shut down the server.
4. Start the server. The server automatically creates a new `cert.kdb` file and a corresponding entry in the Tivoli Storage Manager database. If you do not issue the **DELETE KEYRING** command, the server attempts, on startup, to create the key database with the previous password.
5. Redistribute the new `.arm` file to all backup-archive clients that are using TLS. If you are using TLS 1.2, use the `cert256.arm` file. Use the `cert.arm` file if the TLS protocol you use earlier than 1.2. Reinstall any third-party certificates on the backup-archive client. If you are using an LDAP directory server to authenticate passwords, add the root certificate that was used to sign the LDAP server's certificate. If the root certificate is already a default trusted certificate, you do not have to add it again.

What to do next

If the `cert.kdb` key database file does not exist, the server creates it. One or both of the `SSLTCPSPORT` and `SSLTCPADMINPORT` options must be in the server options file when the server is started. The server generates a changeable password and also generates a self-signed certificate that can be extracted for clients and IBM business partners servers to use. If the `cert.kdb` file exists and the server did not create it, an out-of-sync condition occurs, preventing the server from setting up SSL communications.

Chapter 6. Resolving storage agent problems

can back up and restore client data directly to and from SAN-attached storage by using the storage agent.

Checking the server activity log for storage agent information

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

Storage agents start and manage many sessions to the server. Review the server activity log file for messages from the storage agent. To review the activity log messages, issue the **QUERY ACTLOG** command.

If no messages in the server activity log file are for this storage agent, verify the communication settings:

- Issue **QUERY SERVER F=D** on the server and verify that the high-level address (HLA) and low-level address (LLA) set for the server entry representing this storage agent are correct.
- In the device configuration file specified in the `dsmsta.opt` file, verify that the **SERVERNAME** as well as the HLA and LLA are set correctly in the **DEFINE SERVER** line.

Check for any error messages on the server for this storage agent.

Resolving an error caused by reading or writing to a device

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

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

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

Storage agents are particularly vulnerable if path information is changed or not correct. Issue the **QUERY PATH F=D** command on the server. For each of the storage agent's paths, verify that the settings are correct. In particular, verify that the device listed matches the system device name. If the path information is not correct, update the path information by issuing the **UPDATE PATH** command.

Resolving problems caused by changing storage agent options

Changes to options in the storage agent option file might cause operations to fail, even though they had previously succeeded.

Review any changes to the storage agent option file. Try reverting the settings to their original values and retrying the operation. If the storage agent now works correctly, try reintroducing changes to the storage agent option file one-at-a-time and retry storage agent operations until the option file change that caused the failure is identified.

Resolving problems caused by changing server options or settings

Changes to options in the server option file or changes to server settings using the **SET** commands might affect the storage agent.

Review any changes to server option settings. Try reverting the settings to their original values and retrying the operation. If the storage agent now works correctly, try reintroducing changes to the storage agent option file one-at-a-time and retry storage agent operations until the option file change that caused the failure is identified.

Review server settings by issuing the **QUERY STATUS** command. If any settings reported by this query have changed, review the reason for the change and, if possible, revert it to the original value and retry the storage agent operation.

Storage agent LAN-free setup

LAN-free data movement is the direct movement of client data between a client computer and a storage device on a SAN, rather than on a LAN. You might be experiencing problems with the storage agent that are related to your LAN-free setup.

Resolving the issue of data being sent directly to the server

The client summary statistics do not report any bytes transferred LAN-free.

Before you begin

The client reports the bytes sent LAN-free by issuing the **"ANE4971I LAN-free Data Bytes: xx KB"** command. Similarly, the server does not report any instance of **"ANR0415I Session SESS_NUM proxied by STORAGE_AGENT started for node NODE_NAME"** for this node and storage agent, indicating that the LAN-free proxy operation was done for this client node.

The client will only attempt to send data LAN-free with the storage agent if the primary storage pool destination in the server storage hierarchy is LAN-free. A server storage pool is LAN-free-enabled for a given storage agent if one or more paths are defined from that storage agent to a SAN device.

About this task

To determine if the storage pool destination is configured correctly, perform the following procedures:

Procedure

1. Issue the **QUERY NODE** *nodeName* command to report the policy domain to which this node is assigned.
2. Issue the **QUERY COPYGROUP** *domainName policySetName mgmtclassName F=D* command for the management classes that this node would use from their assigned policy domain. Note that this command reports information for backup files. To query copy-group information for archive files, issue the **QUERY COPYGROUP** *domainName policySetName mgmtclassName TYPE=ARCHIVE F=D* command.
3. Issue the **QUERY STGPOOL** *stgpoolName* command, where *stgpoolName* is the destination reported from the previous **QUERY COPYGROUP** queries.
4. Issue the **QUERY DEVCLASS** *deviceClassName* command for the device class used by the destination storage pool.
5. Issue the **QUERY LIBRARY** *libraryName* command for the library reported for the device class used by the destination storage pool.
6. Issue the **QUERY DRIVE** *libraryName F=D* command for the library specified for the device class used by the destination storage pool. If no drives are defined to this library, review the library and drive configuration for this server and issue the **DEFINE DRIVE** command to define the needed drives. If one or more of the drives report "ONLINE=No" evaluate why the drive is offline and, if possible, update it to online by issuing the **UPDATE DRIVE** *libraryName driveName ONLINE=YES* command.
7. Issue the **QUERY SERVER** command to determine the name of the storage agent as defined to this server.
8. Issue the **QUERY PATH** *stgAgentName* command, where *stgAgentName* is the name of the storage agent defined to this server and reported in the **QUERY SERVER** command. Review this output and verify that one or more paths are defined for drives defined for the device class used by the destination storage pool. If no paths are defined for this storage pool, issue the **DEFINE PATH** command to define the needed paths. Also, review this output and verify that the path is online. If paths are defined but no paths are online, update the path to online by issuing the **UPDATE PATH** *srcName destName SRCTYPE=SERVER DESTTYPE=DRIVE ONLINE=YES* command.

Resolving a disqualified LAN-free-enabled storage pool

The server disqualifies a storage pool from being a LAN-free-enabled storage pool if it was configured for simultaneous-write operations.

In this case, data from the client is sent directly to a server that will not be using a LAN-free storage pool.

Issue the **QUERY STGPOOL** *stgpoolName F=D* command for the destination storage pool for this client. If the storage pool is set for simultaneous-write operations, the "Copy Storage Pool(s):" value references one or more other storage pool names and IBM Tivoli Storage Manager interprets the simultaneous-write operation to be a higher priority than the LAN-free data transfer. Because simultaneous-write operations are considered a higher priority operation, this storage pool is not reported as LAN-free-enabled and as such, the client sends the data directly to the server. The storage agent does not support simultaneous-write operations.

Ensuring that data is transferred using a LAN-free environment

The storage agent and client are both able to manage failover directly to the server, depending upon the LAN-free configuration and the type of error encountered.

Because of this failover capability, it might not be apparent that data is being transferred over the LAN when it was intended to be transferred LAN-free. It is possible to set the LAN-free environment to limit data transfer to only LAN-free.

To test a LAN-free configuration, issue the **UPDATE NODE** *nodeName* **DATAWRITEPATH=LAN-FREE** command for the client node whose LAN-free configuration you want to test. Next, try a data storage operation such as backup or restore. If the client and storage agent attempt to send the data directly to the server using the LAN, the following error message is received:

```
ANR0416W Session sessionNumber for node nodeName not allowed to operation
using path data transfer path
```

The *operation* reported indicates either READ or WRITE, depending upon the operation attempted. The path is reported as LAN-free.

If this message is received when you are trying a LAN-free operation, evaluate and verify the LAN-free settings. Generally, if data is not sent LAN-free when the client is configured to use LAN-free, the storage pool destination for the policy assigned to this node is not a LAN-free enabled storage pool, or the paths are not defined correctly.

Chapter 7. Using trace to resolve problems

IBM Tivoli Storage Manager can, at times, experience problems that you can resolve through trace.

Starting an extended trace of the Operations Center

AIX Linux Windows

By default, the Operations Center log contains data from a basic trace of Operations Center events. IBM Software Support might request that you start an extended trace.

About this task

To start an extended trace of the Operations Center, complete one of the following procedures:

Related concepts:

“Log files overview” on page 99

Related tasks:

“Viewing the Operations Center log from within the Operations Center” on page 100

Tracing the Operations Center by enabling logging functions from within the Operations Center

AIX Linux Windows

From within the Operations Center, you can enable logging functions and start an extended trace that adds troubleshooting data to the Operations Center log.

About this task

In the following procedure, you enable groups of logging functions and start an extended trace.

Attention: Ensure that you disable the groups after the trace. Otherwise, the performance of the Operations Center can be affected.

Procedure

To trace the Operations Center, complete the following steps:

1. On the Operations Center menu bar, hover over the question-mark icon and select **About Operations Center**.
2. Click **Installation Details**.
3. Click the **Logging** tab.
4. From the list of logging groups, select only the rows that IBM Software Support request that you select, and click **Enable**.
5. Confirm that you want to enable the logging groups, and click **Close**.

6. Re-create the problem that you are troubleshooting. The Operations Center is automatically traced, and a new version of the Operations Center log is created.
7. Return to the list of logging groups by repeating step 1 on page 113 - step 3 on page 113.
8. Select all rows that are enabled, and click **Disable**.
9. Confirm that you want to disable the logging groups, and click **Close**.

What to do next

For the location and names of the Operations Center log files, see “Log files overview” on page 99.

Related tasks:

“Viewing the Operations Center log from within the Operations Center” on page 100

“Tracing the Operations Center by enabling functions in the logging configuration file”

Tracing the Operations Center by enabling functions in the logging configuration file

AIX

Linux

Windows

If the problem that you are troubleshooting prevents you from opening the Operations Center, you can open and modify the logging configuration file and start an extended trace that adds data to the Operations Center log.

About this task

In the following procedure, you enable groups of logging functions and start an extended trace.

Attention: Ensure that you disable the groups after the trace. Otherwise, the performance of the Operations Center can be affected.

Procedure

To trace the Operations Center, complete the following steps:

1. Stop the Operations Center web server.
2. Go to the following directory:

AIX

Linux

`installation_dir/ui/Liberty/usr/servers/guiServer`

Windows

`installation_dir\ui\Liberty\usr\servers\guiServer`

where *installation_dir* represents the directory in which Tivoli Storage Manager is installed.

3. Save a copy of the logging configuration file, `OpsCntrLog.config`, to another location for later use.
4. Open the original `OpsCntrLog.config` file in a text editor.
5. In the text editor, enable only the logging groups that IBM Software Support request that you enable, by replacing the word `OFF` with the word `ON` for each relevant group.
6. Save and close the file.
7. Start the Operations Center web server.

8. Re-create the problem that you are troubleshooting. The Operations Center is automatically traced, and a new version of the Operations Center log is created.
9. Stop the Operations Center web server.
10. Return to the `guiServer` directory.
11. Disable the logging groups by replacing the edited `OpsCntrLog.config` file with the previously saved copy.
12. Start the Operations Center web server.

What to do next

For the location and names of Operations Center log files, see “Log files overview” on page 99.

Related tasks:

“Tracing the Operations Center by enabling logging functions from within the Operations Center” on page 113

Enabling a trace for the server or storage agent

You can issue trace commands from the following places: the server console, storage agent console, administrative client connected to either the server or storage agent, server options file (`dsmserv.opt`), or the storage agent options file (`dsmsta.opt`).

Before you begin

Trace commands apply to the server or storage agent to which the command was submitted. Trace commands in the options files are used to trace the applications during startup and initialization or to provide a default set of trace classes. There is one trace class (**ADDMSG**) that is always enabled by default, whether it appears on the options file or not. It is best to trace to a file. Typically, the tracing for the server or storage agent will generate a large amount of output.

About this task

Perform the following steps to enable trace classes for the server or storage agent:

Procedure

1. Determine the trace classes to enable. To have trace messages issued for a given trace class, that trace class needs to be enabled either prior to beginning the trace or after the tracing has begun.
2. Issue the **TRACE ENABLE** *traceClassName* command to enable one or more trace classes. Note that *traceClassName* might be a space-delimited list of trace classes. For example, this command could be entered as **TRACE ENABLE TM SESSION**. The **TRACE ENABLE** command is cumulative, such that extra trace classes can be enabled by issuing **TRACE ENABLE** numerous times. For example, if you wanted to add the PVR trace class in addition to those that are already enabled, issue: **TRACE ENABLE PVR**. To stop having trace messages issued for a given trace class, that trace class needs to be disabled either prior to beginning the trace or after the tracing begins.
3. Issue the **TRACE DISABLE**<*traceClassName*> command to disable one or more trace classes. Note that *trace class name* might be a space delimited list of trace classes. For example, this command could be entered as **TRACE DISABLE TM SESSION**. Additional trace classes can also be disabled by issuing **TRACE DISABLE**.

For example, if you wanted to remove the PVR trace class in addition to those that were already disabled, issue: **TRACE DISABLE PVR**. By issuing **TRACE DISABLE** without specifying any trace classes, all currently enabled trace classes are disabled.

4. Tracing can occur to the console or to a file. Perform the following tasks to begin tracing:
 - For tracing to the console, issue: **TRACE BEGIN**
 - For tracing to a file with no size limitation, issue: **TRACE BEGIN fileName**
 - For tracing to a file with a size limitation, issue: **TRACE BEGIN fileName MAXSIZE= maximum size in megabytes**

Note: The *fileName* can be a fully-qualified path such as /opt/tmp or c:\temp. If a full path is not given, the trace file will be located in the same directory as the running executable file.

5. Perform the operation that is causing the problem.
6. Issue the **TRACE END** command to stop trace messages from being issued. If tracing is being done to a file, ending the trace writes any remaining trace messages to the file and closes the file.

What to do next

It is possible to enable tracing and begin it using the server or storage agent options file. The commands and syntax discussed are the exact same for the server or storage agent options file, and they are generally used to trace startup and initialization of the server. For example, if the following lines were added to the server's option file, tracing would be started for the DB, TM, and LOG trace classes, and the trace messages written to the file MYTRACE.OUT.

```
TRACE ENABLE DB TM LOG
TRACE BEGIN MYTRACE.OUT BUFSIZE=4096
```

Remember: If you are conducting a trace due to a server crash, do not set the **BUFSIZE** parameter.

Related reference:

“Trace classes for a server or storage agent” on page 117

Enabling a stack trace for 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.

Note: Depending on the frequency of the failure, stack trace can flood the activity log file, which can cause problems when you are trying to view the activity log file. You might want to disable stack trace after it completes.

IBM Software Support might find it helpful to enable stack trace on specific messages that are 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 that is connected to either the server or storage agent.

To get a stack trace when a specific message is issued by the server or storage agent, enable the message for stack trace. Issue the **MSGSTACKTRACE ENABLE <messageNumber>** command to enable one or more messages for stack trace.

Remember: *<messageNumber>* might be a space-delimited list of message numbers.

This command can be entered as **MS ENABLE 2017**. The **MSGSTACKTRACE ENABLE** command is cumulative, such that extra messages are enabled by issuing the **MSGSTACKTRACE ENABLE** command more times. If you want to add message 985, in addition to the messages 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 that are issued by the server or storage agent, the stack trace for these messages must be disabled. Issue the **MSGSTACKTRACE DISABLE <messageNumber>** command to disable one or more messages.

The *<messageNumber>* might be a space-delimited list of message numbers. For example, this command can be entered as **MSGSTACKTRACE DISABLE 2017 985**. Extra messages can also be disabled by issuing **MS DISABLE**. For example, if you want to remove message number 7837 in addition to the messages that are already disabled, issue the **MSGSTACKTRACE DISABLE 7837** command.

The following messages are enabled by default for stack trace.

435 437 486 661 685 727 728 780 781 782
784 785 786 790 793 794 860 881 882 883
884 1032 1078 1092 1117 1156 1227 5010 5015 5019
5021 5093 5099 5100 5267 6753 7823 7837 9600 9601
9602 9604 9605 9606 9607 9608 9999

Trace classes for a server or storage agent

The server and storage agent provide aggregate trace classes. These trace classes are a shortcut for using many related trace classes by specifying the aggregate trace class name for the **TRACE ENABLE** command.

The trace classes that are listed in Table 10 are those trace classes that are most typically requested or used for diagnosing problems. This table does not include all possible trace classes that are available. The trace class name is used with the **TRACE ENABLE** and **TRACE DISABLE** commands.

Table 10. Server or storage agent trace classes

Trace classes	Description	Usage
ADDMSG	Issues console messages such as ANR and ANE messages 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.
ADMCMD	Traces related to command processing.	Use this trace class to debug the command interpreter, including the PARALLEL and SERIAL command handling.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
AF	This trace class displays information about user data that is stored on sequential media devices. AF is an aggregate trace class that uses AFCREATE, AFMOVE, AFLOCK, AFTXN, and AFCOPY. Issue TRACE DISABLE AFLOCK unless the locking information is explicitly requested or needed.	Use this trace class to diagnose problems about reading or writing user files to sequential media volumes.
AFCREATE	This trace class displays information about storing user data on sequential media volumes.	Use this trace class to diagnose writing user data on sequential media volumes.
AFMOVE	This trace class displays operations that move user data with sequential media volumes. Move operations are completed by MIGRATION, RECLAMATION, MOVE DATA, and MOVE NODEDATA server processes.	Use this trace class to diagnose problems with the data movement server processes.
AS	This trace class displays information volume selection and assignment, coordination of drives (mount points), and management of data placement on volumes. This aggregate trace class uses ASALLOC, ASRTRV, ASDEALLOC, ASMOUNT, ASVOL, ASTXN, and ASSD. The typical method is to issue TRACE DISABLE ASTXN , unless the locking information is explicitly requested or needed.	Use this trace class to diagnose many different problems about volumes, mount points, or data read and write operations.
ASALLOC	This trace class displays information about reserving and allocating space on sequential media volumes for storing data. This space is for storing data on behalf of a client session or for server data movement operations such as MIGRATION, RECLAMATION, MOVE DATA, or MOVE NODEDATA.	Diagnose problems where the server or storage agent report no space available but space is supposed to be available in the storage hierarchy.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
ASDEALLOC	This trace class displays information about releasing and de-allocating space on sequential media volumes for storing data. Typical deallocation operations on the server are EXPIRATION, MIGRATION, RECLAMATION, MOVE DATA, MOVE NODEDATA, AUDIT VOLUME, DELETE VOLUME, and DELETE FILESPACE.	Use this trace class to diagnose during the deletion of data.
ASMOUNT	This trace class displays information about drive (mount point) selection and assignment for sequential media devices.	Diagnose situations where sessions or processes are waiting for mount points or cases where an operation fails because no mount point is available. Also useful in cases where a mount point is pre-empted.
ASRTRV	This trace class displays information about reading data from sequential media volumes.	Use this trace class to diagnose problems about data such as RESTORE or RETRIEVE client by the client, or MIGRATION, RECLAMATION, STORAGE POOL BACKUP, AUDIT VOLUME, GENERATE BACKUPSET, EXPORT, MOVE DATA, or MOVE NODEDATA by the server.
ASTXN	This trace class displays information about transactions that are used to make database updates to information for sequential media volumes, storage pools, device classes, and other attributes.	Use this trace class to diagnose stoppages, database operations, failures reported for sequential media operations, or general data storage problems.
ASVOL	This trace class displays information about volume selection and assignment for sequential media volumes.	Use this trace class to diagnose situations where sessions or processes are waiting for volumes, or cases where an operation fails because no volume is available. Also useful in cases where volume access is pre-empted.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
ASSD	This trace class displays information about sequential stream data operations. These operations use sequential media device classes, volumes, or mount points but do not store data in the storage hierarchy. Server processes that complete sequential stream data operations are BACKUP DB, EXPORT/IMPORT, and GENERATE BACKUPSET.	Use this trace class to diagnose server processes that complete sequential stream data operations.
BF	Information about user data (files) stored in the storage hierarchy. This aggregate trace class uses BFCREATE , BFRTV , BFSALVAGE , BFLOCK , BFAGGR , BFREMOTE , BFSAGGR , and BFTRG .	Use this trace class to diagnose general data read-or-write problems for client operations and server processes.
BFAGGR	This trace class displays information about server aggregation of user data. The server aggregates many smaller user files into a larger file in the storage hierarchy to optimize performance for data movement operations such as MIGRATION , MOVE DATA , and MOVE NODEDATA .	Use this trace class to diagnose general data read-or-write problems for client operations and server processes, or both.
BFCREATE	This trace class displays information about client operations that store data in the storage hierarchy. Typically, these client operations are BACKUP , ARCHIVE , or SPACE MANAGE operations by the client.	Use this trace class to diagnose failures or other problems while a client is storing data.
BFREMOTE	Traces the first stage of NDMP (Network Data Management Protocol) backup and restore processes.	This trace class is used to identify NDMP-related backup or restore operations. These trace classes are specific to the functions that implement the NDMP protocol. The SPID trace class provides more detailed tracing, including tracing all NDMP file history records that are sent by the NDMP file server.
BFRTV	This trace class displays information about client operations that read data from the storage hierarchy.	Use this trace class to diagnose failures or other problems while a client is reading data.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
BFSAGGR	This trace class displays information about the storing, retrieving, and moving of super aggregates. An object larger than 10 GB is stored as a super aggregate.	Use this trace class to diagnose problems that are related to storing or retrieving objects larger than 10 GB.
BITVECTOR	Diagnoses problems where the server reports problems with disk storage pools.	Use this trace class to display information about reserving and allocating space on volumes in disk storage pools.
BKSET/OBJSET	Trace class for backup set functions. The BKSET and OBJSET trace classes are synonymous.	Use this trace class to debug problems in the GENERATE BACKUPSET command or during a client restore operation from a backup set.
BLKDISK	Trace class for viewing disk I/O activity to storage pool, database, and log volumes.	Use this trace class to view I/O activity to disk to diagnose performance and disk I/O errors.
BRNODE	Trace class for the BACKUP and RESTORE NODE commands, which are used during NDMP operations.	Use this trace class to debug problems in the BACKUP and RESTORE NODE commands.
COLLOCATE	This trace class displays information about collocation processing on storage pools. COLLOCATEDDETAIL trace class can also be used to get more detailed information about the collocation processing. For example, information about the files that are being processed for a collocation group. Files that are being processed for a collocation group can cause many output trace statements.	Use this trace class to diagnose problems with collocation processing.
CRC	This trace class displays information about generating and managing cyclic redundancy checks (CRCs) on the server or storage agent. CRC is an aggregate trace class that uses CRCDATA , CRCPROTO , and CRCVAL .	Use this trace class to diagnose data corruption issues where CRC processing did not report data corruption.
CRCDATA	This trace class displays information about generating and managing CRCs for data that is 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.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
CRCPROTO	This trace class displays information about generating and managing CRCs for data that is exchanged between the client and either the server or storage agent where this node is configured with VALIDATEPROTOCOL=ALL or VALIDATEPROTOCOL=DATAOnly on the server.	Use this trace class to diagnose data corruption issues where CRC processing did not report data corruption.
CRCVAL	This trace class displays information about generating and comparing CRC values.	Informational for showing CRC values during processing.
CRYPTO	This trace class displays information about advanced encryption standard (AES) operations and some general encryption settings.	Use this trace class to isolate and identify encryption-related problems.
DBCLI	Traces the general set of interactions.	Use this trace class to trace the general set of DB2 interactions and the DB2 command-line interface.
DBCONN	Traces connection activities.	Use this trace class to trace IBM Tivoli Storage Manager connections to DB2 connections. This trace class shows such things as the creation of connection handles and the assignment of connections to transactions.
DBDBG	Traces debugging processes. You might use this trace class first when you are debugging a database issue.	Use this trace class to show function entry or exit, exit return codes, and the statements that are built and are being run.
DBITXN	Traces database transaction-related activities. Transaction-related activities concern transaction latch acquisition and release, dbTxnDesc allocation and release, and transaction commit processing from the prepare and commit phase functions.	Use this trace class to trace transaction-related activities for the database interface.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
DBNETDB	This trace class displays information about LAN-free operations and the negotiation and management of information between the server and storage agent.	Use this trace class to diagnose LAN-free problems when the server and storage agent are at different levels. They function better when they are at the same level. You can also use this trace class to diagnose problems with a storage agent that is obtaining configuration information from the server.
DBRC	Traces the return codes from functions in the database component.	Use this trace class to trace the return codes.
DEDUP	Traces the general logic path tracing for data deduplication processing. Does not typically include error paths.	Use DEDUP to trace general logic paths for data deduplication processing.
DEDUP1	Traces error paths for data deduplication processing.	Use DEDUP1 to trace error paths for data deduplication processing.
DEDUP2	Traces the fingerprinting and digital signatures path.	Use DEDUP2 to trace fingerprinting and digital signature paths.
DELTA	Trace class for logical group functions. The DELTA and GROUP trace classes are synonymous.	Use this trace class to debug problems with logical groups, whether delta-base groups (subfile backup) or peer groups (Windows SYSTEM OBJECT or image backups). Group processing is relevant during just about any operation that references backup objects. The backup objects can include client backup and restore, expiration, deletion (DELETE FILESPACE , DELETE VOLUME), export/import, backup set generation and restore, no-query restore, database audit, and others.
DF	This trace class displays information about user data that is stored on disk volumes. DF is an aggregate trace class that enables DFCREATE , DFRTRV , DFMOVE , DFLOCK , DFTXN , and DFCOPY . Issue the TRACE DISABLE DFLOCK command unless the locking information is explicitly requested or needed.	Use this trace class to diagnose problems about reading or writing user files to disk volumes.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
DFCREATE	This trace class displays information about storing user data on disk volumes.	Use this trace class to diagnose writing user data on disk volumes.
DFMOVE	This trace class displays operations that move user data by using disk volumes. Move operations are completed by the MIGRATION , MOVE DATA , and MOVE NODEDATA server processes.	Use this trace class to diagnose problems with the data movement server processes.
DFRTRV	This trace class displays information about reading user data from disk volumes.	Use this trace class to diagnose reading user data from disk volumes.
DS	This trace class displays information about volume selection, space reservation, assignment, and management of data placement on disk volumes. DS is an aggregate trace class that enables DSALLOC , DSRTRV , DSDEALLOC , and DSVOL . Issue TRACE DISABLE DSTXN unless the locking information is explicitly requested or needed.	Use this trace class to diagnose many different problems about disk volume data read-and-write operations.
DSALLOC	This trace class displays information about reserving and allocating space on disk volumes for storing data. The data storage might be completed on behalf of a client session or for server data movement operations such as MIGRATION , MOVE DATA , or MOVE NODEDATA .	Diagnose problems where the server or storage agent report that no space is available, but there appears to be space available in the storage hierarchy.
DSDEALLOC	This trace class displays information about releasing and de-allocating space on disk volumes. Typical deallocation operations on the server are EXPIRATION , MIGRATION , MOVE DATA , MOVE NODEDATA , AUDIT VOLUME , DELETE VOLUME , and DELETE FILESPACE .	Use this trace class to diagnose during the deletion of data.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
DSRTRV	This trace class displays information about reading data from disk volumes.	Use this trace class to diagnose problems about reading data such as RESTORE or RETRIEVE client by the client, or MIGRATION , STORAGE POOL BACKUP , AUDIT VOLUME , GENERATE BACKUPSET , EXPORT , MOVE DATA , or MOVE NODEDATA by the server.
DSVOL	This trace class displays information about volume selection and assignment for disk volumes.	Use this trace class to diagnose situations where sessions or processes are waiting for volumes, or cases where an operation fails because no volume is available.
ICVOLHST	Trace class for volume history functions.	Use this trace class to debug problems with creating volume history entries, for example; during EXPORT , BACKUP DB , or GENERATE BACKUPSET) or deleting volume history entries, for example; during DELETE VOLHISTORY).
IMFS	Trace class for file space functions.	Use this trace class to debug problems that are related to inventory file spaces (for example, during DELETE FILESPACE).
LANFREE	This trace class displays general information about LAN-free operations on either the server or storage agent. Also shows error information for LAN-free-related operations. LANFREE is an aggregate trace class that enables LNFVERB, LNFMEM, LNFENTRY, and LNFDATA.	Any LAN-free failure.
MMS	This trace class displays information about tape libraries and the server or storage agent that uses these libraries. MMS is an aggregate trace class that enables MMSBASE, MMSTXN, MMSLIB, MMSDRIVE, MMSOP, MMSMAN, MMSSCSI, MMSFLAG, MMSACSLs, and MMSSHARE. Include NA and PVR trace classes when you are tracing MMS.	Used to diagnose problems with tape libraries, library volume inventories, or other general library issues.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
MONITOR	This trace class displays information about alert monitoring.	Use this trace class to determine why an alert might not be generated.
NA	This trace class displays information about path information for the server or storage agent. This information relates to the DEFINE PATH, UPDATE PATH, DELETE PATH, and QUERY PATH commands. This trace class is also used to identify issues that are related to operations that involve NDMP file servers, for example; DEFINE DATAMOVER, UPDATE DATAMOVER, BACKUP NODE, and RESTORE NODE commands. This aggregate trace class uses NALOCK, NAPATH, NAMOVER, NADISK, and NACONFIG. It might be best to include MMS and PVR trace classes when you are tracing NA.	Use this trace class to diagnose problems with paths to devices.
PRODCONS	If there are problems with getting work dispatched to batches, PRODCONS displays information about the problem and whether it is in the PC object or in replication.	Use PRODCONS to trace the internal workings of the producer/consumers objects that are used in the server.
PROXYNODE	This trace class displays information about proxynode sessions and the commands that are 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 you are analyzing proxynode session problems.
PVR	This trace class displays information about sequential media devices and the server or storage agent use of these devices. PVR is an aggregate trace class that enables PVRVOL, PVRCLASS, and PVRMP. The PVR trace class contains everything in the PVRIO aggregate trace class and the PVRNOIO trace class.	Use this trace class to diagnose problems with tape drives, failures when reading or writing tape volumes, or other tape-volume-related issues.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
PVRIO	This trace class displays tracing of read, write, or POS operations for sequential media devices and the server or storage agent use of these devices.	Use this trace class to diagnose problems with tape drive failures when reading or writing tape volumes.
PVRNOIO	This trace class displays PVRVOL, PVRCLASS, and PVRMP information.	Use this trace class to diagnose problems with tape drive mounts or other tape-volume-related issues.
REPL	REPL is an aggregate trace class that enables REPLBATCH, REPLCMD, REPLFS, REPLINV, REPLPROC, REPLSTATS, and REPLSESS.	Use this trace class to diagnose problems with replication.
REPLBATCH	This trace class displays tracing related to batch processing, where individual files are sent from the source server to the target server.	Use this trace class to diagnose replication problems with batch processing.
REPLCMD	This trace class displays tracing related to command parsing and the resolution of file space replication rules.	Use this trace class to diagnose replication problems with command parsing and the resolution of file space replication rules.
REPLFS	This trace class displays tracing related to the iteration of the file spaces to decide what is to be replicated, updated, or deleted.	Use this trace class to diagnose replication problems with iterating file spaces to decide what is to be replicated, updated, or deleted.
REPLINV	This trace class displays tracing related to the inventory updates (IM tables) as part of replication.	Use this trace class to diagnose replication problems with inventory updates.
REPLPROC	This trace class displays tracing of the overall replication process. This trace class is the main thread and dispatcher.	Use this trace class to diagnose replication problems with the replication process.
REPLSESS	This trace class displays tracing related to the establishment of sessions for replication, including the session management on both source and target servers.	Use this trace class to diagnose replication problems with the establishment of sessions.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
REPLSTATS	This trace class displays tracing related to updating the statistics as replication runs. Also includes insertion or update of history records in the replication history table.	Use this trace class to diagnose replication problems with statistical updates.
RETPROT	Trace class for the archive retention protection functions.	Use this trace class to debug problems when you are using the RETINIT and RETMIN parameters in the archive copy group. You can also use this trace class for problems that are caused by using the VB_SignalObject verb (only supported by the client API) to signal an object's event or to hold or release an object. Finally, you can use this trace class for 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 that are related to schedule commands like DEFINE/UPDATE/QUERY SCHEDULE or DEFINE ASSOCIATION . Also use this trace class to debug problems that are related to the central scheduler background processes, such as the schedule manager and schedule prompter.
SESSION	This trace class displays information about sessions that are connected to the server, including all verbs that are sent and received by the server.	This trace class is used for protocol violations, transaction processing errors, or in cases where the client is stopped and not responding.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
SESSREMOTE	Traces communication between the server and the client during NDMP backup and restore operations.	This trace class is used to identify NDMP-related backup or restore operations that are initiated when you are using the Tivoli Storage Manager web or command-line client.
SHRED	This trace class displays information that is related to data-shredding operations on the server.	This trace class is used to diagnose problems with data shredding. Data shredding is only applicable if one or more storage pools on the server have a nonzero value for the SHRED attribute. Activity that is related to data shredding occurs primarily during the EXPIRE INVENTORY, DELETE FILESPACE, DELETE VOLUME, MOVE DATA, MIGRATE, and SHRED DATA commands. Other trace classes that report activity related to data shredding are BFDESTROY, DFDESTROY, DSALLOC, DSDEALLOC, and CRCDATA .
SPI/SPID	Traces the server NDMP protocol interface.	The SPI and SPID trace classes are used to identify issues that are 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 that are sent by the NAS file server.
SSLDATA	Detailed Secure Sockets Layer (SSL) trace is used to display byte-level information about data that is sent or received between the backup-archive client and the server.	Use the SSLDATA trace class to debug the session data corruption issues that might be caused by SSL that is running through the SSLTCP or SSLTCPADMIN server options. This trace is a byte-level trace so it can collect a large amount of data.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
SSLINFO	General SSL trace is used to display setup and characteristics of SSL sessions between the backup-archive client and the server.	Use the SSLINFO trace class to debug session connection and handshake errors that might be caused by the SSL that is running through the SSLTCP or SSLTCPADMIN server options. This trace class can be used in tandem with the TCPINFO and SESSION trace classes.
TBREORG	This trace class collects information about table and index reorganization activities that are initiated by the server.	Use the TBREORG trace class to debug server-initiated reorganization activity.
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	This trace class collects information about TCP/IP used between the client and either server or storage agent. TCP is an aggregate trace class. It enables TCPINFO and TCPERROR.	Use this trace class to debug session connection errors or data corruption issues that might be caused by the network.
TCPDATA	The detailed TCP/IP trace is 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	The general TCP/IP trace is 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	This trace class provides information about events that are sent to a TEC server. These events correspond to the 'Tivoli' event receiver.	To debug connection issues that are incurred by TEC event logging.
TOC	This trace class is used for the Table Of Contents (TOC) component, which is used during file-level NDMP operations. TOC is an aggregate trace class that enables TOCBUILD, TOCLOAD, TOCREAD, and TOCUTIL.	Use this trace class to debug problems during file-level NDMP operations, such as an NDMP backup with the TOC=YES parameter, or an NDMP restore with the FILELIST parameter.
TOCBUILD	Tables Of Contents (TOC) build functions.	Use this trace class to debug problems during an NDMP backup with the TOC=YES parameter.

Table 10. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
TOCLOAD	Table Of Contents (TOC) load functions.	Use this trace class to debug problems while you are displaying files and directories on the client graphical user interface (GUI).
TOCREAD	Table Of Contents (TOC) read functions.	Use this trace class to debug problems during a QUERY TOC command or while you are trying to load a TOC to display files and directories on the client GUI.
TOCUTIL	Table Of Contents (TOC) utility functions.	Use this trace class to debug problems that are related to TOC component initialization or TOC retention.
UNICODE	This trace class displays information about code page conversions and Unicode filespace operations.	Use this trace class to debug problems that are related to code page conversion problems or Unicode filespace problems.
XI	This trace class displays general information for the IMPORT and EXPORT commands.	Use this trace class to debug problems that are related to IMPORT and EXPORT commands.

Show commands for the server or storage agent

SHOW commands are unsupported diagnostic commands that are used to show 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 shows, 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 must be used only with the recommendation of IBM Software Support. The **SHOW** commands that are included here are a portion of the available **SHOW** commands.

Table 11. Server or storage agent SHOW commands

SHOW command	Description	Recommendation
AGGREGATE	Shows information about an aggregate object in the server storage hierarchy. The syntax is SHOW AGGRegate <i>aggrID-high 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 that are 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. If you have problems restoring or retrieving files, expiring or moving data, backing up primary storage pools, copying active data to active data pools, or auditing volumes, you might issue this command.
ASQUEUED	Shows the mount point queue. The syntax is SHOW ASQueued .	To use a drive, a client session or server process, you 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 is stopped and waiting for a mount point.
ASVOL	Shows assigned volumes. The syntax is SHOW ASVo1 .	As sequential media volumes are assigned for use by a session or a process, they are tracked in an in-memory list. You can view this list to determine the state of in-use volumes, and stoppages or deadlock situations where a session or process is stuck waiting for a volume or holding a volume and waiting for something else.

Table 11. Server or storage agent SHOW commands (continued)

SHOW command	Description	Recommendation
BFOBJECT	Shows 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 <p>The syntax is SHOW BFOBJECT.</p>	This command helps you determine the existence and attributes of a bitfile object in the server's storage hierarchy. If you have problems restoring, retrieving, expiring, or auditing the object, you might issue this command.
CMD DEDUPDELETEINFO	Shows the status of background deletion threads for dereferenced deduplicated objects.	Issue this command to check the status of the background deletion process for deduplicated objects. When a file is deleted or moved out of a deduplicated storage pool, the extents are queued to a background processor for attempted removal from the storage pool. This command is useful for checking the backlog of queued extents and the status of each deletion thread.
CONFIGURATION	The CONFIGURATION command is a summary SHOW command that actually issues many different show commands and queries. The syntax is SHOW CONFIGURATION .	Issue this command to provide general configuration and other information about the server to IBM service.
DB2CONNECTIONS	The DB2CONNECTIONS command shows the defined DB2 connections from the various connection pools. This command does not require any additional parameters. The syntax is SHOW DB2CONNECTIONS .	Issue this command to show how many DB2 connections are defined, in-use, and free in total and within a particular 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.
DBVARS	Shows database global attributes. The syntax is SHOW DBVARS .	Issue this command to view the current state and attributes of the server database.

Table 11. Server or storage agent SHOW commands (continued)

SHOW command	Description	Recommendation
DEDUPOBJECT	Shows data deduplication information for files. When you issue this command, you must specify the objectID parameter. Issue the SHOW VERSION command to determine the value of this parameter. The syntax is SHOW DEDUPObject .	Issue this command to show data deduplication information, such as: <ul style="list-style-type: none"> • The bit file ID for each extent • The owning bit file ID • The offset and length of the owning bit file • The digest type and value of the data deduplication object
DEVCLASS	Shows information about device classes. The syntax for this command is SHOW DEVClass .	Issue this command to show 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	Shows all backup group leaders for an object in the server inventory. The syntax is SHOW GROUPLeaders objID-high objID-low . <i>objID-high</i> and <i>objID-low</i> are the high-order and low-order 32-bit words of the 64-bit object ID of the object that is 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. If you have problems restoring, retrieving, expiring, or auditing the object, you might issue this command.
GROUPMEMBERS	Shows all backup group members for an object in the server inventory. The syntax is SHOW GROUPMembers objID-high objID-low . <i>objID-high</i> and <i>objID-low</i> are the high-order and low-order 32-bit words of the 64-bit object ID of the object that is 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. If you have problems restoring, retrieving, expiring, or auditing the object, you might issue this command.

Table 11. Server or storage agent SHOW commands (continued)

SHOW command	Description	Recommendation
INVOBJECT	Shows 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 that is being queried. The high-order word is optional; if not specified, a value of zero is assumed. The object can be a backup object, an archive object, a space-managed object, and so on.	Issue this command to determine the existence and attributes of an object in the server inventory. You might issue this command if you are having problems restoring, retrieving, expiring, or auditing the object. The INVOBJECT command reports the following items: <ul style="list-style-type: none"> • New information for archive retention protected objects. • Whether the archive object is in deletion hold. • Whether the object uses event-based retention.
LIBINVENTORY	Shows 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 it is left out, the command returns the inventory information for all libraries.	Issue this command if there is a problem with the library inventory information. The command shows current in-memory properties of the library inventory.
LIBRARY	Use the LIBRARY command to show 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 it is left out, the command returns information for all the libraries.	This command is useful to gather a quick view of all in-memory information about a library and its drives. This output can be gathered for any problem that is related to libraries or drives; for example, mounting problems.
LOCK	Shows 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.

Table 11. Server or storage agent SHOW commands (continued)

SHOW command	Description	Recommendation
MEMTREND	The MEMTREND command reports the memory that is used by the server, in megabytes. It is recorded at hourly intervals for the last 50 hours. This command is set in the server code. It is not configurable. The command also shows 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, a memory leak might have occurred. For the measurements to be valid, the measurement period (the last 50 hours) must be normal, steady state 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. This command helps to determine the server's memory usage trend.
MP	Shows mount points. The syntax is SHOW MP .	Issue this command to determine which volume is in-use by a mount point and other attributes for the assigned mount points. SHOW LIBRARY and SHOW DEVCLASS have useful complementary information with this command to show the current state of drives and current devclass mount point counts.
NASDEV	Shows the SCSI devices that are attached to a network-attached storage (NAS) file server that is associated with a NAS datamover definition. The syntax is SHOW NASDev .	Create a Network Data Management Protocol (NDMP) connection to the specified NAS file server and show the attached SCSI devices on the file server. This command requires a NAS node and datamover definition only.
NASFS	Show the file systems on a NAS file server that is associated with a NAS datamover definition. The syntax is SHOW NASFs .	Create an NDMP connection to the specified NAS file server and show the file systems that are defined on the file server. Any file systems that are shown might be backed up by IBM Tivoli Storage Manager. This command requires a NAS node and datamover definition only.

Table 11. Server or storage agent SHOW commands (continued)

SHOW command	Description	Recommendation
NASINFORMATION	Shows configuration information about the NAS file server that is associated with a NAS datamover definition. The syntax is SHOW NASInformation .	Create an NDMP connection to the specified NAS file server and show general configuration information that is retrieved from the file server. This command is useful for identifying basic communication problems with NAS file servers such as authentication errors. This command requires a NAS node and datamover definition only.
NASWORKLOAD	Shows the workload of NAS files 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 and backup and restore operations.
REPLICATION	Shows all known replication servers and their globally unique identifier (GUID) and all running replication processes. The processes might include the individual statistics of each file space and the status of each replication session.	Issue this command if replication is not progressing or if replication is not working correctly.
RESQUEUE	Shows the resource queue. The syntax is SHOW RESQueue .	Use the resource queue to monitor common resources on the server. If a resource is stopped or holding a resource for an unreasonable amount of time, the resource monitoring algorithms for the server cancel the resource user. This command is used to show information about transactions, locks, and other resources that are used by a storage agent on the database server that it is configured to use.
SESSIONS	Shows information about sessions that are 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 command is also useful in cases where a session is canceled and is still shown in the QUERY SESSION .

Table 11. Server or storage agent SHOW commands (continued)

SHOW command	Description	Recommendation
SLOTS	Shows the current state of the specified library's slot information; for example, which volumes are in the library and in which slot). The syntax is SHOW SLOTS <i>libraryName</i> .	The information that is shown is the information that is saved directly from the library hardware to in-memory values. This information can be used to determine whether this information is out-of-sync, incorrect, or if the values returned from the library hardware itself are incorrect. Alternatively, issue this command to determine the drive element numbers for a SCSI library if QUERY SAN is unavailable for a particular library.
SSPOOL	Shows information storage pools. The syntax is SHOW SSPool .	Issue this command to show the states and attributes of defined storage pools.
THREADS	Shows information about all threads that are known to the server. The syntax is SHOW THReads . Important: On some operating systems (as an example: HP), the information that is reported is obtained without serialization. On a busy system, information can be inconsistent, multiple threads might report holding the same mutex, or a thread might report that it is waiting on a mutex that is held by another thread that does not claim to hold it.	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 is stopped so that you can see whether there are threads waiting for resources that are held by another thread.

Table 11. Server or storage agent SHOW commands (continued)

SHOW command	Description	Recommendation
TOCSETS	Shows all Table Of Contents (TOC) sets known to the server. The syntax is SHOW TOCsets DELETE=setNum TOUCH=setNum . The DELETE parameter causes the specified TOC set number to be deleted. The TOUCH parameter updates the last used date of the specified TOC set number. A TOC set is retained for the TOC retention period that follows the last used date (see SET TOCRETENTION command).	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 to provide file and directory names to the client GUI. This command shows the status of the TOC set; for example, building or loading, and how much temporary database space is in use for each TOC set. You might issue this command if you are experiencing problems with an NDMP backup with the TOC=YES parameter, or have problems restoring files from an NDMP backup, or if TOC sets are being retained in the server database too long or not long enough.
TOCVARS	Shows information about the TOC component of the server. The syntax is SHOW TOCvars .	Issue this command to determine the status of the TOC component. You might issue this command if you are experiencing problems completing an NDMP backup with the TOC=YES parameter, or have problems restoring files from an NDMP backup.
TXNTABLE	Shows 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.

Table 11. Server or storage agent SHOW commands (continued)

SHOW command	Description	Recommendation
VALIDATE LANFREE	Validates whether the definitions are in place on the server so that a client can complete LAN-free data movement operations. In cases where these definitions are not present or are incorrect, it might be difficult to determine whether the LAN-free environment is configured correctly. The syntax is VALIDATE LANFREE nodeName storageAgent . 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 the storage pool is able 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 show object IDs.
VOLINUSE	Shows whether the volume specified is in the server's in-use list. The VOLINUSE command shows extra information that might be helpful, including whether the volume is pending removal from the in-use list. The syntax is SHOW VOLINUSE volumeName . If the volume must be removed from the in-use list, you can specify the following parameter to remove the volume from the list: SHOW VOLINUSE volumeName REMOVE=YES .	Issue this command to determine whether a volume is on the in-use list and, if necessary, to remove it from that list. Operations that are associated with this volume might fail if the volume is removed from the in-use list.

Enabling a trace for the Tivoli Storage Manager device driver

Tracing is available for the IBM Tivoli Storage Manager device driver. The Tivoli Storage Manager device driver can be traced from the server console, an administrative client, or from a shell running on the system where the device driver is installed.

The tracing instructions apply to the Tivoli Storage Manager device driver on all platforms where the device driver is supported. For devices that use device drivers other than the Tivoli Storage Manager device driver, the ability to trace and instructions on how to trace those device drivers is provided by the device vendor.

Related reference:

“Tracing from the server console” on page 141

“Tracing data from a command shell for AIX, Oracle Solaris, and Windows” on page 142

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 two drivers: one for library-autochanger devices and one for tape devices. You might choose which one you want to trace. The following syntax is for the command:

```
DDTRACE START [ LIBRARYDD | TAPEDD]
[flags=EE |, FULL |, SYSLOG | BASE ]
DDTRACE GET [ LIBRARYDD | TAPEDD]
DDTRACE END [ LIBRARYDD | TAPEDD]
```

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.

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 LIBRARYDD Starts tracing the library-autochanger.

DDTRACE START LIBRARYDD TAPEDD Traces both the library and the tape drives.

Whichever of these you use, specify the same ones for all commands in the start-get-end series.

The **FLAGS** parameter is optional and usually not required. The following values are for the **FLAGS** parameter:

- EE** Traces all device driver routine entries and exits.
- FULL** Turns on more debug tracing and provides more detail. Because the memory buffer size is fixed, however, fewer events are traced. Does not trace routine entry and exit points.

SYSLOG

On some platforms, **SYSLOG** directs the trace statements to be written to the system log in addition to the memory buffer. This offering is most useful in debugging kernel stoppages or in when the trace wraps in the memory buffer.

BASE **BASE** is the default and cannot be specified with any other flags. It is only used to turn off the **EE**, **FULL**, and **SYSLOG** flags without turning off trace.

Tracing data from a command shell for AIX, Oracle Solaris, and Windows

AIX

Solaris

Windows

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

The stand-alone `ddtrace` utility is installed in the devices directory, which is the same directory as the `mttest`, `lbtest`, and `optest` utilities. Its syntax and options are identical to the **DDTRACE** server command. For example:

```
$ ddtrace start librarydd tapedd flags=EE - Start tracing both the library and tape drivers, and get additional entry/exit trace.
```

```
$ ddtrace get librarydd tapedd - Get the trace from memory and write it to the file ddtrace.out.
```

```
$ ddtrace end librarydd tapedd - Stop tracing to memory.
```

The main use of this stand-alone utility is primarily for cases when the driver needs to be traced during the IBM Tivoli Storage Manager server initialization. The `ddtrace` utility writes the memory buffer to the “`ddtrace.out`” file in the current directory. If the file exists, it appends to the file and does not overwrite it.

Tracing to detect a code page conversion failure

The IBM Tivoli Storage Manager server uses operating system functions to convert between Unicode and the server code page. If the system is not set up correctly, the conversion fails.

Procedure

Perform the following steps to attain more information on the failure:

1. Begin tracing the UNICODE trace class.
2. Repeat the action that caused the error message to occur.
3. Check the server README file for any platform-specific requirements for language installation.
4. Ensure that the locales indicated by the problem code pages are installed and any requirements that are listed in the README file are installed.

Tracing data for the client

You can enable tracing on the client or client application programming interface (API) by altering the client options file.

About this task

Perform the following steps to enable tracing on the client or client API:

Procedure

1. Determine the trace classes to enable from the following table:

Trace Class Name	Description	When to use	Additional Notes
SERVICE	Display general processing information for the client.	Useful in many cases. Generally recommended for protocol violations, transaction processing errors, or in cases where the client is stopped and not responding.	
VERBINFO	Collect information regarding the client-server protocol used by IBM Tivoli Storage Manager.	To debug protocol violations, transaction processing errors, or in cases where the client is stopped and not responding.	
VERBDETAIL	Detailed information regarding the client-server protocol used by Tivoli Storage Manager. This displays internal memory buffers containing the verbs sent and received by the client.	To debug session data corruption issues that might be caused by the network.	This generates a large amount of output.

2. Enable the trace by adding the following text to the client options file:
`traceflag <trace class name>`.
Attention: `<trace class name>` might be a comma-delimited list of trace classes. For example, this text could be entered as `traceflag service,verbinfo,verbdetail`.
3. Configure trace to begin and issue the trace messages to a file by adding the following text to the client options file: `tracefile <file name>`.
4. Perform the operation that is causing the problem.

Tip: Tracing might also be configured and started by invoking the client from a command prompt and specifying the flags above. For example, `dsm -traceflags=service -tracefile=file.out`.

Client and Journal Daemon traceflags

To run journal-based backup, you must use the Journal Daemon process. This process is used to track file system changes and maintain change journal databases.

The Journal Daemon uses the same tracing mechanism as the client, but the trace settings are specified in the journal configuration file (`tsmjbbd.ini`) as follows:

```
[JournalSettings]
TraceFlags=all_jbb
;
; the following two settings allow tracefile segmentation
```

```

;
TraceMax=100
TraceSegMax=1
tracefile=tracefiles\trace.out

```

Journal Daemon specific trace settings:

- BTREEDB - low-level BTREE database base class
- CACHEDB - disk cache backup and Windows 2003 exclude cache processing
- DBPERF - low-level database operation performance
- DBSTATS - performance tracking of database query, insert/update, delete, and tree walk operations
- FILEOPS - internal database activity
- JBBCOMM - listening thread
- JBBDAEMON - process manager
- JBBFILEMON - file system monitor
- JBBDBACCESS - database controller thread
- JBBDBINFO - low-level database access
- JBBNPCOMM - named pipe communications
- JBBSERVICE - Windows platform-specific service tracing
- JBBVERBINFO - detailed verb information
- ALL_JBB - aggregate traceflag that includes all of the above settings

Trace Settings for the backup-archive client specified in dsm.opt:

- JOURNAL - journal based backup tracing

Client trace classes

The client provides individual and aggregate trace classes. Aggregate trace classes are a shortcut for enabling many related trace classes by simply specifying the aggregate trace class name. For the documented trace classes, there might be references to trace classes that are enabled as part of an aggregate trace class, but are not explicitly discussed on their own.

The trace classes in Table 12 are the trace classes that are usually requested or used for diagnosing problems. The trace class name must be used with the TRACEFLAG options in the dsm.opt file.

Table 12. Trace classes

Trace Class	Description	Recommendation
ALL_BACK	Displays general backup processing information for the client. Aggregate of TXN, INCR, POLICY, and PFM trace classes and implicitly included in the SERVICE trace class.	Use this trace class for problems that are 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 that are related to reading and writing data and obtaining file attribute information.

Table 12. Trace classes (continued)

Trace Class	Description	Recommendation
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 that are 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 that are 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 that are related to all aspects of NDMP backup and restore operations.
ALL_SESS	Displays all session and verb information that is sent between the client and the server. Aggregate of SESSION, VERBINFO, SESSVERB, VERBADMIN, and VERBDETAIL trace classes. All of the trace classes in this aggregate are implicitly included in the SERVICE trace class, except VERBDETAIL.	Use this trace class for problems that are 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 that relates 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 that are related to volume snapshots that are used in online image backup and open file-support operations.
ALL_WAS	Displays Web Application Server (WAS) processing information for the client. Aggregate of several WAS-related trace classes and implicitly included in the SERVICE trace class.	Use this trace class for problems that are 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.

Table 12. Trace classes (continued)

Trace Class	Description	Recommendation
CLIENTTYPE	Displays client type on each trace output line.	Use this trace class for tracing situations when more than one client component is involved, such as the client acceptor and the file system agent.
COMPRESS	Displays compression information. Part of the SERVICE trace aggregate.	Use this trace class to determine how much data is compressed on a per-file basis.
DELTA	Displays adaptive subfile backup processing information. Part of the SERVICE trace aggregate.	Use this trace class to determine errors in adaptive subfile backup and restore operations.
DIROPS	Displays directory read and write operations. Part of the SERVICE and ALL_FILE trace aggregates.	Use this trace class when problems occur in a read or write directory.
DOMAIN	Displays incremental domain processing information. Part of the SERVICE trace aggregates.	Use this trace class for determining how DOMAIN statements are resolved during backup processing, such as problems in resolving the ALL-LOCAL domain.
ENCRYPT	Displays data encryption information. Part of the SERVICE trace aggregate.	Use this trace class to determine whether 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 the error codes that are generated by the operating system.
FILEOPS	Displays file read and write operations. Part of the SERVICE and ALL_FILE trace aggregates.	Use this trace class when problems occur in a file open, read, write, or close operation.
FIOATTRIBS	Displays comparisons of file attributes between the local client version and the active version on the 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 to determine whether files are candidates for incremental backup, especially with the FIOATTRIBS trace class.

Table 12. Trace classes (continued)

Trace Class	Description	Recommendation
INCLEXCL	Displays include-exclude status for the object that is being processed. This flag is also used for the Preview function.	Use this trace class to determine which object (usually file or directory) is included or excluded during backup-archive/preview.
MEMORY	Displays memory allocation and free requests. This trace class writes a large amount of information into the trace file and is not included in any aggregate classes.	Use this trace class to determine memory leaks, memory spikes, and other memory-related problems.
OPTIONS	Displays current processing options. Part of the SERVICE trace aggregate.	Use this trace class to determine which options are in effect for the current session, and for problems in accepting processing options from server client-options sets.
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 server passwords from local storage, for example, PASSWORDACCESS=GENERATE errors.
PID	Displays process ID on each trace statement. Part of the SERVICE trace aggregate.	Use this trace class to diagnose problems that might involve multiple processes.
POLICY	Displays policy information available to the backup-archive client. Part of the SERVICE and ALL_BACK trace aggregates.	Use this trace class to see which policies are available during a backup or archive operation.
SCHEDULER	Displays general processing information for the scheduler. An aggregate that includes most of the client trace classes that are listed in this table. Aggregate of all trace classes except MEMORY, THREAD_STATUS, and *DETAIL classes.	Useful in many cases. This trace class is used for diagnosing scheduler problems when the nature of the problem is unknown. If the SCHEDULER trace flag is used, it generally is not necessary to specify any other trace flags because it already includes most of the basic trace classes.

Table 12. Trace classes (continued)

Trace Class	Description	Recommendation
SERVICE	Displays general processing information for the client. An aggregate that includes most of the client trace classes that are 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 with the SERVICE trace flag.	Useful in many cases. This trace class is used when the nature of the problem is unknown. If the SERVICE trace flag is used, it is not 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 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 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 you are diagnosing problems related to threading.
TXN	Displays transaction processing information. Part of the SERVICE and ALL_BACK trace aggregates.	Use this trace class when you are diagnosing problems related to transaction processing problems on the server, and for such problems as transaction stops and retries.
VERBDETAIL	Displays detailed verb information pertinent to client-server sessions. Part of the ALL_SESS trace aggregates.	Use this trace class to determine the contents of verbs that are sent between the client and server.

Table 12. Trace classes (continued)

Trace Class	Description	Recommendation
VERBINFO	Displays verb information pertinent to client-server sessions. Part of the SERVICE and ALL_SESS trace aggregates.	Use this trace class with the SESSION traceflag to give session context to general processing errors or to determine session problems like session timeouts and protocol violations.
WIN2K	Displays Windows system object or system state processing. Part of the SERVICE trace aggregates. Only valid on the Windows backup-archive client.	Use this trace class to determine errors with backup or restore of the system state information.

Enabling a backup-archive client trace

There are two methods of tracing that are available for the backup-archive client.

The first method is to configure trace parameters prior to starting the backup-archive client. The second is to enable tracing while the client is running. Choose which method of tracing to enable.

Enabling a client trace using the command line

You can trace the available backup-archive client by enabling client trace on the command line.

About this task

Complete the following steps to enable client tracing on the command line:

Procedure

1. Determine the trace classes to enable.
2. Choose which trace classes to enable by adding the following text to the `dsm.opt` client options file: `traceflags <trace class name>`
3. Use a minus sign (-) in front of a trace class to turn off tracing for a trace class. Make sure that the trace classes that have tracing turned off 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 can be entered as `traceflags service,verbdetail`

4. Choose the location of the trace messages output by adding the following text to the client options file: `tracefile <file name>`.

The *tracefile* name must be fully qualified, for example:

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

AIX **HP-UX** **Linux** **Solaris** `tracefile /home/spike/trace.out`

Mac OS X `tracefile trace.txt`

5. Set a maximum size for the trace file 1 - 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 that is specified with `tracemax`, "Continued at beginning of file" is written to the end of the trace file and tracing continues from the top of the file. The end of the tracefile is indicated with "END OF DATA." You can locate the end of the trace by searching for this string. If you specify a `TRACEMAX` size of 1001 or higher and `TRACESEGSIZE` is not specified, then the trace file is automatically split into multiple segments of 1000 MB per segment (see `TRACESEGSIZE` discussion).

You can choose to allow the client split the trace into smaller segments (1 - 1,000 MB per segment) by specifying the following variable in the client options file: `tracesegsize <trace segment size in MB>`

When trace is split into small segments, you to easily manage large amounts of trace data, avoiding the problems that are associated with compressing large files and eliminating the task of using a separate "file splitter" utility. For example, issue the following command to specify a trace segment size of 200 MB: `tracesegize 200`

A trace file segment name is specified with the `tracefile` option, plus an extension that indicates 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. When you are specifying the segment size, do not use any comma separators:

Correct: `tracemax 1000`

Incorrect: `tracemax 1,000`

If you use the `TRACESEGSIZE` option, the trace file segments are named by using the name that is specified in the option file with an extra extension using the segment number. For example, `trace.out.1`

6. Perform the operation that exhibits the problem.

What to do next

Tracing might also be configured and started by starting 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 144

Enabling a trace while the client is running

You can trace the available backup-archive client while the client is running.

Before you begin

- The backup-archive client must be installed to use dynamic tracing.
- The DSMTRACELISTEN YES option must be in effect when the client is started.
 - **AIX** **HP-UX** **Linux** **Solaris** This option is specified in the system options file (`dsm.sys`) in the stanza that the client uses. Users must be logged in as root to use `dsmtrace`.
 - **Windows** This option is specified in the client options file (usually `dsm.opt`). Users must be logged in as a member of the Administrators group.

When the client starts, it starts a separate “trace listener” thread. This thread “listens” on a named pipe, waiting to be contacted by the `dsmtrace` utility. To make the named pipe name unique, the client process ID (PID) is a part of the pipe name. When you use `dsmtrace` to configure tracing, it contacts the client through the named pipe on which the client is listening and passes to it the preferred trace configuration operation. The client then passes the results of the operation back to `dsmtrace` through another similarly named output pipe. `dsmtrace` displays the results to the console. The client starts the trace listener thread only when client option DSMTRACELISTEN YES is in effect. If DSMTRACELISTEN NO is in effect, then the listener thread is not started and dynamic tracing is not available to that client. DSMTRACELISTEN NO is the default value.

About this task

The steps for gathering a client trace are as follows:

Procedure

1. Stop the backup-archive client.
2. Configure the client options file with the preferred trace options.
3. Restart the backup-archive client and reproduce the problem.
4. Stop the backup-archive client.
5. Remove the trace options from the backup-archive client options file.
6. Send the resulting trace file to IBM technical support for analysis.

You can also use the `dsmtrace` utility to start, stop, and configure client tracing dynamically without having to stop the client or modify the options file. Dynamic tracing is especially useful when you must trace only the beginning of a long-running backup-archive client operation, or when you must start tracing after the backup-archive client is running for some time.

The `dsmtrace` utility includes the following features:

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

The following table summarizes the availability of this feature:

Table 13. Availability of the dsmtrace utility

Client Component	AIX, HP-UX, Linux, or Oracle Solaris Program Name	Windows Program Name
Backup-Archive Client (command line)	dsmc	dsmc.exe
Backup-Archive Client (GUI)	N/A	dsmagent.exe
Client acceptor	dsmcad	dsmcad.exe
Remote Client Agent	dsmagent	dsmagent.exe
Scheduler Service	N/A	dsmcsvc.exe
Journal Service	N/A	tsmjbbd.exe
Data Protection for Domino® (command line)	domdsmc	domdsmc.exe
Data Protection for Domino (GUI)	N/A	domdsm.exe
Data Protection for Microsoft Exchange (command line)	N/A	tdpexcc.exe
Data Protection for Microsoft Exchange (GUI)	N/A	tdpexc.exe
Data Protection for Microsoft SQL Server (command line)	N/A	tdpsqlc.exe
Data Protection for Microsoft SQL Server (GUI)	N/A	tdpsql.exe

Note:

- The center column in Table 13 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.

Example

The following example shows you how to enable client trace while the client is running:

1. Identify the process PID of the backup-archive client that you want to trace (make sure that DSMTRACELISTEN YES is in effect). Issue the following command to show all running instances of the client: `dsmtrace query pids`

Example output:

```
D:\tsm>dsmtrace query pids
```

```
IBM Tivoli Storage Manager
dsmtrace utility
dsmtrace Version 5, Release 3, Level 0.0
dsmtrace date/time: 10/24/2004 21:07:36
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.
```

```
PROCESS ID  PROCESS OWNER  DESCRIPTION                               EXECUTABLE NAME
4020        andy             Backup-Archive Client (CLI)                dsmc.exe
```

```
D:\tsm>
```

Important: Linux The threading model for some versions of Linux is to run each thread as a separate process, which means that when you query process information, you might see several processes for each instance of the client. The process that you must 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, issue the **PS** command to identify the parent dsmc process:

```
linuxppc:~ # ps -ef | grep dsmc
```

```
root 28967 1151 0 Oct22 pts/16 00:00:00 dsmc
root 28968 28967 0 Oct22 pts/16 00:00:00 dsmc
root 28969 28968 0 Oct22 pts/16 00:00:00 dsmc
root 28970 28968 0 Oct22 pts/16 00:00:00 dsmc
root 24092 24076 0 08:15 pts/93 00:00:00 grep dsmc
```

```
linuxppc:~ #
```

Notice that the parent for processes 28969 and 28970 is 28968. The parent for 28968 is 28967. The parent for 28967 is 1151, but the 1151 process does not appear in this display output. Process 1151 is the process that started 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 you are tracing an API application, the `-pipenameprefix` option must be included.

- AIX HP-UX Linux Solaris Use prefix `/tmp/TsmTraceTargetAPI`
- Windows Use prefix `\\.pipe\TsmTraceTargetAPI`

3. After sufficient trace data is collected, disable the tracing by issuing the following command:

```
dsmtrace disable 4020
```

Example output:

```
C:\program files\tivoli\tsm\baclient>dsmtrace disable 4020
```

```
IBM Tivoli Storage Manager
dsmtrace utility
  dsmtrace Version 5, Release 3, Level 0.0
  dsmtrace date/time: 10/24/2004 21:47:43
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.

ANS2802I Tracing has been disabled.
```

Other examples of enabling client trace while the client is running are defined in the following list:

dsmtrace query pids

This command displays all running processes whose names are listed in the table in the Background section.

dsmtrace query pids -filter=*

This command displays all running processes.

dsmtrace query pids -filter=dsm*

This command displays all running processes whose name begins with "dsm"

dsmtrace query pids -filter=dsm?

This command displays all running processes whose name begins with "dsm" plus one other character.

dsmtrace enable 2132 -traceflags=service -tracefile=c:\trace.txt

This command turns on SERVICE tracing for process 2132. Trace output is written to file c:\trace.txt.

dsmtrace enable 2132 -traceflags=-extrc

This command turns off extrc tracing for process 2132 (presumably tracing is already running for this process).

dsmtrace enable 4978 -traceflags=fileops -tracefile=/tmp/dsmtrace.out -tracemax=1000 -tracesegsize=200

This command turns on FILEOPS tracing for process 4978. The trace is written to files /tmp/dsmtrace.out.1, /tmp/dsmtrace.out.2, and so on, with each file being no larger than 200 MB. After 1000 MB are written, tracing wraps back to /tmp/dsmtrace.out.1.

dsmtrace query trace 4978 -on

This command displays basic trace information and lists trace flags that are turned on for process 4978.

dsmtrace disable 4978

This command disables tracing for process 4978.

dsmtrace disable 364 -pipenameprefix=/tmp/TsmTraceTargetAPI

This command disables tracing for API application process 364.

Known trace problems and limitations

The known problems and limitations of trace processes are gathered to help you resolve problems that you might encounter when you are running a trace process.

- If tracing is not currently active for a process and `dsmttrace` is used only with the `-TRACEFLAGS` option, for example, `dsmttrace enable 2346 -traceflags=service`, then you still see the following message:

```
ANS2805I Tracing has been enabled.
```

In this case, the trace flags were enabled, but tracing is not active until a trace file is specified by using the `-TRACEFILE` option.

- Do not use the `dsmttrace 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 IBM 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 Oracle Solaris only) a residual named pipe in the `/tmp` directory.

- **Windows** You must be logged in with a local administrative account 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 that begin 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 might not start. Any old FIFOs with process numbers that do not match the FIFOs 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 that you need to identify is the `dsmc` parent process.
- When multiple instances of the same program are running, you must identify the PID of the instance that you want to trace. In such a situation, information such as process information from the operating system might be available to help you identify the required PID. For example, if you want to trace `dsmc` that is being run by user 'andy' and there are two instances of `dsmc`, one owned by user 'andy' and the other owned by user 'kevin', you can use the process owner to identify which process to trace.
- If an options file contains a false option and the client does not start, you might see some named pipe errors in the `dsmerror.log` file. These error messages can be safely ignored.

Trace options

Trace has several options that you can employ.

DSMTRACEListen

DSMTRACEListen No | Yes

- No** The client does not start the trace listener thread and dynamic tracing is not available. The default is No.
- Yes** The client starts the trace listener thread and dynamic tracing is available.

Windows The DSMTRACEListen option is specified in the client options file (usually dsm.opt).

HP-UX The DSMTRACEListen option is specified in the system options file (dsm.sys) in the stanza that the client uses and cannot be specified from the command line.

dsmtrace

dsmtrace enable <pid> <options>

Use this command to start or modify tracing for a process.

pid The process ID (PID) for the client. Use dsmtrace query pids or your operating system facilities to identify the correct PID.

options

The client trace options.

dsmtrace disable <pid>[<options>]

Use this command to stop tracing for a process. The trace file closes and the trace flags, maximum trace size, maximum trace segment size, and trace file name are all cleared.

<pid> The PID for the client. Use **dsmtrace query pids** or your operating system facilities to identify the correct PID.

<options>

The client trace options.

dsmtrace help

This command displays basic syntax for dsmtrace.

dsmtrace query pids [-Filter=<spec>]

<spec>

The client process name filter specification, which can include the wildcard characters "?" (match exactly one character) or "*" (match zero or more characters).

If no filter is specified, then the default behavior is to display process information for any running instances of the program names listed in the table in the Background section above.

Important: **AIX** **HP-UX** **Linux** **Solaris** When using the FILTER, put the * symbol before and after the search text. This adjustment is necessary because the executable file name often includes the path in front of it, and in some cases, the executable file name might have additional characters at the end of it. For example:

- /opt/tivoli/tsm/client/ba/bin/dsmc
- domdsmc_DominoUserID

Thus, instead of `-filter=dsmc` or `-filter=domdsmc`, use `-filter=*dsmc*` or `-filter=*domdsmc*`.

dsmtrace query trace *<pid>* [*<options>*] [*<displayType>*] [-ALL | -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:

- ALL** Displays all trace flags and, for each flag, indicates whether it is turned on or off. The information shown with the `-BASic` display type is also included.
- ON** Displays the names of the trace flags that are turned on. The information shown with the `-BASic` display type is also included.
- Off** Displays the names of the trace flags that are turned off. The information shown with the `-BASic` display type is also included.
- BASic** Displays the name of the trace file and the maximum trace and trace segment sizes. This display type also indicates whether tracing is enabled or disabled.

-PIPENameprefix

-PIPENameprefix=*<pipeNamePrefix>*

The `-PIPENameprefix` option must be used when tracing application programming interface (API) applications:

- **AIX** **HP-UX** **Linux** **Solaris** Use prefix `/tmp/TsmTraceTargetAPI`
- **Windows** Use prefix `\\.\pipe\TsmTraceTargetAPI`

-TRACEFILE

-TRACEFILE=*<traceFileName>*

The `-TRACEFILE` option must specify a valid file name to which the trace is written. If tracing is already running, then this option has no effect.

-TRACEFLags

-TRACEFLags=*<traceFlags>*

Specify one or more trace flags. Typically, the trace flag `SERVICE` is used. Separate multiple trace flags with a comma. Trace flags can also be turned off by prefixing the flag name with a minus sign. When combining trace flags that you want to turn on, with trace flags that you want to turn off, put the flags that you want to turn off at the end of the list. For example, if you want to turn on `SERVICE` tracing except for `VERBDETAIL`, specify `-TRACEFLags=SERVICE,-VERBDETAIL`. If tracing is already running, then this option can be used to turn on additional trace flags or turn off trace flags.

-TRACEMax

-TRACEMax=<maximumTraceSize>

This option limits the maximum trace file length to the specified value (by default the trace file grows indefinitely). When the maximum length is reached, then the trace wraps back to the beginning of the file. Specify a value in MB between 1 and 4095. If tracing is already running, this option has no affect.

-TRACESegsize

-TRACESegsize=<maximumTraceSegmentSize>

This option is used when you anticipate a large trace file and you want the trace file to be written in smaller, more easily-managed segments. Each segment is no larger than the specified size. When this option is used, a segment number is appended to the trace file name for each segment. Specify a value in MB between 1 and 1000. If tracing is already running, this option has no affect.

Note:

- To turn tracing on for a process, you must use the -TRACEFLAGS and -TRACEFILE options (and -PIPENAMEPREFIX when tracing an API application).
- To modify trace flags for an existing process, use -TRACEFLAGS (and -PIPENAMEPREFIX when tracing an API application).
- If you need to modify the trace file name, maximum trace size, or maximum trace segment size, then you need to first disable tracing altogether (see the **dsmtrace disable** command).

Determining if data is encrypted or compressed during backup-archive by using a trace

You must perform several steps to determine whether the data during backup-archive is compressed or encrypted, or both.

Procedure

1. Add the trace options that are 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_256BIT ? "AES_256BIT" :
endSendObjExOut.encryptionType & DSM_ENCRYPT_AES_128BIT ? "AES_128BIT" :
endSendObjExOut.encryptionType & DSM_ENCRYPT_DES_56BIT ? "DES_56BIT" :
"NONE");
}
...

```

What to do next

For more information, see the *API Function Calls* in *Using the Application Programming Interface*.

Tracing data for the API

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

To enable tracing for the IBM Tivoli Storage Manager API, add the following lines to the `dsm.opt` file or another file designated as the client options file:

```

TRACEFILE trace file name
TRACEFLAGS trace flags

```

trace file name

The name of the file where you want to write the trace data.

trace flags

The list of trace flags to enable. Separate each trace flag by a space. The following trace flags are specific to the Tivoli Storage Manager API:

api Information about the API function calls

api_detail
Detailed information about the API function calls

You can also specify other backup-archive client and Tivoli Storage Manager API trace flags. Refer to the backup-archive client documentation for a list of any available trace classes. For example:

- TRACEFILE /log/trace.out
- TRACEFLAGS api api_detail verbinfo verbdetail time stamp

Important: If you do not have write permission for the file pointed by the TRACEFILE option, dsmSetup or dsmInitEx/dsmInit fail with return code DSM_RC_CANNOT_OPEN_TRACEFILE (426).

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

Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on an AIX or Linux system

AIX

Linux

By using Tivoli Monitoring for Tivoli Storage Manager, you can create and configure agent instances that monitor Tivoli Storage Manager servers. Stop all agent instances, modify the configuration files, and restart the agent instances to activate tracing of monitoring agents for servers on AIX or Linux.

About this task

Before you activate tracing, you can also open the Tivoli Enterprise Portal Agent log workspace and view the agent activities. The Agent Log workspace contains information about any Tivoli Storage Manager server that has an agent instance that is configured to monitor it. You can see output from the trace file without having to enable the trace file by using the Agent Log attribute group.

Complete the following steps to activate tracing:

Procedure

1. From a command line, go to the following directory:

```
cd install_dir/itm/bin
```

where *install_dir* is the monitoring agent installation directory. The default installation directory is /opt/tivoli/tsm/reporting. If you installed the monitoring agent on to an existing IBM Tivoli Monitoring server, go to the bindirectory. The default installation directory is /opt/IBM/ITM.

2. Stop monitoring agent instances by completing one of the following steps:
 - Stop the monitoring agents on the CandleManage graphical user interface by issuing the following commands:
 - a. Run the CandleManage program by issuing the following commands:

For example:

instancename15001111103143325000.log, and hostname1500DsmQuery.out

Tracing the Tivoli Monitoring for Tivoli Storage Manager agent on a Windows operating system

Windows

By using Tivoli Monitoring for Tivoli Storage Manager, you can create and configure agent instances that monitor Tivoli Storage Manager servers. To activate tracing of monitoring agents for servers running on Windows operating systems, stop all agent instances, modify the configuration file, and restart the agent instances.

About this task

Before activating tracing, you can also open the Tivoli Enterprise Portal Agent log workspace and the Agent Log attribute group, and view the agent activities. The Agent Log workspace contains information about any Tivoli Storage Manager server that has an agent instance configured to monitor it.

Complete the following steps to activate tracing:

Procedure

1. Stop the monitoring agent instances by completing the following steps:
 - a. On the Tivoli Monitoring server, click **Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services**.
 - b. Select each monitoring agent instance, right-click, and select **Stop**.

2. Locate the directory where the configuration file is stored:

```
install_dir\itm\tmaitm6
```

For example:

```
C:\IBM\itm\tmaitm6
```

3. To activate agent tracing, ensure that the following value is set in the `kskenv_agentInstance` file:

```
KSK_TRACE=1
```

4. The application programming interface (API) can also be traced, but is not required unless requested by a Tivoli support representative. To activate tracing for the API, ensure that the following value is set in the `kskenv_agentInstance` file:

```
KSK_APITRACE=1
```

5. Start the Tivoli Monitoring for Tivoli Storage Manager agent instances by completing the following steps:
 - a. On the Tivoli Monitoring server, click **Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services**.
 - b. Select each monitoring agent, right-click, and select **Start**.

Results

The trace results are located in the same directory as the configuration file:

```
install_dir\itm\tmaitm6\logs
```

The API trace results are located in the following directory:

install_dir\itm\tmaitm6

For example:

C:\IBM\itm\tmaitm6\logs

C:\IBM\itm\tmaitm6

The log file that contains the trace information has the format *aaaappptttt.log*, and the API trace has the format *aaaapppttttDsmQuery.out*, where:

aaaa is the agent instance name

pppp is the server port number

tttt is the timestamp

For example:

instancename15001111103143325000.log, and hostname1500DsmQuery.out

Chapter 8. Resolving data storage problems

If you are experiencing a problem in storing or retrieving data, several methods are available to help you resolve the problem.

Resolving unreadable data problems

You might receive unreadable data during import or node replication processes related to a lack of code page conversion during these processes.

If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and uses extended ASCII characters can be affected.

To resolve the issue, update the fields with the appropriate **UPDATE** commands after the import or node replication operation.

Checking the server activity log to resolve data storage issues

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

Issue the **QUERY ACTLOG** command to check the activity log. Often, other messages that are issued can offer additional information about the cause of the problem and how to resolve it.

Checking HELP for messages issued for a data storage problem

Check **HELP** for any messages issued by IBM Tivoli Storage Manager.

The Tivoli Storage Manager messages provide additional information in the **Explanation**, **System Action**, or **User Response** sections of the message. Often, this supplemental information about the message might provide the necessary steps necessary to resolve the problem.

Recreating the data storage problem

If a problem can be easily or consistently recreated, it might be possible to isolate the cause of the problem to a specific sequence of events.

Data read or write problems might be sequence-related, in terms of the operations being performed, or might be an underlying device error or failure.

Typical problems related to the sequence of events occur for sequential volumes. One example would be that a volume is in use for a client backup and that volume is preempted by a data restore from another client node. This situation might surface as an error to the client backup session that was preempted. However, that client backup session might succeed if it was retried or if it was not preempted in the first place.

Resolving data storage errors related to reading or writing to a device

If there is an error due to reading or writing data from a device, many systems and devices record information in a system error log file. For example, the `errpt` file for AIX and the Event Log file for Windows.

If a device or volume used by is reporting an error to the system error log file, it is likely a device issue. The error messages recorded in the system error log file might provide enough information to resolve the problem.

Changing the storage hierarchy to resolve data storage problems

The storage hierarchy includes the defined storage pools and the relationships between the storage pools on the server.

The storage pool definitions are also used by the storage agent. If attributes of a storage pool were changed, the change might affect data store and retrieve operations. Review any changes to the storage hierarchy and storage pool definitions. Issue the **QUERY ACTLOG** command to see the history of commands or changes that might affect storage pools. Also, use the following **QUERY** commands to determine if any changes were made:

- **QUERY STGPOOL F=D**

Review the storage pool settings. If a storage pool is **UNAVAILABLE**, then data in that storage pool cannot be accessed. If a storage pool is **READONLY**, then data cannot be written to that pool. If either situation is the case, review why these values were set and consider issuing the **UPDATE STGPOOL** command to set the pool to **READWRITE**. Another consideration is to review the number of scratch volumes that are available for a sequential media storage pool.

- **QUERY DEVCLASS F=D**

The storage pools can be influenced by changes to device classes. Review the device class settings for the storage pools, including checking the library, drive, and path definitions. Issue the **QUERY LIBRARY**, **QUERY DRIVE**, and **QUERY PATH** commands for sequential media storage pools.

Changing the server policies to resolve data storage problems

The server policy attributes that directly relate to data storage are the copy group destinations for backup and archive copy groups. Similarly, the management class, **MIGDESTINATION**, also impacts where data is stored.

Review any changes to the server storage policies. Issue the **QUERY ACTLOG** command to view the history of commands or changes that might affect storage policies. Also, use the following **QUERY** commands to determine if any changes were made:

- **QUERY COPYGROUP F=D**

Review the **DESTINATION** settings for the **TYPE=BACKUP** and **TYPE=ARCHIVE** copy groups. Also review the "Migration Destination" for management classes used by HSM clients. If storage pool destinations were changed and resulting data read or write operations are now failing, either evaluate the changes made and correct the problem or revert to the previous settings.

- **QUERY NODE F=D**

Assigning a node to a different domain might impact data read-and-write operations for that client. Specifically, the node might now be going to storage

pool destinations that are not appropriate, based on the requirements of this node. For example, it might be assigned to a domain that does not have any TYPE=ARCHIVE copy group destinations. If this node tries to archive data, it fails.

Resolving a data storage backup or copy problem that occurs only with a specific node

If you cannot backup or copy data to a specific node, you might not have an active data pool listed in your active destinations. These are specified in the node policy domain.

Issue the `QUERY NODE nodeName F=D` command to verify that the node that is storing the data is authorized. The `QUERY NODE` command finds the policy domain name to which the node is assigned. Issue the `QUERY DOMAIN domain_name` where *domain_name* is the output gathered from the previous `QUERY NODE` command. Look in the `ACTIVEDESTINATION` parameter for the list of active data ports. If the active data pool into which you want to store data is not on the list, issue the `UPDATE DOMAIN` command to add the active data pool to the list.

Resolving a data storage problem that occurs only for a specific volume

If problems occur only for a specific storage volume, there might be an error with the volume itself, whether the volume is sequential media or DISK.

If your operation is a data write operation, issue the `UPDATE VOLUME volumeName ACCESS=READONLY` command to set this volume to READONLY, then retry the operation. If the operation succeeds, try setting the original volume back to READWRITE by issuing the `UPDATE VOLUME volumeName ACCESS=READWRITE` command. Retry the operation. If the operation fails only when using this volume, consider issuing the `AUDIT VOLUME` command to evaluate this volume and issue the `MOVE DATA` command to move the data from this volume to other volumes in the storage pool. After the data is moved off of this volume, delete the volume by issuing the `DELETE VOLUME` command.

Hints and tips for storage

The hints and tips that are gathered here are from actual problem experiences. You might find that one of the solutions is right to fix your IBM Tivoli Storage Manager problem.

Device driver hints and tips

Device driver problems might be attributed to the operating system, the application using the device, the device firmware, or the device hardware itself.

Whenever a device problem is encountered, ask “Has anything been changed?”

If the adapter firmware changed, this change might cause a device to exhibit intermittent or persistent failures. Try reverting back to an earlier version of the firmware to see if the problem continues.

If cabling between the computer and the device was changed, this change often accounts for intermittent or persistent failures. Check any cabling changes to verify that they are correct.

A device might exhibit intermittent or persistent failures if the device firmware was changed. Try reverting back to an earlier version of the firmware to see if the problem continues.

For SCSI connections, a bent pin in the SCSI cable where it connects to the computer (or device) can cause errors for that device or any device on the same SCSI bus. A cable with a bent pin must be repaired or replaced. Similarly, SCSI buses must be terminated. If a SCSI bus is improperly terminated, devices on the bus might exhibit intermittent problems, or data that is transferred on the bus might be or appear to be corrupted. Check the SCSI bus terminators to ensure that they are correct.

Remember: If the “hints and tips” information does not adequately address your device driver issue or this is the initial setup of your system's device drivers, check that your hardware devices are supported. See the Support Portal.

Adjusting to operating system changes

Operating system maintenance can change kernel levels, device drivers, or other system attributes that can affect a device.

Similarly, upgrading the version or release of the operating system can cause device compatibility issues. If possible, revert the operating system back to the state prior to the device failure. If reverting is not possible, check for device driver updates that might be needed based on this fix level, release, or version of the operating system.

Adjusting to changes in the HBA or SCSI adapter connecting to the device

A device driver communicates to a given device through an adapter.

If it is a fibre channel-attached device, the device driver uses a host bus adapter (HBA) to communicate. If the device is SCSI attached, the device driver uses a SCSI adapter to communicate. In either case, if the adapter firmware was updated or the adapter itself was replaced, the device driver might have trouble using the device.

Work with the vendor of the adapter to verify that it is installed and configured appropriately. The following list shows the other possible steps:

- If the adapter was changed, try reverting back to the previous adapter to see if the issue is resolved.
- If other hardware in the computer was changed or the computer was opened, reopen the computer and check to make sure that the adapter is properly seated in the bus. By opening and changing other hardware in the computer, the cards and other connections in the computer might have loosened, which might cause intermittent problems or total failure of devices or other system resources.

Resolving a loose cable connection

If a connection is loose from the device to the cable, or from the cable to the device, problems with a device might occur.

Check the connections and verify that the cable connections are correct and secure.

For SCSI devices, check that the SCSI terminators are correct and that there are no bent pins in the terminator itself. An improperly terminated SCSI bus might result in difficult problems with one or more devices on that bus.

Resolving error messages in the system error log

A device might try to report an error to a system error log.

The following are examples of various system error logs:

- **AIX** errpt
- **Windows** Event Log

The system error logs can be useful because the messages and information that is 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 actions based on the messages that are issued to the error log.

Supporting 32-bit or 64-bit Linux kernel modules for 32-bit or 64-bit applications

Linux

The Linux kernel modules control the bit mode of the Linux SCSI generic device driver, all the different Host Bus Adapter (HBA) drivers, and other settings.

All of these kernel modules support only applications that have the same bit mode with running kernel modules. In other words, 64-bit kernel modules support only 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. A segmentation fault also happens if 64-bit application invokes a 32-bit kernel module on a 32-bit Linux system.

To avoid these segmentation faults, ensure that the bit mode of the Linux kernel module and its applications are the same by verifying that the 32-bit applications can invoke only 32-bit kernel modules on 32-bit Linux systems and 64-bit applications can invoke only 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 run only 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 run only on an AMD64/EM64T system with 32-bit Linux operating system.

A 64-bit Tivoli Storage Manager server that issues the **QUERY SAN** command requires a 64-bit host bus adapter (HBA) application programming interface (API) on an AMD64/EM64T system. If an AMD64 system is equipped with a Qlogic HBA, it might create a problem since, by default, Qlogic provides a 32-bit HBA API only on the AMD64 system. You must install the 64-bit HBA API on the system before you issue the 64-bit **QUERY SAN** command.

Adjusting to HBA driver changes on the Linux 2.6.x kernels

The most distinct change for HBA drivers on the Linux 2.6.x kernels is that all drivers have “ko” as a new suffix.

The following list shows the driver names and locations in 2.6.x kernels:

Adaptec

The driver (aic7xxx.ko) is located in the /lib/modules/kernel-level/drivers/scsi/aic7xxx/ directory.

Emulex

The driver (lpfcdd.ko) is located in the /lib/modules/kernel-level/drivers/scsi/lpfc/ directory.

Qlogic

Its driver names are qla2xxx.ko, qla2100.ko, qla2200.ko, qla2300.ko, qla2322.ko, and so on. There is a certain order to load the HBA drivers. The qla2xxx.ko is a base driver and should be loaded first. After loading the qla2xxx.ko driver, the system should then load the qla2300.ko driver if it is equipped with a Qla2300 card. All of drivers are located in the /lib/modules/kernel-level/drivers/scsi/qla2xxx/ directory.

Enabling multiple LUN support on Linux kernels

Linux

To configure SCSI devices with multiple LUNs on a Linux system, the Linux kernel must be set to enable multiple LUN support.

Multiple LUN support on some Linux distributions, however, is not a default option and requires users to manually add this option to the running kernel. Perform the following steps to set up and enable multiple LUN support on IA32 architecture:

1. Add one parameter to a boot loader configuration file.
 - For LILO boot loader:
 - a. Add append="max_scsi_luns=128" to the /ect/lilo.conf file.
 - b. Run lilo.
 - For GRUB boot loader:
 - a. Add max_scsi_luns=128 after the kernel image list at /etc/grub.conf file for RedHat distribution.
 - b. Add max_scsi_luns=128 after the kernel image list at /boot/grub/menu.1 file for SuSE distribution.
2. Restart the system.

Using Tivoli Storage Manager to perform a ddtrace on Linux or HP-UX

HP-UX

Linux

The passthru device driver can be traced by issuing the **DDTRACE** command.

To enable trace, issue the following commands from the server console or admin client:

- Linux:

```
trace enable lpdd <other server trace class names>
trace begin <file name>
```

- HP-UX:

```
trace enable pvrhppdd <other server trace class names>
trace begin <file name>
```

Select one of the following three options:

- ddtrace start librarydd tapedd (to trace both library and drive)
- ddtrace start librarydd (library trace only)
- ddtrace start tapedd (drive trace only)

Remember: **DDTRACE GET** and **DDTRACE END** are not required.

The IBM Tivoli Storage Manager passthru device driver trace cannot be enabled through the ddtrace utility.

Using the HP-UX passthru driver

HP-UX

If the autoconf utility did not claim your device, ensure that the correct drivers are loaded to the kernel.

For HP-UX 11i v2, the sctl driver, as well as the HP-UX stape, sdisk, and schgr native drivers are required for device configuration for the IBM Tivoli Storage Manager passthru device driver. For HP-UX 11i v3 on IA64, the esctl driver, as well as the HP-UX estape, esdisk, and eschgr native drivers are required for device configuration.

Refer to the HP-UX operating system documentation for information on how to load drivers to the kernel.

Related reference:

“Device driver hints and tips” on page 167

“Adjusting to changes in the HBA or SCSI adapter connecting to the device” on page 168

“Resolving a loose cable connection” on page 169

Updating the device information for host systems on a dynamic SAN without restarting

When devices in a SAN environment change, the information about this changed environment is not automatically sent to host systems attached to the SAN.

If the device information is not updated on the host systems that are attached to the SAN, the device paths that are previously defined no longer exist. If you use the existing device information to define device paths, backup, or restore data, these operations might fail. To avoid these types of failures, use a different method for each operating system to update the device information on the SAN without restarting host systems.

AIX Issue the **CFGMR** command to force the operating system to reconfigure itself. Then, run SMIT to reconfigure your IBM Tivoli Storage Manager devices.

HP-UX Run the **autoconf** utility with the **-f** option to issue the **IOSCAN** command, which forces the operating system to rescan SCSI buses and fibre channels.

Linux There is no system command to reconfigure the operating system. To rescan 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 the **autoconf** utility or the **TSMSCSI** command to reconfigure 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 14. HBA adapters and corresponding drivers for architectures Linux architectures

HBA adapters	HBA driver name	Available architectures
Adaptec 7892	aix7xxx	IA32, AMD64
Qlogic 22xx	qla2200	IA32, AMD64
Qlogic 23xx	qla2300	IA32, AMD64
Qlogic 2362	qla2362	EM64T
Emulex	lpfcdd	IA32, iSeries, pSeries

Setting the multiple LUN options to “on” for Adaptec SCSI and Qlogic Fibre-Channel HBA BIOS settings on Linux

By default, Adaptec SCSI adapters set the multiple logical unit number (LUN) option to “off” in their BIOS, which makes the SCSI adapter driver unable to probe a SCSI unit with multi-LUN properly.

Procedure

The multiple LUN option must be turned on. Complete the following steps to turn on the multiple LUN options:

1. Press the Ctrl and A keys at the same time.
2. Select **SCSI Device Configuration** in the **Configure/View Host Adapter Setting**.
3. Change No to Yes for Bios Multiple LUN support.

Turning on the tape enable option:

By default, Qlogic Fibre host bus adapters set the tape enable option as off in their BIOS, which affects the running of some SCSI commands on several SCSI tape devices.

Procedure

The tape enable option must be turned on. Complete 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 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 usually installed inside a specific computer and used for storage by a Tivoli Storage Manager server on that computer.

Disk subsystem

An external disk subsystem connects to a computer through a SAN (storage area network) or some other mechanism. Generally, disk subsystems are outside of the computer to which they are attached and might be in close proximity, or they might be located much farther away. These subsystems might also have some method of caching the input/output requests to the disks. If data is cached, despite a bypass cache request, which can occur on remote file systems and certain disk subsystems, input/output failures can result. The failures are due to a difference between the Tivoli Storage Manager tracking and what is actually resident in a file system. Remote file systems and disk subsystems exhibiting these characteristics are not supported. Disk subsystems often have their own configuration and management software. A disk subsystem must report the results synchronously.

The server might define hard disk drives and disk subsystems that are 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 then looks like another file on that drive or file system.

Bypassing cache during write operations

Database, recovery log, and storage pool volumes are opened with the appropriate operating system settings to require data write requests to bypass any cache and be written directly to the device.

By bypassing cache during write operations, IBM Tivoli Storage Manager maintains the integrity of client attributes and data. Bypassing the cache is required. If an external event, such as a power failure, causes the server or the computer where the server is installed to halt or break while the server is running, the data in the cache might or might not be written to the disk. If the Tivoli Storage Manager data in the disk cache is not successfully written to the disk, information in the server database or recovery log might not be complete. Also, data that was supposed to be written to the storage pool volumes might be missing.

Hard disk drives installed on the computer where the server is installed and running have less of an issue with bypassing cache. In this case, the operating system settings that are used when Tivoli Storage Manager opens volumes on that hard disk drive generally manage the cache behavior appropriately and honor the request to prevent caching of write operations.

Typically, the use and configuration of caching for disk subsystems is a greater issue because disk subsystems often do not receive information from the operating system about bypassing cache for write operations. Disk subsystems also might ignore this information when a volume opens. Therefore, the caching of data write operations might result in corruption of the server database or loss of client data, or both. The problems are dependent upon which Tivoli Storage Manager volumes are defined on the disk subsystem and the amount of data lost in the cache. Disk subsystems should be configured to not cache write operations when a Tivoli Storage Manager database, recovery log, or storage pool volume is defined on that disk. Another alternative is to use nonvolatile cache for the disk subsystem. Nonvolatile cache employs a battery backup or some other sort of scheme to allow the contents of the cache to be written to the disk if a failure occurs.

Moving existing data to other volumes prior to altering or moving the database

The size and location of IBM Tivoli Storage Manager storage pool volumes (files) can not change after they are defined and used by the server.

If the size is changed or the file is moved, internal information that Tivoli Storage Manager uses to describe the volume might no longer match the actual attributes of the file. If you need to move or change the size of a Tivoli Storage Manager storage pool volume, move any existing data to other volumes prior to altering or moving the database.

FILE directory mapping between storage agents and servers for shared files

IBM Tivoli Storage Manager servers and storage agents can access the same data in FILE device classes by defining a set of directories to 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 you issue 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 that are used in the FILE device class. The default location is the current working directory of the server at the time that the command is issued. You can specify the directories for AIX, Linux, Oracle Solaris, or HP-UX.

Do not specify multiple directories from the same file system because you can cause incorrect space calculations. For example, if the directories /usr/dir1 and /usr/dir2 are in the same file system, the space check counts each directory as a separate file system. The space check does a preliminary evaluation of available space during storage operations. If space calculations are incorrect, the server might 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 the following table for the file name extension that is created by the server for scratch volumes, depending on the type of data that is stored.

Table 15. File name extensions for scratch volumes

For scratch volumes used to store this data:	The file extension is:
Client data	.BFS
Export	.EXP
Database backup	.DBV

For each storage agent that shares **FILE** access, the **PATHS** defined to each **DRIVE** seen by the storage agent must provide access to the same set of directories. When the **PATHS** are defined, the directories for each storage agent must match in number and order 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, which can result in successful LAN-restores and mount failures for the LAN-free restore operations.

Tape drives and libraries hints and tips

Problems with tape drives and libraries might be with software on the computer trying to use the device, connections to the device, or the device.

Whenever a device problem is encountered, ask, “Has anything been changed?” Suspect anything on the computer trying to use the device. Or look at the device itself, especially if the device worked prior to a given change then stopped working after that change.

- If the adapter firmware changed, a device might exhibit intermittent or persistent failures. Try reverting back to an earlier version of the firmware to see if the problem continues.
- If cabling between the computer and the device was changed, intermittent or persistent failures might occur. Check any cabling changes to verify that they are correct.
- If the device firmware has changed, a device might exhibit intermittent or persistent failures. Try reverting back to an earlier version of the firmware to see if the problem continues.

Adjusting to operating system changes

Operating system maintenance can change kernel levels, device drivers, or other system attributes that can affect a device. Similarly, upgrading the version or release of the operating system can cause device compatibility issues.

If possible, revert the operating system to the state prior to the device failure. If you cannot revert the operating system, check for device driver updates that might be needed based on the fix level, release, or version of the operating system.

Adjusting to device driver changes

A device driver upgrade might result in a tape drive or library device not working. These issues can also occur as a result of the type of driver that you use.

When working with IBM libraries or drives, as opposed to using other vendor libraries and drives, the type of device driver that you choose is important. IBM libraries and drives should use the IBM device driver, while other vendor libraries and drives should use the IBM Tivoli Storage Manager device driver.

Revert to the previous (or earlier) version of the device driver to see if the problem was introduced by the newer version of the driver.

Adjusting to a replaced adapter or other hardware changes

A small computer system interface (SCSI) connection to the device uses a SCSI adapter. A fibre-channel (optical) connection to the device uses a host bus adapter (HBA).

In either case, the cause of the problem might be from a changed adapter or an open computer where other hardware was changed or fixed.

Remember: The connecting point for the device to the computer is known as an adapter. Another term for adapter is *card*.

See the following information to help you adjust to a replaced adapter or hardware:

- If the adapter was changed, revert to the previous adapter to see if the issue is resolved.
- If hardware in the computer was changed or the computer opened, check the computer to ensure that the adapter is properly seated in the bus. By opening and changing other hardware in the computer, the cards and other connections in the computer might have been loosened. Loosening the connections might cause intermittent problems or total failure of the devices or other system resources.

Resolving a loose cable connection

Problems might occur to the device if a connection is loose from the computer to the cable, or from the cable to the device.

Check the connections and verify that the cable connections are correct and secure.

For SCSI devices, check that the SCSI terminators are correct and that there are no bent pins in the terminator itself. An improperly-terminated SCSI bus might result in problems with one or more devices on that bus.

Using error messages to resolve a device malfunction

A device might report an error to a system error log where you can try to find the cause of the problem.

Examples of various system error logs are:

- `errpt` for AIX
- Event Log for Windows

The system error logs can be useful because the messages and information recorded might help to report the problem, or the messages might include recommendations on how to resolve the problem. Check the appropriate error log and take any recommended actions based on messages issued to the error log.

SAN hints and tips

Problems with a SAN (storage area network) might be with software on the computer trying to use the device, connections to the device, or the device.

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

To better understand how to diagnose problems with a SAN, review the following terminology and typical abbreviations:

Fibre channel

Fibre channel denotes a fibre-optic connection to a device or component.

Host bus adapter

A host bus adapter (HBA) is used by a given computer to access a SAN. An HBA is similar in function to a network adapter in how it provides access for a computer to a LAN (local area network) or WAN (wide area network).

SAN A SAN is a network of shared devices that are typically accessed using fibre. Often, a SAN is used to share devices between many different computers.

Know your SAN configuration

Understanding the SAN configuration is critical in SAN environments. Various SAN implementations have limitations or requirements on how the devices are configured and set up.

The three SAN configurations are point-to-point, arbitrated loop, and switched fabric.

Point-to-point

The devices are connected directly to the host bus adapter (HBA).

Arbitrated loop

Arbitrated loop topologies are ring topologies and are limited in terms of the number of devices that are supported on the loop and the number of devices that can be in use at a given time. In an arbitrated loop, only two devices can communicate at the same time. Data being read from a device or written to a device is passed from one device on the loop to another until it reaches the target device. The main limiting factor in an arbitrated loop is that only two devices can be in use at a given time.

Switched fabric

In a switched fabric SAN, all devices in the fabric will be fibre native devices. This topology has the greatest bandwidth and flexibility because all devices are available to all HBAs through some fibre path.

Ensuring that your HBA works with your SAN

The host bus adapter (HBA) is a critical device for the proper functioning of a SAN. The problems that might occur with HBAs range from improper configuration to outdated BIOS or device drivers.

For a given HBA, check the following items:

BIOS HBAs have an embedded BIOS that can be updated. The vendor for the HBA has utilities to update the BIOS in the HBA. Periodically, the HBAs in use on your SAN should be checked to see if there are BIOS updates that should be applied.

Device driver

HBAs use device drivers to work with the operating system to provide connectivity to the SAN. The vendor typically provides a device driver for use with their HBA. Similarly, the vendor provides instructions and any necessary tools or utilities for updating the device driver. Periodically the device driver level should be compared to what is available from the vendor and, if needed, updated to pick up the latest fixes and support.

Configuration

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

Related reference:

“HBA configuration issues”

HBA configuration issues

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

The HBA vendor typically provides information about the settings for your HBA and the appropriate values for these settings. Similarly, the HBA vendor might provide a utility and other instructions on how to configure your HBA. The following settings typically affect the use of IBM Tivoli Storage Manager with a SAN:

- Storage area network (SAN) topology

The HBA must be set appropriately, based on the currently used SAN topology. For example, if your SAN is an arbitrated loop, the HBA must be set for this configuration. If the HBA connects to a switch, this HBA port must be set to “point-to-point” and not “loop.”

With Tivoli Storage Manager SAN Device Mapping, you can complete SAN discovery on most of the systems 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 because of a restart or other reason.

Go to the Support Portal 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 must also be set to this same value. Or the HBA must 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.

AIX **Linux** **Solaris** To help with problem determination, you can use the `dsmsanlist` module to obtain information about devices on a storage area network (SAN). The `dsmsanlist` module is installed by default when the Tivoli Storage Manager server or the Tivoli Storage Manager storage agent is installed.

Fibre-channel switch configuration issues

A fibre-channel switch supports many different configurations. The ports on the switch must be configured appropriately for the type of SAN that is set up and for the attributes of the SAN.

The vendor for the switch usually provides information about the appropriate settings and configuration based upon the SAN topology being deployed. Similarly, the switch vendor should provide a utility and other instructions on how to configure it. The following settings typically affect how IBM Tivoli Storage Manager uses a switched SAN:

Fibre-channel link speed

In many SAN topologies, the SAN is configured with a maximum speed. For example, if the fibre-channel switch maximum speed is 1 GB/sec, the host bus adapter (HBA) should also be set to this same value. Or the HBA should be set for automatic (AUTO) negotiation if the HBA supports this capability.

Port mode

The ports on the switch must be configured appropriately for the type of SAN topology being implemented. For example, if the SAN is an arbitrated loop, the port should be set to `FL_PORT`. For another example, if the HBA is connected to a switch, the HBA options should be set to “point-to-point” and not “loop.”

Data gateway port settings

A data gateway in a SAN translates fibre channel to SCSI for SCSI devices attached to the gateway.

Data gateways are popular in SANs because they allow the use of SCSI devices, therefore it is important that the port settings for a data gateway are correct.

The vendor for the data gateway usually provides information about the appropriate settings and configuration based upon the SAN topology being deployed and SCSI devices used. Similarly, the vendor might provide a utility and other instructions on how to configure it. The following settings can be used for the fibre channel port mode on the connected port on a data gateway:

Private target

Only the SCSI devices attached to the data gateway are visible and usable from this port. For the available SCSI devices, the gateway simply passes the frames to a given target device. Private target port settings are typically used for arbitrated loops.

Private target and initiator

Only the SCSI devices attached to the data gateway are visible and usable from this port. For the available SCSI devices, the gateway simply passes the frames to a given target device. As an initiator, this data gateway might

also initiate and manage data movement operations. Specifically, there are extended SCSI commands that allow for vendor-acquired data movement. By setting a given port as an initiator, it is eligible to be used for vendor-acquired data movement SCSI requests.

Public target

All SCSI devices attached to the data gateway, as well as other devices available from the fabric, are visible and usable from this port.

Public target and initiator

All SCSI devices attached to the data gateway as well as other devices available from the fabric are visible and usable from this port. As an initiator, this data gateway might also initiate and manage data movement operations. Specifically, there are extended SCSI commands that allow for vendor-acquired data movement. By setting a given port as an initiator, it is eligible to be used for vendor-acquired data movement SCSI requests.

SAN configuration between devices

Devices in a SAN, such as a data gateway or a switch, typically provide utilities that display what that device sees on the SAN.

It is possible to use these utilities to better understand and troubleshoot the configuration of your SAN. The vendor for the data gateway or switch typically provides a utility for configuration. As part of this configuration utility, there is usually information such as:

- How this device is configured
- 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 and the SCSI devices that are available in the SAN.

Switch

A switch reports information about the SAN fabric.

The fibre-channel link error report

Most SAN devices provide monitoring tools that can be used to report information about errors and performance statistics.

The vendor for the device should provide a utility for monitoring. If a monitoring tool is available, it typically reports errors. The following errors are experienced more frequently:

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

These errors are recoverable, where the error handling is usually provided by firmware or hardware. In most cases, the method to recover the device is to retransmit the failing frame. The fibre-channel link is still active when these errors are encountered. Applications using a SAN device that encounter this type of link error usually are not aware of the error unless it is a solid error. A solid error is one where the firmware and hardware recovery cannot successfully retransmit the data after repeated attempts. The recovery for these type of errors is typically very fast and will not cause system performance to degrade.

Link failure (loss of signal, loss of synchronization, NOS primitive received)

This error indicates that a link is actually “broken” for a period of time. It is likely due to a faulty gigabit interface connector (GBIC), media interface adapter (MIA), or cable. The recovery for this type of error is disruptive. This error appears in the application using the SAN device that encountered this link failure. The recovery is at the command exchange level and involves the application and device driver having to perform a reset to the firmware and hardware, which causes the system to run degraded until the link recovery is complete. These errors should be monitored closely, as they typically affect multiple SAN devices.

Remember: Fibre-channel link errors are often caused by a customer engineer (CE) action to replace a SAN device. As part of the maintenance performed by the CE to replace or repair a SAN device, the fibre cable might be temporarily disconnected. If the fibre cable is disconnected, the time and duration of the error should correspond to when the service activity was performed.

Common SAN device errors

Several SAN-specific messages can be issued when you experience problems with your storage agent SAN devices.

See Table 16 for errors that are generated for SAN devices.

Table 16. Common SAN device errors

Error	Explanation
ANR8302E I/O error on drive <i>TSMDRIVE01 (/dev/mt9)</i> (OP=WRITE, Error Number=5, CC=205, KEY=FF, ASC=FF, ASCQ=FF, SENSE=**NONE**, Description=General SCSI failure). Refer to Appendix D in the 'Messages' manual for appropriate action	<p>This message is often issued for SAN device errors. The CC=205 reports that the device driver detects a SCSI adapter error. If a SAN-attached device encounters a link reset that is caused by link loss, it is 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 might 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 that 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 error.</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 error.</p>

Table 16. Common SAN device errors (continued)

Error	Explanation
<p>ANR8963E: Unable to find path to match the serial number that is 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 was removed or replaced in the SAN. The following steps might resolve this error:</p> <ul style="list-style-type: none"> • Device Removed <p>If the device was removed from the SAN, delete the server definitions that refer to this device. Issue the <code>QUERY PATH F=D</code> command to determine any paths that reference the device. Then issue the DELETE PATH command to remove these paths.</p> • Device Replace <p>A SAN Device was replaced with a new device as a result of maintenance or an upgrade. Perform the following procedures:</p> <ul style="list-style-type: none"> – Try not to delete the drive or drive path definition after you replace the drive. – Issue one of the following server commands: <ul style="list-style-type: none"> - <code>UPDate DRive <libraryName> <driveName> SERIAL=AUTODetect</code> <p>This command force-records the new serial number into the server database. Because the drive is replaced, the element number stays the same.</p> - <code>UPDate PATH <sourceName> <driveName> SRCT=SERVER DESTT=DRIVE LIBRARY=<libraryName> DEVIce=xxxxx AUTODetect=Yes</code> <p>This command force-records the new serial number into the database. Because the drive is replaced, the element number stays the same.</p> – If the drive or drive path is deleted, redefine this new, replaced drive. You must restart the Tivoli Storage Manager server so that the element number/serial number map for the library is refreshed. This mapping occurs only at initialization. <p>Issue the QUERY PATH F=D command to find any paths that are defined on the server that reference this device, then issue the following command to update the path information:</p> <pre>UPDATE PATH AUTODetect=Yes</pre>
<p>ANR8972E: Unable to find element number for drive <i>driveName</i> in library <i>libraryName</i></p>	<p>If the ELEMent parameter is set to <code>AUTODetect</code> when defining the drive, Tivoli Storage Manager automatically gets the drive's element number. However, if the library does not provide an element number/serial number map, this message is issued.</p> <p>Perform the following steps to correct this error:</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.

AIX
Linux
Solaris
 To help with problem determination, you can use the `dsmsanlist` module to obtain information about devices on a storage area

network (SAN). The `dsmsanlist` module is installed by default when the Tivoli Storage Manager server or the Tivoli Storage Manager storage agent is installed.

Related concepts:

“SAN device mapping errors” on page 184

SAN device mapping hints and tips

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

The following items illustrate the advantages of IBM Tivoli Storage Manager SAN device discovery and device mapping:

Tivoli Storage Manager can display all the devices on the SAN

The **QUERY SAN** server command shows all the devices that are seen by the server via all the Fibre Channel host bus adapters (HBAs) installed on the system. The parameters that are shown are device type, vendor name, product model name, serial number, and the device name. If `FORMAT=DETAIL` is specified for the query, additional information such as World Wide Name (WWN), port, bus, target, and LUN are displayed. This information helps identify all the tape, disk, and data mover devices on the SAN. For AIX, the data mover is not apparent and is not shown.

Tivoli Storage Manager can update the device path automatically when a device's path changes

Tivoli Storage Manager does not require persistent binding for the devices it sees through the HBA. Instead, the server uses the SNIA (Storage Networking Industry Association) HBAAPI to discover and obtain the serial number for all the devices on the SAN. It can also determine each device's path. By comparing a device's serial number that is recorded in the Tivoli Storage Manager database with the serial number obtained from the device in real time, a change in a device's path is detected. If the path was changed, SAN discovery automatically obtains the new path for the device. The Tivoli Storage Manager database is also updated with the new path information.

The HBAAPI wrapper library is the wrapper that is used by the 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 included with the server package (except on AIX):

- **Windows** `hbaapi.dll`
- **AIX** `/usr/lib/libhbaapi.a` (provided by AIX with HBAAPI installation)
- **Linux** 32-bit: `libhbaapi32.so`
- **Linux** 64-bit: `libhbaapi64.so`
- **Solaris** 32-bit: `libhbaapi32.so`
- **Solaris** 64-bit: `libhbaapi64.so`

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

Disabling SAN device mapping:

Occasionally you must disable SAN device mapping to circumvent or isolate a problem when you are troubleshooting device problems.

About this task

Perform the following step to disable SAN device mapping and device discovery:

Procedure

Issue the **setopt SANDISCOVER OFF** server command. The **setopt SANDISCOVERY** commands can be issued as many times as needed.

Tip: Another way to disable/enable SAN discovery is to enter the following option in the `dsmserv.opt` file:

SANDISCOVERY OFF disables SAN discovery.

SANDISCOVERY ON enables SAN discovery.

SANDISCOVERY ON is the default for the AIX, Linux, Oracle Solaris, and Windows platforms.

Platform-specific information:

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

AIX The **QUERY SAN** command does not show any Gateway devices because Gateway devices are not apparent 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. Ensure that you load the Emulex driver before you load the ioctl module.

Tip: See the list of supported HBAs and required driver levels by operating system.

SAN device mapping errors

The errors that are most often generated during SAN device mapping can be related to SAN discovery, SAN device malfunction, libraries that are not valid, and other SAN-related issues.

ANR1745I: Unable to discover SAN devices. Function is busy.

This error message appears if there is another active SAN discovery.

The server is not able to perform SAN discovery. Try again after the other SAN discovery is completed.

ANR1786W, ANR1787W, or ANR1788W

You might see error messages ANR1786W, ANR1787W, or ANR1788W when there is 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 server is unable to perform SAN discovery, go to the Support Portal 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 HBAAPI error on the SAN.

“Get HBA target mapping failed” means that the HBA encountered an error while gathering device mapping information by sending various SCSI commands.

Verify that all SAN devices are working properly (for example, a SAN Data Gateway might be hung and might need rebooted). If all devices appear functional, verify the firmware of device on the SAN, and the HBA driver, are at the appropriate levels. If the result is that the server is not able to perform SAN discovery, go to the Support Portal to verify that the HBA driver is up-to-date and at a supported level.

Tip: For IBM tape devices, make sure that the latest firmware is installed. Firmware before 4772 for IBM 3580 tape devices causes problems with Qlogic HBAAPI.

ANR1790W SAN discovery failed

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

Verify that all SAN devices are working properly (for example, a SAN Data Gateway might be hung and might need rebooted). If all devices appear functional, verify that the firmware of device on the SAN, and the HBA driver, are at the appropriate levels.

Tip: For IBM tape devices, make sure that the latest firmware is installed. Firmware before 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 server to communicate with the SNIA HBAAPI.

The HBAAPI wrapper libraries are in the same directory as the Tivoli Storage Manager executable file (unless the full path is given as shown below). The following list shows the HBA wrapper files that are shipped with the server package (except on AIX and Linux zSeries). Error message ANR1791W indicates that the HBAAPI wrapper file is either missing or might 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)
- **Linux** 32-bit: libhbaapi32.so

- **Linux** 64-bit: libhbaapi64.so
- **Solaris** 64-bit: libhbaapi64.so

The result is that the server is not able to perform SAN discovery.

ANR1792W HBAAPI vendor library failed to load or is missing

Error message ANR1792W indicates that the vendor's library file failed to load. Verify the validity of the library files.

AIX, HP-UX, Linux, or Oracle Solaris systems (except on Linux zSeries) store their HBAAPI libraries in the location that is specified by the `/etc/hba.conf` file. Windows files are stored in the `C:\winnt\system32` directory. The following examples are of vendor library files:

- `C:\winnt\system32\qlsdm.dll` (QLogic's Windows file)
- `/usr/lib/libHBAAPI.a` (Emulex's AIX file)
- `/usr/lib/libqlsdm.so` (Qlogic's Linux file)
- `/usr/lib/libemulexhbaapi.so` (Emulex's Linux 32-bit file)
- `/usr/lib64/libemulexhbaapi.so` (Emulex's Linux 64-bit file)
- `/usr/lib/libqlsdm.so` (Qlogic's Oracle Solaris file)
- `/opt/JNIsnia/Solaris/Jni/64bit/JniHbaLib.so` (JNI's Oracle Solaris file)

The result is that the 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 operating systems 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 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 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 server is not able to correct the device path automatically if the path was changed. This command only has to be issued one time.

The difference between the two commands is that **SANDISCOVERY OFF** polls the device and marks the inactive path offline. **SANDISCOVERY PASSIVE** does not poll the device and does not mark the inactive path offline.

setopt SANDISCOVERY ON

This command enables the SAN discovery. The **SETOPT SANDISCOVERY ON** command can be issued as many times as necessary.

Another way to disable/enable SAN discovery is to put the following option in the `dsmserv.opt` file:

SANDISCOVERY OFF or SANDISCOVERY PASSIVE

These two commands can disable the SAN discovery.

SANDISCOVERY ON

This command enables the SAN discovery.

AIX **Linux** **Solaris** **Windows** **SANDISCOVERY** is defaulted to ON.

Go to the Support Portal to verify the platform/HBA vendor/driver level support level before setting **SANDISCOVERY ON** to enable the SAN discovery.

AIX **Linux** **Solaris** To help with problem determination, you can use the `dsmsanlist` module to obtain information about devices on a storage area network (SAN). The `dsmsanlist` module is installed by default when the server or the storage agent is installed.

ANR2034E QUERY SAN: No match found using this criteria

Error message ANR2034E is issued when the server tries to collect configuration information for the SAN and finds nothing.

The result is that the server is unable to perform SAN discovery.

The following are possible reasons for not finding information about the SAN:

- The system or OS level is unsupported.
- This environment is not a SAN environment.
- There might be a problem with the SAN.
- HBAAPI might return the zero value of the number of HBAs on the system.
- HBAAPI might 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 that 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 reinstalling it corrects the problem.
- Check the fibre-channel cable connection to the HBA.
- Check the fibre-channel cable connection from the HBA to the SAN device (switch, data gateway, or other device).
- Check the Gigabit Inter-phase Converter (GBIC).

- On the SAN device (switch, data gateway, or other device) try a different target port. Sometimes the SAN devices might have a specific port failure.
- Halt the server, restart the system, and restart the server. If there were configuration changes in the SAN, sometimes the operating system, device driver, or HBA requires a system restart before they can communicate with the SAN.
- Recycle the destination port on the SAN device.
- Re-seat the HBA card.
- Replace the HBA.

ANR8226E Error detecting version of HBA-API library

Error message ANR8226E is only displayed for AIX.

The server attempted to determine the level of the devices.common.IBM.fc.hba-api fileset and encountered an error. Error message ANR8226E indicates that an error occurred while trying to detect the HBA-API libraryFileset version on AIX.

The result is that the server is not able to perform SAN discovery.

AIX

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

Due to problems in AIX HBA-API code, the minimum fileset devices.common.IBM.fc.hba-api levels 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 server is not able to perform SAN discovery.

Related reference:

“SAN device mapping hints and tips” on page 183

SAN devices are missing from the display of QUERY SAN server command:

The possible reasons for the **QUERY SAN** server command not showing all the devices can be due to configuration or vendor support issues.

Ensure that the SANDISCOVERY server option set to ON.

Refreshing the SAN configuration:

The **QUERY SAN** server command might not be displaying all the devices because of the SAN configuration.

You might have to refresh the SAN because the configuration was changed (add/remove device) and the system configuration needs to be updated.

Update configuration on AIX:

For IBM devices:

Issue the **cfgmgr** command to configure new devices and view the new configuration. The special file name for IBM tape devices (not the IBM Tivoli Storage Manager devices) is `/dev/rmtX` for tape drives and `/dev/smcX` for medium changers.

Tip: Special file name: `/dev/rmt0`, `/dev/smc0`

For the Tivoli Storage Manager devices:

To update the special files, use **smitty > devices > Tivoli Storage Manager Devices > remove all defined devices**, then **discover devices supported by Tivoli Storage Manager**. The special file name is `/dev/mtX` for tape drives and `/dev/lbX` for medium changers.

Tip: Special file name: `/dev/mt0`, `/dev/lb0`

Alternatively, you can reinstall the IBM device driver. Tivoli Storage Manager device driver updates all the current special file name.

Update configuration on Windows:

With the plug and play, the Windows registry is updated and the device name might change without the need to restart the computer or have the device driver's involvement. The Tivoli Storage Manager server detects the change in a special file name and updates the new special file name when it accesses the tape devices (during server initialization or normal operation). The correct device name is updated in the Tivoli Storage Manager database. The special file name is `/dev/mtA.B.C.D` for both Tivoli Storage Manager devices and IBM devices, and `/dev/lbA.B.C.C` for both Tivoli Storage Manager devices and IBM medium changers. The special file name `TapeX` is only for IBM tape drives and `ChangerX` is only for IBM medium changes.

Tip: Special file name: `mt0.1.0.0`, `lb0.0.1.0`, `Tape0`, and `Changer0`.

Update configuration on Linux:

The host bus adapter (HBA) gets the most up-to-date configuration information as a result of the RSCN. Sometimes, the computer must be restarted to be able to pick up the configuration changes.

For IBM devices:

Issue the **lin_taped** command to reconfigure devices. The device information can be retrieved from the `/proc/scsi/IBMTape` file for tape devices and `/proc/scsi/IBMchanger` file for medium changers. The special file name is `/dev/IBMTapeX` for tape devices and `/dev/IBMChangerX` for medium changers.

Tip: Special file name: `/dev/IBMTape0`, `/dev/IBMChanger0`

For the Tivoli Storage Manager devices:

Users can issue **autoconf**, the Tivoli Storage Manager device driver auto configure script. This script resides in the `/opt/tivoli/tsm/devices/bin` directory (or in the same directory as the `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/tsm SCSI/mt0`, `/dev/tsm SCSI/lb0`

Alternatively, you can reinstall the IBM device driver. Tivoli Storage Manager device driver updates all the current special file names.

With the Linux pass-thru device driver for the Tivoli Storage Manager devices, the HBA driver and the generic driver must be reloaded to get all the current special file names. You have to run the `autoconf` script so that the Tivoli Storage Manager device driver can create configuration files (`/dev/tsm SCSI/lbinfo` and `/dev/tsm SCSI/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 HBA API 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 HBA API 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-HBA API wrapper library `libhbaapi64.so` is in the same directory as `dsmserv` or in the `/opt/tivoli/tsm/server/bin` directory. The wrapper library, `libhbaapi64.so`, is a link to the `/usr/lib64/libzfcphbaapi.so` file.

Update the configuration on Oracle Solaris:

The HBA gets the most up-to-date configuration information as result of the RSCN. Most of the time, the computer must be restarted to pick up the configuration changes. Reinstall the IBM device driver package or the Tivoli Storage Manager device driver package and run `autoconf`. Or issue the `rem_drv` and `add_drv` commands to reconfigure devices and update the special file name:

For IBM devices:

The IBM device driver has already completed configuring the devices after installing the device driver package. The `/opt/IBMtape/tapelist -l` command can show all IBM device information on the system. The special file name is `/dev/rmt/Xst` for tape devices and `/dev/rmt/Xsmc` for medium changers.

Tip: Special file name: `/dev/rmt/0st`, `/dev/rmt/0smc`

For Tivoli Storage Manager devices:

Modify `/usr/kernel/drv/mt.conf` and `/usr/kernel/drv/lb.conf` and ensure that `name="mt" parent="pseudo" instance=16383`; and `name="lb" parent="pseudo" instance=16383`; are not commented out in the `mt.conf` and `lb.conf` files. Ensure that each entry for a device in `/usr/kernel/drv/lb.conf` and `/usr/kernel/drv/mt.conf` is correct. After running `autoconf`, device information can be found in the `mtinfo` and `lbinfo` files in the `/opt/tivoli/tsm/devices/bin` directory or the same directory as the `autoconf` script. The special file name is `/dev/rmt/Xmt` for tape drives and `/dev/rmt/Xlb` for medium changers.

Tip: Special file name: /dev/rmt/0mt, /dev/rmt/01b

Important:

Ensure that the two pseudo devices, /devices/pseudo/mt@16383:tsmmtctl and /devices/pseudo/lb@16383:tsmlbctl, are in the /devices/pseudo directory.

Ensure that the pseudo device special files, /dev/tsmmtctl and /dev/tsmlbctl are linked to their corresponding pseudo devices in /devices/pseudo/mt@16383:tsmmtctl and /devices/pseudo/lb@16383:tsmlbctl.

Resolving configuration problems that cause SAN device absence:

The possible reasons for the **QUERY SAN** server command not displaying all the devices can be due to a configuration problem with the HBA hardware, HBA driver level, or operating system level.

About this task

Perform the following steps to resolve configuration issues:

Procedure

1. Go to the Support Portal. Verify the platform/HBA vendor/driver level support level to make sure that the HBA driver level and operating system level are compatible and supported by IBM Tivoli for SAN discovery.
2. Use the HBA vendor utility to check to see whether 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.
3. AIX Linux Solaris Use the dsmsanlist module to obtain information about devices on a storage area network (SAN). The dsmsanlist module is installed by default when the Tivoli Storage Manager server or the Tivoli Storage Manager storage agent is installed.

Verifying vendor support for any particular device in the SAN:

Many devices or combinations of devices might not be supported in a given storage area network (SAN). These limitations arise from the ability of a given vendor to certify their device using Fibre Channel Protocol.

For a given device, verify with the device vendor that it is supported in a SAN environment. Vendor support includes all hardware associated with the SAN, which means verifying that this device is supported with the vendors of the HBAs, hubs, gateways, and switches that make up the SAN environment.

NDMP filer-to-Tivoli Storage Manager server operation hints and tips

IBM Tivoli Storage Manager defaults to the standard network data management protocol (NDMP) control port of 10000. If this port is in use by another application (such as a second Tivoli Storage Manager server), all filer-to-server operations fail.

To avoid conflicts with other applications, use the `NDMPCONTROLPORT` server option to specify a different port for your server.

During filer-to-server operations, Tivoli Storage Manager uses the following items:

- Up to two extra TCP/IP ports.
- A control port that is used internally by Tivoli Storage Manager during both backup and restore operations.
- A data port during NDMP backup operations to a Tivoli Storage Manager native storage pool.

The data port is an ephemeral port that is acquired at the beginning of NDMP backup operations to a Tivoli Storage Manager native storage pool. If a port is not available, an error message is issued and backup of NAS devices to Tivoli Storage Manager native pools is not possible. To avoid conflicts with other applications, you can control which port is acquired for use as the data port during NDMP backup operations by setting the `NDMPPORTRANGELOW` and `NDMPPORTRANGEHIGH` server options. A data port is not needed by the Tivoli Storage Manager server because NAS restores from Tivoli Storage Manager native pools.

Firewall issues with NDMP filer-to-Tivoli Storage Manager server backup and restore

A firewall might prevent the network-attached storage (NAS) file server from contacting the Tivoli Storage Manager server on the acquired data port during NAS backup operations to a native storage pool. If you must modify the data port that is selected by the Tivoli Storage Manager server, use the `NDMPPORTRANGELOW` and `NDMPPORTRANGEHIGH` server options.

A firewall might prevent the Tivoli Storage Manager server from contacting the NAS file server on the configured data port during NAS restore operations from a native storage pool. If a firewall prevents Tivoli Storage Manager from accessing the NAS file server, the outbound connection from Tivoli Storage Manager fails.

Resolving SCSI device problems

Tape drives and libraries might report information back to IBM Tivoli Storage Manager about the error encountered. This information is reported in one or more of the messages.

If messages ANR8300, ANR8301, ANR8302, ANR8303, ANR8943, or ANR98944 are issued, the data that Tivoli Storage Manager reports from these devices might help to determine the steps that are needed to resolve the problem. Generally, when the server reports device 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 product information at Messages, return codes, and error codes. This appendix documents information about standard errors that might be reported by any SCSI device. You can also use this information with documentation provided by the vendor for the hardware to help determine the cause and resolution for the problem.

Resolving sequential media volume (tape) errors through messages ANR0542W or ANR8778W

Problems occurring with sequential media volumes can be revealed through error messages ANR0542W and ANR8778W.

ANR0542W Retrieve or restore failed for session *sessionNumber* for node *nodeName* - storage media inaccessible

Error message ANR0542W is often related to an issue with the drive or connection to the drive that was selected to read this tape volume.

Perform the following steps to verify that IBM Tivoli Storage Manager can access this volume:

- Issue the `QUERY LIBVOL libraryName volumeName` command.
- For a 349X library, issue the `mtlib -l /dev/lmcp0 -qV volumeName` command. The device is typically `/dev/lmcp0`, but if it is different, then substitute the correct library manager control point device.

The following steps might possibly resolve this problem:

1. If `mtlib` does not report this volume, then it appears that this volume is out of the library. In this case, put the volume back into the library.
2. If the volume is not reported by `QUERY LIBVOL`, then the server does not know about this volume in the library. Issue the **CHECKIN LIBVOL** command to synchronize the library inventory in the server with the volumes that are actually in the tape library.
3. If both commands successfully report this volume, then the cause is likely a permanent or intermittent hardware error. There might be an error with the drive itself or an error with the connection to the drive. In either case, review the system error logs and contact the vendor of the hardware to resolve the problem.

ANR8778W Scratch volume changed to private status to prevent re-access

Review the activity log messages to determine the cause of the problem involving this scratch volume. Also, review the system error logs and device error logs for an indication that there was a problem with the drive used to try to write to this scratch volume.

If this error was caused by a drive requiring cleaning or some other hardware-specific issue that was resolved, any volumes that were set to private status as a result of this might be reset to scratch by issuing the `AUDIT LIBRARY libraryName` command.

Appendix A. Getting call stack information from a core file

You can use the sample `gt` script shell provided here to get the call stack for each running thread from a core file.

The input parameters are the path/name of the executable file (default `./dsmserv`) and the path/name of the core file (default `./dsmcore`). The output file is `dsm_gdb.info`.

Restriction: Files named `dsm_gdb.cmd` and `dsm_gdb.info` are overwritten when you run this script.

```
#!/bin/ksh
#
# If you see the following error:
# ./dsm_gdb.cmd:9: Error in source command file:
# No symbol table is loaded. Use the "file" command.
# then comment out the line that prints buildStringP
#
# if you see other errors, you're on your own ...
exe=${1:-"./dsmserv"} # get parm 1 (executable file path/name), set default
core=${2:-"./dsmcore"} # get parm 2 (core file path/name), set default
echo " "
# look for the executable file ... quit if not found
if [[ -f $exe ]]; then
echo "using executable file:" $exe
else
echo "didn't find executable file ("$exe") ... exiting"
exit
fi
# look for the core file, if not found, look for ./core ... quit if not found
if [[ -f $core ]]; then
echo "using core file:" $core
else
if [[ -f ./core ]]; then
echo "didn't find core file ("$core") but found ./core ... renaming to" $core
mv ./core $core
echo "using core file:" $core
else
echo "didn't find core file ("$core") ... exiting"
exit
fi
fi
echo " "
# make gdb command file to get thread info
nl="\0134\0156" # octal codes for \n (so echo won't think it's \n)
echo "# dsm gdb command file" >|dsm_gdb.cmd
echo "define doit" >>dsm_gdb.cmd
echo "info registers" >>dsm_gdb.cmd # show register values
echo "echo" $nl >>dsm_gdb.cmd
echo "where" >>dsm_gdb.cmd # show function traceback
echo "echo" $nl"======"$nl >>dsm_gdb.cmd
echo "end" >>dsm_gdb.cmd
echo "echo" $nl"======"$nl$nl >>dsm_gdb.cmd
echo "x/s buildStringP" >>dsm_gdb.cmd
echo "echo" $nl"======"$nl$nl >>dsm_gdb.cmd
echo "info threads" >>dsm_gdb.cmd # show thread info
echo "echo" $nl"======"$nl >>dsm_gdb.cmd
echo "thread apply all doit" >>dsm_gdb.cmd
echo "quit" >>dsm_gdb.cmd
echo "invoking gdb to get thread info (watch for errors) ..."
echo "if you see:"
```

```
echo ". warning: The shared libraries were not privately mapped; setting a"
echo ". breakpoint in a shared library will not work until you rerun the program"
echo "that's ok."
echo "if you see:"
echo ". ./dsm_gdb.cmd:x: Error in source command file:"
echo "then type 'quit', edit this script, and read the comments at the top"
gdb -se $exe -c $core -x ./dsm_gdb.cmd >|dsm_gdb.info
rm dsm_gdb.cmd # done with this now
exit
```

Appendix B. Running the `tsmdiag` utility

You can diagnose problems with a Tivoli Storage Manager server by running the `tsmdiag` utility on the system where the Tivoli Storage Manager server is installed. After you collect the diagnostic data, you can send the information to IBM Software Support.

Procedure

To run the `tsmdiag` utility, complete the following steps:

- AIX** **HP-UX** **Linux** **Solaris** Change the `tsmdiag` directory permissions by issuing the following command:

```
chmod -R 757 /opt/tivoli/tsm/server/bin/tsmdiag
```
- Issue the `tsmdiag` command from the following directory:
 - AIX** **HP-UX** **Linux** **Solaris** By using a DB2 instance user ID, you must issue the `tsmdiag` command from the `/opt/tivoli/tsm/server/bin/tsmdiag` directory.
 - Windows** By using an administrator user ID, you must issue the `tsmdiag` command from the `\server\tsmdiag` directory.

For example, the following command collects a default set of diagnostic information files from a Tivoli Storage Manager server on a local host. This command is run by an administrator that is called `admin` with an administrator password `admin01` on a Tivoli Storage Manager server. This server is running on TCP/IP port 1501 of a local host.

```
tsmdiag -id admin -pa admin01 -tcpport 1501
```

- Obtain the results file from the following directory:
 - AIX** **HP-UX** **Linux** **Solaris** `/opt/tivoli/tsm/server/bin/tsmdiag/results/tsmdiag_results<year>-<month>-<day>-<hour>-<minute>-<second>.tar`
 - Windows** `C:\Program Files\tivoli\tsm\server\tsmdiag\results\tsmdiag_results<year>-<month>-<day>-<hour>-<minute>-<second>.zip`
- Submit the results file with the problem report to IBM Software Support.

Example `tsmdiag` commands

The following command connects to a Tivoli Storage Manager server called `MYSERVER` on TCP/IP port 1501. When a DB2 administrator called `admin` runs the following command, a default set of diagnostic information files is collected. Diagnostic information about the performance of the server `MYSERVER` is also collected.

```
tsmdiag -id admin -pa admin01 -tcpport 1501 -servername MYSERVER -performance
```

The following command connects to a Tivoli Storage Manager server on the default TCP/IP port 1500. When an administrator called `admin` runs the following command, the command collects a default set of diagnostic information files. This command also provides results from the Tivoli Storage Manager server **SHOW** commands and some diagnostic information about the status of the Tivoli Storage Manager server.

```
tsmdiag -id admin -pa admin01 -hang
```

tsmdiag utility options

You use the `tsmdiag` utility to help diagnose problems with a Tivoli Storage Manager server component. When you run the utility, you can specify options that determine the type of diagnostic information that is provided.

You can specify the following options when you issue the `tsmdiag` command:

id *adminName*

The administrator or root user ID of the server on which to run the `tsmdiag` command. This option is mandatory.

-pa *adminPwd*

The password for the administrator or root user ID. This option is mandatory.

-tcpserveraddress *ipAddress*

Specifies the TCP/IP name or address of the server on which to run the `tsmdiag` command. The option is optional. The default value is `localhost`.

-tcpport *portNumber*

Specifies the TCP/IP port of the server on which to run the `tsmdiag` command. This option is optional. The default value is `1500`.

AIX

HP-UX

Linux

Solaris

-servername

The name of the server on which to run the `tsmdiag` command. This option is optional. The default value is `SERVER1`.

-crash Specifies whether to report if the server has failed. This option is optional. The default value is `off`.

-dbc corrupt

Specifies whether to report if database corruption exists. This option is optional. The default value is `off`.

-db growth

Specifies whether to report if there is excessive large database growth on the server. This option runs the `serverReorgInfo.pl` script and `tsmdiag_dedup_stats.pl` script, which produce additional diagnostic information. The `serverReorgInfo.pl` script takes over 1 hour to run. This option is optional. The default value is `off`.

-hang Specifies whether to report if the server is hanging. This option is optional. The default value is `off`.

-performance

Specifies whether to report if there are performance problems with the server. This option runs the `tsmdiag_sysmonv6.pl` script, which produces extra diagnostic information. The `tsmdiag_sysmonv6.pl` script can take 1.5 hours to run. This option is optional. The default value is `off`.

-v Specifies that the report output is produced in verbose format. This option is optional. The default value is `off`.

-? Specifies the usage information for the `tsmdiag` utility. If you issue the `tsmdiag ?` command, a list of the preceding options is displayed.

Appendix C. IBM Global Security Kit return codes

The server and client use the IBM Global Security Kit (GSKit) for SSL (Secure Sockets Layer) processing between the server and the backup-archive client. Some messages that are 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 that is installed on your system, or you can use one of the following methods:

- For Windows, issue the following commands:

```
regedit /e gskitinfo.txt "HKEY_LOCAL_MACHINE\software\ibm\gsk8\"
notepad gskitinfo.txt
```

CAUTION:

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

- For the 64-bit AIX server, issue the following command from the command line:
`gsk8ver_64`

See Table 17 for the GSKit SSL return codes.

The server uses the GSKit Key Management API to automatically create the key management database and server private and public keys. Some messages that are issued for this processing might include GSKit Key Management return codes. See Table 18 on page 204 for the key management return codes.

Table 17. IBM Global Security Kit SSL general return codes

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000000	0	GSK_OK	The task completes successfully. Issued by every function call that completes successfully.
0x00000001	1	GSK_INVALID_HANDLE	The environment or SSL handle is not valid. The specified handle was not the result of a successful <code>open()</code> function call.
0x00000002	2	GSK_API_NOT_AVAILABLE	The dynamic link library (DLL) was unloaded and is not available (occurs on Microsoft Windows systems only).
0x00000003	3	GSK_INTERNAL_ERROR	Internal error. Report this error to IBM Software Support.
0x00000004	4	GSK_INSUFFICIENT_STORAGE	Insufficient memory is available to complete the operation.
0x00000005	5	GSK_INVALID_STATE	The handle is not in a valid state for operation, such as completing an <code>init()</code> operation on a handle twice.

Table 17. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000006	6	GSK_KEY_LABEL_NOT_FOUND	Specified key label is not found in key file.
0x00000007	7	GSK_CERTIFICATE_NOT_AVAILABLE	Certificate is not received from the 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.
0x0000000d	13	GSK_INVALID_PARAMETER	Invalid parameter.
0x0000000e	14	GSK_ERROR_UNEXPECTED_INT_EXCEPTION	Invalid parameter. 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 the 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 is 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. Check that 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	Both the password and the stash-file name were not specified. The key file is not 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.

Table 17. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
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 that is used for an LDAP (lightweight directory access protocol) 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 that is 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	The SSLv2 and the SSLv3 are not enabled.
0x00000193	403	GSK_ERROR_NO_CERTIFICATE	The required certificate was not received from the 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 is not found.
0x00000198	408	GSK_ERROR_BAD_KEYFILE_PASSWORD	The specified key file password is incorrect. The key file cannot 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 is received from the partner.
0x0000019f	415	GSK_ERROR_BAD_PEER	Did not receive a valid SSL protocol from the partner.
0x000001a0	416	GSK_ERROR_PERMISSION_DENIED	Report this error to IBM Software Support.

Table 17. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
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 cannot be created. Report this error to IBM Software Support.
0x000001aa	426	GSK_ERROR_INITIALIZATION_FAILED	Initialization failed. Report this internal error to service.
0x000001ab	427	GSK_ERROR_LDAP_NOT_AVAILABLE	Not able to access the specified user registry when a certificate is being validated.
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 that is specified by the caller.
0x000001af	431	GSK_ERROR_PKCS11_TOKEN_NOTPRESENT	A PKCS #11 token is not present in the slot.
0x000001b0	432	GSK_ERROR_PKCS11_TOKEN_BADPASSWORD	The password/pin to access the PKCS #11 token is not valid.
0x000001b1	433	GSK_ERROR_INVALID_V2_HEADER	The SSL header received was not a properly formatted SSLv2 header.
0x000001b2	434	GSK_CSP_OPEN_ERROR	Cannot open the hardware-based cryptographic service provider. Either the CSP name is not specified correctly or a failed attempt was made to access the specified CSP certificate store.
0x000001b3	435	GSK_CONFLICTING_ATTRIBUTE_SETTING	Attribute setting conflict between PKCS11, CMS key database, and Microsoft Crypto API.
0x000001b4	436	GSK_UNSUPPORTED_PLATFORM	The requested function is not supported on the platform that the application is running. For example, the Microsoft Crypto API is not supported on platforms other than Windows 2000.

Table 17. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x000001b6	438	GSK_ERROR_INCORRECT_SESSION_TYPE	Incorrect value is returned from the reset session type callback function. Only GSKit <code>gsk_sever_session</code> , <code>gsk_sever_session_with_cl_auth</code> , or <code>gsk_sever_session_with_cl_auth_crit</code> is allowed.
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 <code>reset_cipher()</code> , and the connection uses SSLv2.
0x0000025a	602	GSK_MISC_INVALID_ID	A valid ID was not specified for the <code>gsk_secure_soc_misc()</code> function call.
0x000002bd	701	GSK_ATTRIBUTE_INVALID_ID	The function call does not have a valid ID. This issue might also be caused by specifying an environment handle when a handle for an 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 a numeric attribute is set, the specified value is not valid for the specific attribute that is 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.
0x000002c5	709	GSK_CIPHER_INVALID_WHEN_FIPS_MODE_OFF	The particular cipher is not allowed when FIPS mode of operation is off.
0x000002c6	710	GSK_CIPHER_INVALID_WHEN_FIPS_MODE_ON	No approved FIPS ciphers are selected in FIPS mode of operation.
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 is started, so it cannot be restarted.
0x00000645	1605	GSK_TRACE_OPEN_FAILED	Trace file cannot be opened. The first parameter of <code>gsk_start_trace()</code> must be a valid full path file name.

Table 18. IBM Global Security Kit key management return codes

Return code (hex)	Return code (decimal)	Constant	Explanation
0x00000000	0	GSK_OK	The task completes successfully. This message is issued by every function call that completes successfully.
0x00000001	1	GSK_INVALID_HANDLE	The environment or SSL handle is not valid. The specified handle was not the result of a successful open() function call.
0x00000002	2	GSK_API_NOT_AVAILABLE	The DLL (dynamic link library) was unloaded and is not available (occurs on Microsoft Windows systems only).
0x00000003	3	GSK_INTERNAL_ERROR	Internal error. Report this error to IBM Software Support.
0x00000004	4	GSK_INSUFFICIENT_STORAGE	Insufficient memory is available to complete the operation.
0x00000005	5	GSK_INVALID_STATE	The handle is in an incorrect state for operation, such as completing an init() operation on a handle twice.
0x00000006	6	GSK_KEY_LABEL_NOT_FOUND	Specified key label is not found in key file.
0x00000007	7	GSK_CERTIFICATE_NOT_AVAILABLE	Certificate is not received from the 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 is corrupted or the password ID is incorrect.
0x0000006b	107	GSK_KEYFILE_CERT_EXPIRED	The default key in the key file has an expired certificate.

Table 18. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x0000006c	108	GSK_ERROR_LOAD_GSKLIB	An error occurred while one of the GSK dynamic link libraries is loaded. Check GSK was installed correctly.
0x0000006d	109	GSK_PENDING_CLOSE_ERROR	This message indicates that a connection is trying to be made in a GSK environment after the GSK_ENVIRONMENT_CLOSE_OPTIONS was set to GSK_DELAYED_ENVIRONMENT_CLOSE and gsk_environment_close() function was called.
0x000000c9	201	GSK_NO_KEYFILE_PASSWORD	Both the password and the stash-file name were not specified, so the key file is not 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.
0x000000cd	205	GSK_ERROR_LDAP_INVALID_CREDENTIALS	A Password that is used for an LDAP query is not correct.
0x000000ce	206	GSK_ERROR_BAD_INDEX	An index into the Fail Over list of LDAP servers was not correct.
0x000000cf	207	GSK_ERROR_FIPS_NOT_SUPPORTED	This installation of GSKit does not support FIPS mode of operation.
0x0000012d	301	GSK_CLOSE_FAILED	Indicates that the GSK environment close request was not properly managed. Cause is most likely due to attempting a gsk_secure_socket*() command after a gsk_close_environment() 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	SSLv2 and SSLv3 are not enabled.
0x00000193	403	GSK_ERROR_NO_CERTIFICATE	The required certificate was not received from the partner.
0x00000194	404	GSK_ERROR_BAD_CERTIFICATE	The received certificate was formatted incorrectly.

Table 18. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
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 is not found.
0x00000198	408	GSK_ERROR_BAD_KEYFILE_PASSWORD	The specified key file password is incorrect. The key file cannot be used. The key file might also be corrupt.
0x00000199	409	GSK_ERROR_BAD_KEY_LEN_FOR_EXPORT	In a restricted cryptography environment, the key size is too long to be supported.
0x0000019a	410	GSK_ERROR_BAD_MESSAGE	An incorrectly formatted SSL message was received from the partner.
0x0000019b	411	GSK_ERROR_BAD_MAC	The MAC was not successfully verified.
0x0000019c	412	GSK_ERROR_UNSUPPORTED	Unsupported SSL protocol or unsupported certificate type.
0x0000019d	413	GSK_ERROR_BAD_CERT_SIG	The received certificate contained an incorrect signature.
0x0000019e	414	GSK_ERROR_BAD_CERT	Incorrectly formatted certificate is received from the partner.
0x0000019f	415	GSK_ERROR_BAD_PEER	An SSL protocol that is not valid is received from the partner.
0x000001a0	416	GSK_ERROR_PERMISSION_DENIED	Report this error to IBM Software Support.
0x000001a1	417	GSK_ERROR_SELF_SIGNED	The self-signed certificate is not valid.
0x000001a2	418	GSK_ERROR_NO_READ_FUNCTION	The read() failed. Report this error to IBM Software Support.
0x000001a3	419	GSK_ERROR_NO_WRITE_FUNCTION	The write() failed. Report this error to IBM Software Support.
0x000001a4	420	GSK_ERROR_SOCKET_CLOSED	The partner closed the socket before the protocol completed.
0x000001a5	421	GSK_ERROR_BAD_V2_CIPHER	The specified V2 cipher is not valid.
0x000001a6	422	GSK_ERROR_BAD_V3_CIPHER	The specified V3 cipher is not valid.
0x000001a7	423	GSK_ERROR_BAD_SEC_TYPE	Report this error to IBM Software Support.
0x000001a8	424	GSK_ERROR_BAD_SEC_TYPE_COMBINATION	Report this error to IBM Software Support.

Table 18. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x000001a9	425	GSK_ERROR_HANDLE_CREATION_FAILED	The handle is not created. Report this error to IBM Software Support.
0x000001aa	426	GSK_ERROR_INITIALIZATION_FAILED	Initialization failed. Report this internal error to service.
0x000001ab	427	GSK_ERROR_LDAP_NOT_AVAILABLE	Unable to access the specified user registry when a certificate is being validated
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 that is specified by the caller.
0x000001af	431	GSK_ERROR_PKCS11_TOKEN_NOTPRESENT	A PKCS #11 token is not present in the slot.
0x000001b0	432	GSK_ERROR_PKCS11_TOKEN_BADPASSWORD	The password/pin to access the PKCS #11 token is incorrect.
0x000001b1	433	GSK_ERROR_INVALID_V2_HEADER	The SSL header received was not a properly formatted SSLv2 header.
0x000001b2	434	GSK_CSP_OPEN_ERROR	Could not open the hardware-based cryptographic service provider (CSP). Either the CSP name is not specified correctly or a failed attempt was made to access the specified CSP certificate store.
0x000001b3	435	GSK_CSP_OPEN_ERROR	Some conflicting attributes for SSL operation were defined.
0x000001b4	436	GSK_CSP_OPEN_ERROR	The Microsoft Crypto API is only supported on Microsoft Windows 2000 with Service Pack 2 applied.
0x000001b5	437	GSK_CSP_OPEN_ERROR	System is running in IPv6 mode without setting a PEERID.
0x000001f5	501	GSK_INVALID_BUFFER_SIZE	The buffer size is negative or zero.
0x000001f6	502	GSK_WOULD_BLOCK	Used with nonblocking I/O. Refer to the nonblocking section for usage.
0x00000259	601	GSK_ERROR_NOT_SSLV3	SSLv3 is required for reset_cipher(), and the connection uses SSLv2.

Table 18. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x0000025a	602	GSK_MISC_INVALID_ID	An ID that is not valid was specified for the <code>gsk_secure_soc_misc()</code> function call.
0x000002bd	701	GSK_ATTRIBUTE_INVALID_ID	The function call has an ID that is not valid. This issue might also be caused by specifying an environment handle when a handle for an 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 a numeric attribute is set, the specified value is not valid for the specific attribute that is 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.
0x000002c5	709	GSK_CIPHER_INVALID_WHEN_FIPS_MODE_OFF	The particular cipher is not allowed when FIPS mode of operation is off.
0x000002c6	710	GSK_CIPHER_INVALID_WHEN_FIPS_MODE_ON	No approved FIPS ciphers are selected in FIPS mode of operation.
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 is started so it cannot be started again.
0x00000645	1605	GSK_TRACE_OPEN_FAILED	Trace file cannot be opened. The first parameter of <code>gsk_start_trace()</code> must be a valid, full-path file name.

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

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

Accessibility features

The IBM Tivoli Storage Manager family of products includes the following accessibility features:

- Keyboard-only operation using standard operating-system conventions
- Interfaces that support assistive technology such as screen readers

The command-line interfaces of all products in the product family are accessible.

Tivoli Storage Manager Operations Center provides the following additional accessibility features when you use it with a Mozilla Firefox browser on a Microsoft Windows system:

- Screen magnifiers and content zooming
- High contrast mode

The Operations Center and the Tivoli Storage Manager server can be installed in console mode, which is accessible.

The Operations Center help system is enabled for accessibility. For more information, click the question mark icon on the help system menu bar.

Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

IBM and accessibility

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

Notices

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

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

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

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
United States of America*

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.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

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

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

*IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
US*

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

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

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

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

COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

Portions of this code are derived from IBM® Corp. Sample Programs.

© Copyright IBM® Corp. _enter the year or years_. All rights reserved.

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 www.ibm.com/legal/copytrade.shtml.

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

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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

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

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

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

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

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights Except as expressly granted in this permission, no other permissions,

licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

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

See Tivoli Storage Manager glossary.

To view glossaries for other IBM products, see IBM Terminology.

Index

Special characters

\$\$_TSMDBMGR_\$\$ hidden user ID 81

A

- accessibility features 209
- active tasks
 - delay in canceling 101
- administrative commands
 - DELETE KEYRING 107
- administrators
 - locked 13
- AIX JFS2
 - image backup 31
 - snapshot-based backup-archive 31
- alerts
 - delay in closing or assigning 100
- allocating additional memory 55
- ANR1221E
 - error message 88
- ANR2317W
 - error message 89
- API
 - option file 35
- application programming interface (API)
 - instrumentation 33
 - tracing 159
- assign multiple alerts
 - delay 100
- automatic deployment
 - troubleshooting 62

B

- backup application
 - files automatically excluded 21
 - files excluded by EXCLUDE DIR 23
 - files excluded by include/exclude statements 21
 - files excluded due to incremental copy frequency 56
 - include/exclude due to compression, encryption, and subfile backup statements 24
 - include/exclude statements coded wrong 25
 - platform-specific include/exclude statements 25
- BACKUP DB
 - ANR2971E with SQL code 79
 - common errors 80
 - incorrect environment variables 77
- backup-archive client
 - help 1
 - SHOW commands 45

C

- cache
 - bypass during write operations 174
- cancel multiple tasks
 - delay 101
- certificate authority 104

- client
 - authentication failure 10
 - can problem be reproduced 6
 - error messages
 - examining 5
 - generating errors
 - connected to the server 103
 - identifying when and where problems occur 5
 - image backup 29
 - resolving problems 5
 - scheduler 17
 - server activity log
 - examining 5
 - trace classes 144, 149
- client deployment
 - troubleshooting 62
- client option sets
 - resolving problems 9
 - using 9
- client schedule log 18
- close multiple alerts
 - delay 100
- communication errors
 - resolving 103
- complex password
 - LDAP directory server 12
- complex passwords
 - audit the LDAP directory server 14
- compressed data during backup-archive 158
- copy frequency 56

D

- Daemon traceflags
 - client and journal 143
- data
 - sent to the IBM Tivoli Storage Manager storage agent or server 36
 - unreadable 165
- data storage hints and tips
 - backup or copy problem with specific node 167
 - change the server policies 166
 - change the storage hierarchy 166
 - HELP 165
 - reading or writing to a device 166
 - recreate the problem 165
 - server activity log 165
 - specific volume 167
- database error messages 77
- database ID file missing or incorrect 76
- database manager
 - start problems 71
- database page verification failure 70
- database reorganization 82
- database restore errors 76
- DB2 log files 74
- DB2 memory 73
- DB2 password
 - expired 62
- DB2 version 74

- db2dump directory
 - shutdown resolution 69
- DELETE KEYRING command 107
- device driver
 - 32-bit Linux kernel modules 169
 - 64-bit Linux kernel modules 169
 - Adaptec SCSI requirements 172
 - error messages in the system error log 169
 - HBA changes 168
 - HBA drivers on the Linux 2.6.x kernels 170
 - Linux server running on x86_64 architecture 169
 - loose cable connection 169
 - multiple LUN support on Linux kernels 170
 - operating system changes 168
 - performing ddtrace from version 5.3.2 on Linux 171
 - QLogic fibre-channel HBA BIOS requirements 172
 - SCSI adapter changes 168
 - updating device information 172
- device driver trace
 - from a command shell - AIX, Oracle Solaris, Windows 142
 - from the server console/admin client 141
- diagnostic tips
 - client 5
 - storage agent 109
- disability 209
- documentation
 - to resolve client problems 6
- dsmsanlist 179, 182, 187, 191

E

- encrypted data during backup-archive 158
- encrypted file system 28
- error messages
 - ANR1330E 90
 - ANR1331E 90
 - ANR2968E 78
 - LDAP authenticated passwords 14
- extended tracing 113, 114
- external user repository server
 - stoppage 65

F

- FILE directory mapping 174

G

- GSKit
 - installation problems 60
 - return codes 199
- gt script 195

H

- help
 - server or storage agent 2
- help facilities 1
- help system
 - CLI for server or storage agent 3
 - dsmcutil 2
 - GUI and Web GUI clients 3
 - reporting a problem 3
 - server or storage agent
 - commands 2

- help system (*continued*)
 - server or storage agent (*continued*)
 - messages 3
 - Windows 2
- hints and tips
 - device driver 167
 - disk subsystems 173
 - hard disk drives 173
 - NDMP filer-to-Tivoli Storage Manager server
 - operations 192
 - SAN 177
 - SAN configuration 177
 - SAN device mapping 183
 - tape drives and libraries
 - adapter firmware changes 175
 - cabling between the computer and device changes 175
 - device driver changes 176
 - device firmware changes 175
 - error messages in system error log 177
 - loose cable connections 176
 - operating system changes 176
 - other hardware changed or fixed 176
 - replaced adapter 176

I

- IBM Global Security Kit
 - key management return codes 199
 - return codes 199
- IBM Knowledge Center vii
- image backup
 - client 29
 - error 29, 30
- IMPORT command 57
- INCLEXCL option 21
- Installation Manager
 - logs directory 60
- installation problems 60

J

- journal
 - restarting 40
- journal-based backup (JBB)
 - database viewing utility 40
 - determining 39
 - running in foreground 40

K

- key database file
 - out-of-synch 106
 - password recovery 106
- keyboard 209
- Knowledge Center vii
- known issues
 - with the Operations Center 102

L

- LABEL LIBVOLUME 57
- LAN-free setup
 - storage agent 110
- LDAP directory server
 - password 12

- LDAP-authenticated password
 - problem resolution 10
- limit memory 73
- limitations
 - of the Operations Center 102
- Linux image backup error 29
- Linux Snapshot image backup error
 - error message ANS1258E 30
- locked nodes and administrators 13
- log files 99, 100, 113, 114
 - DB2 upgrade 75
 - installation 60
- logging configuration file 114
- logging groups 113, 114

M

- Microsoft diagnostic information
 - VSS 43
- Microsoft tuning
 - VSS 43
- monitoring agents
 - activating tracing 160, 162
- moving data to other volumes 174

N

- nodes
 - locked 13
- non-root user ID
 - running applications using the API 37
- ntbackup.exe 45

O

- Operations Center
 - known issues 102
 - troubleshooting 99, 100, 113, 114
- options 198

P

- password authentication
 - client setup 11
- performing ddtrace on HP-UX 171
- process ended 86
- process started 85
- process symptoms
 - files not expired 93
 - migration does not run 93
 - migration only uses one process 93
- processes
 - delay in canceling 101
- programs
 - dsm 7
 - dsmadm 7
 - dsmc 7
 - dsmj 7
- publications vii

R

- recovery of individual SQL databases from a VM backup
 - DBCS SQL database names 49
 - displaying active SQL databases 49

- recovery of individual SQL databases from a VM backup
 - (continued)
 - messages 49
 - resolving problems 47
 - saving VSS XML manifest files 50
 - troubleshoot database access 47
- recovery of individual SQL databases from a VM backup
 - determining the status of VSS writers 51
- RELABEL 57
- reorganization
 - database 82
- RESTORE DB
 - ANR2971E with SQL code 79
 - common errors 80
 - incorrect environment variables 77
- restricting DB2 memory 73

S

- SAN
 - configuration 188
 - configuration between devices 180
 - configuration problems 191
 - fibre channel switch configuration 179
 - fibre-channel link error report 180
 - gateway port settings 179
 - host bus adapter configuration 178
 - host bus adapters 178
 - vendor support 191
 - SAN device mapping
 - disabling 184
 - errors 184
 - missing from the display of QUERY SAN 188
 - SAN devices
 - storage agent 181
 - scheduled event
 - status 18
 - scheduler
 - client service restart 19
 - SCSI devices 192
 - Secure Sockets Layer (SSL)
 - determining errors 104
 - general return codes 199
 - sequential media volume
 - tape 193
 - server
 - database 71
 - diagnostic tips
 - change server options or the settings create errors 54
 - checking the server activity log 53
 - code page conversion failure 142
 - failing a scheduled client operation 54
 - recreating the problem 53
 - resolving errors from reading or writing to a device 54
 - resolving failed connections by client or administrators 103
 - resolving server space issues 55
 - process 82
 - process messages 82
 - stoppage or loop errors 64
 - storage pool
 - ANR0522W error message 94
 - collocation 95
 - COPY ACTIVATEDATA command 96
 - high volume usage 95
 - resolving problems 94, 96
 - simultaneous write 96

- server (*continued*)
 - storage pool (*continued*)
 - unable to store data 96
- server activity log
 - checking for errors 18
- server instance
 - configuring 56
- server or storage agent
 - trace classes 117
- server stoppage
 - activity log 68
 - library files 67
 - resolving general problems 63
 - server error file (dsmserv.err) 66
 - system image 66
 - system logs 68
- sessions
 - delay in canceling 101
- SET LDAPPASSWORD command
 - problems related to 12
- shared memory 56
- SHOW commands
 - server or storage agent 131
- Snapshot Difference
 - resolving problems 26
- snapshot directory 28
- SSL (Secure Sockets Layer)
 - determining errors 104
 - general return codes 199
- stack trace
 - server or storage agent 116
- startup problems
 - dsm 7
 - dsmadmc 7
 - dsmc 7
 - dsmj 7
- status
 - scheduled event 18
- storage agent
 - diagnostic tips
 - check the server activity log 109
 - error caused by reading or writing to a device 109
 - problems caused by changing server options 110
 - problems from changing storage agent options 110
 - LAN-free setup
 - data sent directly to server 110
 - storage pool configured for simultaneous write 111
 - testing LAN-free configuration 112
 - SAN devices 181
- summary records 59
- support for API
 - before calling IBM
 - files to gather 34
 - information to gather 33

T

- table reorganization 82
- test flags
 - VSS 42
- trace
 - client
 - backup-archive client 149
 - device driver 140
 - enable client trace on command line 149
 - enable client trace while client is running 151
 - known problems and limitations 155

- trace (*continued*)
 - options 156
 - server or storage agent 115
- trace classes
 - client 144
 - server or storage agent 117
- trace data
 - is it compressed during backup-archive 158
 - is it encrypted during backup-archive 158
- tracing 100, 113, 114
 - agents 160, 162
 - application programming interface (API) 159
 - client 142
 - User ID/Password plug-in 72
- transient errors
 - VSS 42
- troubleshooting
 - Operations Center 99, 100, 113, 114
- tsmdiag 197, 198
- tsmdiag utility 197

U

- uninstall stoppage 62
- upgrade
 - server manually 61
- upgrade problems 60
- utilities
 - tsmdiag 197

V

- Volume Shadow Copy Services
 - Windows 42
- vsreq.exe sample program 44
- VSS
 - Microsoft diagnostic information 43
 - Microsoft tuning 43
 - ntbackup.exe 45
 - test flags 42
 - trace 44
 - transient errors 42
 - vsreq.exe sample program 44
 - Windows 42

W

- Windows
 - VSS 42
- Windows Services
 - server service start/stop 69



Product Number: 5608-E01
5608-E02
5608-E03

Printed in USA