IBM Storage Protect
8.1.21

Problem Determination Guide

Note:

Before you use this information and the product it supports, read the information in “Notices” on page 205.

This edition applies to version 8, release 1, modification 21 of IBM® Storage Protect (product number 5725-W98, 5725-W99, 5725-X15), and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

About this publication

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

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

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

Any new and changed information in this product documentation is indicated by a vertical bar (|) to the left of the change.

Who should read this guide

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

You should be familiar with IBM Storage Protect and the operating systems used for the IBM Storage Protect environment.

Publications

The IBM Storage Protect product family includes IBM Storage Protect Plus, IBM Storage Protect for Virtual Environments, IBM Storage Protect for Databases, and several other storage management products from IBM.

To view IBM product documentation, see IBM Documentation.

Chapter 1. Help facilities

IBM Storage Protect 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

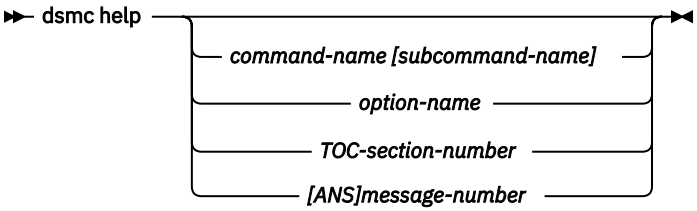

Entering the **HELP** command with no arguments causes help to display the complete table of contents. Either with the initial command or when HELP displays a prompt, you can enter the following parameters.

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.

Windows Access help for the Client Service Configuration Utility (dsmcutil)

To obtain help information for the IBM Storage Protect 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 4296, it assumes the message prefix ANR and reports the help information for ANR4296W. If the message number is specified with the prefix, for example HELP ANR4296, it reports the help information for that message. Issue HELP ANR4296 to view the following example output for that message:

```
ANR4296W A reclamation operation was started on offsite copy storage pool volumes. This
operation mightresult in data movement charges from your cloud storage provider.
Explanation: This operation reclaims data by using a cloud-container storage pool that is based
in the public cloud.When data is moved from a public cloud to an on-premises location, your
organization might incur datamovement charges from the cloud storage provider during reclamation
processing.
System Action: Server operation continues.
User Response: To help minimize data movement charges, you can disable reclamation for offsite
volumes on the copystorage pool. Issue the UPDATE STGPOOL command and specify the
OFFSITERECLAIMLIMIT=0
parameter setting.
```

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 IBM Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect.

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

- Use a minimal options file consisting of only `TCPSERVERADDRESS`, `TCPPORT`, and `NODENAME`.

- If the problem occurs for a file during incremental backup, try to reproduce the problem with a selective backup of just that file.
- If the problem occurs during a scheduled event, try to reproduce the problem by manually running the command.

Resolving issues with the web user interface for IBM Storage Protect backup-archive client

With the web graphical user interface (GUI) for IBM Storage Protect backup-archive client, you can restore, backup, archive, and retrieve data that is saved to the IBM Storage Protect server.

For more information, see the *Using the IBM Storage Protect web user interface for remote client operations* in IBM Documentation.

The following sections cover some common problems that you might encounter when you try to access the backup-archive client web GUI.

The web GUI is not accessible with the URL `https://hostname:9081/bagui`

If you cannot access the web GUI, follow these steps:

1. Ensure that the GUI is fully installed.
2. If the GUI is fully installed, ensure that you have installed the GUI by completing one of the following method depending on your operating system:
   - Windows During installation, you must select **Customized Installation**, where you can select the **Web GUI** option to install.
   - Linux AIX You must have to choose the Web package to install.
3. Complete one of the following action depending on your operating system:
   - Windows Verify that the IBM Storage Protect for BAClient Web Server service is running.
   - Linux Verify the output of the command **`/etc/init.d/webserver status`**.
4. Restart the system.

The web GUI is accessible but you cannot log in

If you cannot log in to the web GUI, ensure that you have configured the remote client agents **dsmcad** and **dsmagent** appropriately.

For more information, see the *Starting the client acceptor service and registering an administrator* in IBM Documentation.

You can log in to the web GUI but operations such as backup or restore fail

When an operation fails, you can download all the relevant logs from the web GUI to analyze and determine the problem.

To determine the reason and resolve a failed operation, use the following steps:

1. In the web GUI, click the information icon by the failed operation to see the corresponding ANS message for the error.
2. If you need additional information, download and analyze the `dsmerror.log` file:
   a. Download the system logs. See *Downloading system logs* in IBM Documentation for instructions.
   b. Open the `ba/dsminfo.txt` file. The `ba/dsminfo.txt` file includes information from the following files: `dsmerror.log`, `dsmwebcl.log`, and `dsmsched.log`.

c. Use the error information from the `dsmerror.log` file to analyze and resolve the problem. See the ***ANS messages* list** for additional help to resolve the relevant messages.

3. If the previous steps do not resolve the problem, start the backup-archive client using one of the following options:
   - Issue **DSMC** commands to start the backup-archive client. For help using the different **DSMC** command options, see *Use options on the DSMC command* in IBM Documentation.
   - Use the Java™ GUI session to start the backup-archive client. For help using Java GUI options, see *Starting a Java GUI session* in IBM Documentation.

Note: If you still cannot start the backup-archive client using these options, resolve this issue before continuing with the web GUI operations. For help, see *Determining why the dsmc, dsmadmc, dsm, or dsmj programs do not start* in IBM Documentation.

Collecting documentation to resolve problems with the client application

IBM Software Support 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: IBM Storage Protect 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 `dsmsched.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®, Linux®, and Mac OS X clients.
  - The include-exclude file. This file contains the objects to include or exclude from client operations. Its location is set by the client `inclexcl` 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 IBM Storage Protect 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**

Linux AIX **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
IBM
Storage Protect 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, IBM Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect 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.
```

Mac OS X Linux AIX Extra setup steps are required for non-root users in order for them to be able to run IBM Storage Protect applications or IBM Storage Protect for Data Protection applications. You receive the ANS1398E error if you try to run IBM Storage Protect 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 an IBM Storage Protect API error. Here is one method for setting up `dsmerror.log` for use by non-root users:

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

Resolving problems with client option sets

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

An administrator for IBM Storage Protect can create a set of client options to be used by a client node on IBM Storage Protect. The client options are defined on the IBM Storage Protect 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    Override    Option Value
                   number
------------      --------    --------    ---------------------------------
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 Storage Protect 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 IBM Storage Protect 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:\lnotes\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, IBM Storage Protect uses the node name in the client options file to contact the server. However, when the scheduler runs as a Windows service, IBM Storage Protect 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 IBM Storage Protect scheduler service on Windows

- With the *validate:yes* parameter, contact the IBM Storage Protect server to authenticate (and store the updated password)

Resolving LDAP-authenticated password problems

Most problems that arise from password authentication might be attributed to the connection between the IBM Storage Protect 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 IBM Storage Protect server. For more information, see *Authenticating users by using an Active Directory database*.

For the LDAP authentication method that is used with the IBM Storage Protect server 7.1.7 and later, a user (LDAPUSER) is no longer required full directory authority for authentication. A user can authenticate with the default user authority in Active Directory. Full directory authority is still required for IBM Security Directory Server LDAPUSER accounts.

The IBM Storage Protect 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, IBM Storage Protect 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 user principal name (upn)
- 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=abbynormal,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.

For more information about configuration, see *Authenticating users by using an Active Directory database* in IBM Documentation.

After you install the Tivoli® Storage Manager V6.3.3 or later server, or the IBM Storage Protect 7.1.3 or later server, 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 following command with **StartTLS** option:

```
openssl s_client -connect ldaphostname:389 -showcerts -CAfile rootca.crt -starttls ldap
```

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 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=tsmserver,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 caused by incorrect LDAP server certificate

If you are experiencing errors such as invalid certificate or other problems that are related to incorrect certificate, you must validate the LDAP server certificate.

Procedure

To validate the LDAP server certificate, complete the following steps:

1. Obtain the correct LDAP certificate from the LDAP server that is stored under `ldap.crt` file by issuing the following command:

   ```
   openssl s_client -starttls ldap -connect ldap.example.com:389 -showcerts </dev/null 2>/dev/
   null | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' | tee ldap.crt | openssl x509
   -text
   ```

2. Add the LDAP certificate that you obtained in the step 1 to the `cert.kdb` file by issuing the following command:

   ```
   gsk8capicmd_64 -cert -add -db cert.kdb -stashed -label ldap_server -file ldap.crt
   ```

3. Validate the LDAP certificate that is added in `cert.kdb` file by issuing the following command:

   ```
   gsk8capicmd_64 -cert -validate -db cert.kdb -stashed -label ldap_server
   ```

   This command helps you check whether the necessary intermediate and root certificates are valid and have not expired.

Collecting trace information to debug issues with LDAP server

You can issue trace commands from the IBM Storage Protect server or server options file (`dsmserv.opt`) to collect the trace information. You can use the information that is exported to a trace log file for debugging the problem. You can also send the file to IBM Software Support for further analysis.

Procedure

Complete the following steps to get the trace information.

1. Perform one of the following options to get the trace information for the LDAP server.

   - To get the trace information from the IBM Storage Protect server, issue the following sequence of commands:

     ```
     export ldap_debug=65535
     export ldap_debug_file=ldaptrc.log
     ```

   - To get the trace information from the `dsmserv.opt` file, issue the following sequence of commands:

     ```
     trace enable ldap ldapcache
     trace begin trace.log
     ```

2. Start the IBM Storage Protect server.
3. Go to the server logs directory and find the trace file that you specified in step 1.

What to do next

You can also turn on and off the LDAP tracing while the IBM Storage Protect server is running by issuing the following commands:

```
ldaptrace enable <level=65535>
<recreate>
ldaptrace disable
```

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 IBM Storage Protect 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 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 IBM Storage Protect server key database file (`cert.kdb`) and verify that the certificates are not expired.

Table 1. Errors that might occur when you authenticate passwords with an LDAP directory server (continued)	
Error messages	**Resolution**
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.
ANR3114E LDAP error 50 (Insufficient access) ANR1885E LDAP directory service initialization: Permission was denied when the LDAP directory entry was accessed as LDAPUSER ANR3103E Failure occurred while initializing LDAP directory services ANR1885E SET LDAPPASSWORD: Permission was denied when the **LDAPUSER** entry was accessed	The **LDAPUSER** does not have full access control rights to the base DN that is specified in the **LDAPURL** option.
ANR3114E LDAP error 116 (Failed to connect to SSL server) ANR3116E LDAP SSL/TLS error 420 (Socket closed)	For Directory Server, the SSL_TIMEOUT_MILLISEC is not set high enough. See technote 1233758 .
ANR3114E LDAP error 4 (Size limit exceeded)	Increase the LDAP server search size limit to accommodate the total number of LDAP-authenticated nodes and administrators.
ANR3114E LDAP error 91 (Connection error) occurred during ldap_sasl_bind. ANR3103E Failure occurred while initializing LDAP directory services.	The LDAP server is not active or is offline.

Resolving client scheduling problems

The administrator for the IBM Storage Protect 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 Storage Protect schedules on numerous client computers.

Procedure

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

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

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

3. Issue the following query to limit the query results to exception cases:

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

Linux AIX If you use the traditional method to manage the scheduler, you can write a shell script to search for and stop running IBM Storage Protect schedulers or client acceptor processes, and then restart the processes. The following example shell script shows you how to recycle the IBM Storage Protect 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"
```

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

AIX 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
Acceptor 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 statis
```

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 Storage Protect 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 IBM Storage Protect 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 Storage Protect uses the files for internal processing.

If particular files must be included in the backup processing, IBM Storage Protect 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 on page 21 for the files that are automatically excluded.

Table 2. Files automatically excluded during backup processing

Platform	Files Excluded

Windows Windows	• 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
	Windows system files are silently excluded from the system drive backup processing and cannot be included.
	To process these Windows system files, you must issue a **DSMC BACKUP SYSTEMSTATE** command.
	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.
	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.
	`(partial contents of the dsmfino.txt file)`
	`==========`
	`SFP`
	`c:\windows\system32\ahui.exe (protected)`
	`c:\windows\system32\apphelp.dll (protected)`
	`c:\windows\apppatch\apphelp.sdb (protected)`
	`c:\windows\system32\asycfilt.dll (protected`

Linux AIX Linux AIX	Client trace file
Mac OS X Mac OS X	• 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:

Mac OS X Linux AIX
```
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:

```
Mac OS X Linux AIX 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:

```
Mac OS X Linux AIX exclude.dir /home/spike
```

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

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

```
Mac OS X Linux AIX 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:

```
Mac OS X Linux AIX 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:

```
Linux AIX exclude.dir /usr/
```

```
Mac OS X exclude.dir /Users/
```

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

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

```
Linux AIX exclude.dir /usr
```

```
Mac OS X exclude.dir /Users
```

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

Determining whether compression, encryption, and subfile backup statements are set to 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.

An INCLUDE statement that is associated with compression, encryption, and subfile backup implies only that if a specified file is a candidate for backup processing and matches the file pattern, the file is considered for compression, encryption, and subfile backup operations. To back up the file, you must remove the EXCLUDE statement from the specified file path.

Restriction: The subfile backup feature is deprecated for client files starting with IBM Storage Protect 8.1.0. However, subfile backup operations are still supported by IBM Tivoli Storage Manager 7 (all levels).

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:

```
Mac OS X Linux AIX exclude /usr/file.o
include.compress /usr/*.o
```

This statement indicates that the /usr/file.o file is excluded from backup processing. The INCLUDE.COMPRESS statement indicates that "if a file is a candidate for backup processing and matches the pattern /usr/*.o, compress the file". Do not interpret the INCLUDE.COMPRESS statement as "back up all files that match the pattern /usr/*.o and compress them". To back up the /usr/file.o file in this example, you must remove the EXCLUDE statement.

Consider the following example:

```
Windows exclude c:\Users\file.o
include.compress c:\Users\*.o
```

This statement indicates that the c:\Users\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 c:\Users*.o, compress the file". Do not interpret the INCLUDE.COMPRESS statement as "back up all files that match the pattern c:\Users*.o and compress them". To back up the c:\Users\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 *:\...\*
```

Mac OS X | Linux | AIX **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 Storage Protect 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
```

Windows | Linux | AIX Resolving Snapshot Difference problems

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

Prerequisites

To use the Snapshot Difference feature, you must first set up a NetApp user ID and password on the client. The user ID and password are necessary for IBM Storage Protect 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.

Restriction: If you upgrade to NetApp Data ONTAP 9.10.1 or later, snapshot difference incremental backups are no longer supported. Instead, use the following methods to protect NetApp filers:

- Run snapshot difference incremental backups with an earlier supported version of ONTAP until it reaches end-of-support by NetApp.
- Use NDMP volume-based backup.
- Run incremental backup operations without the **snapdiff** option.

For more information, see technote 6449354.

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. IBM Storage Protect 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 to 9.9.1
- Windows Common internet files system-attached (CIFS) volumes
- Both traditional and FlexVol filer volumes
- Java GUI
- Linux AIX 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.

Verifying the filer volume type

Windows

IBM Storage Protect 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 IBM Storage Protect from processing any files that use characters that are not in the 7-bit ASCII range. IBM Storage Protect 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. IBM Storage Protect 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. IBM Storage Protect inspects the string to see whether any character is outside of the 7-bit ASCII range. If so, IBM Storage Protect 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. IBM Storage Protect 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 `snapdiff` 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 IBM Storage Protect 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 IBM Storage Protect is not inspecting each file on the volume during the backup operation.

- You explicitly delete a file from IBM Storage Protect inventory by issuing the **DSMC DELETE BACKUP** command. The Snapshot Difference API does not detect that a file was manually deleted from IBM Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect to create a snapshot to use as the source (base) and a full incremental backup is performed.

Running a full incremental backup without the `snapdiff` option solves these limitations. IBM Storage Protect 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`

Linux AIX 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.
```

Linux AIX 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`.

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

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
```

Linux AIX Resolving image backup errors

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

Linux Resolving Linux image backup errors

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-ioctl (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
   ```

Linux Resolving backup failures when you use Linux snapshot image backup

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 IBM Storage Protect 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.

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

During IBM Storage Protect 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 IBM Storage Protect.

Before you begin

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

- The `Control-C` keystroke is issued during an IBM Storage Protect snapshot backup process. The JFS2 snapshot unmount request might fail with a "Device Busy" error, due to the IBM Storage Protect process being in the middle of accessing the snapshot.
- Two IBM Storage Protect 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 IBM Storage Protect 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 is configured to issue snapshot delete requests in a certain order, with the oldest snapshot deletion requested first, and the next oldest snapshot deletion requested next, and so on. If IBM Storage Protect cannot accept the sequence due to concurrent processes creating snapshots for the same file system, AIX fails the delete requests. In the previous examples, IBM Storage Protect 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. `unmount -f /tsm*`
4. `rmdir /tsm*`
5. `snapshot -d /dev/tsm*`

   If the snapshot delete process fails with "Device Busy" or some other error message, issue the `- unmount -f <srcfs>` command to unmount the source file system. 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 IBM Storage Protect process, IBM Storage Protect 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 IBM Storage Protect does not try to delete older snapshots:

- If the snapshot was not created by IBM Storage Protect, then IBM Storage Protect 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 IBM Storage Protect, an error message is generated that asks the user to delete the older snapshot and the operation again.

- If the snapshot is created by IBM Storage Protect but is still mounted, the snapshot is being used by some other IBM Storage Protect process.
- If the snapshot is created by IBM Storage Protect, is not mounted, but is newly created, the snapshot might have been created by some other IBM Storage Protect process.

In all such cases, you might have to perform a manual deletion. If any unused older snapshots are existing, subsequent IBM Storage Protect 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 IBM Storage Protect might stop during snapshot deletion and snapshot query processes. It might also cause data corruption during used-block-image backup. Therefore, IBM Storage Protect will not perform the following tasks:

- Snapshot monitoring
- Snapshot deletion

To use these features, ensure that your operating system level is at AIX 7.1 or later. For more information about supported operating systems, see the following technote: Overview - IBM Storage Protect Supported Operating Systems.

Support solutions for the IBM Storage Protect API

Resources are available to learn about or to diagnose the IBM Storage Protect 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 Storage Protect API?
- What is the exact level of the IBM Storage Protect server?
- What is the IBM Storage Protect server platform and operating system level?
- What is the exact level of the IBM Storage Protect storage agent (if LAN-free environment)?
- What is the IBM Storage Protect 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 Storage Protect application programming interface (API).

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

- The IBM Storage Protect 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 IBM Storage Protect administrator can view this log file for you if you do not have an IBM Storage Protect administrator user ID and password.
- If the API client is configured for LAN-free data movement, collect the options file for the IBM Storage Protect 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 IBM Storage Protect API function calls and are suspected of causing problems.
- The IBM Storage Protect 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`.

Mac OS X | Linux | AIX Verifying that the API uses the correct options file

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.pwdfl
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 31 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 storage agent rather than the server

You must know whether your data is being sent to the IBM Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect storage agent. If **totalLFBytesSent** is 0 0, the data was sent directly to the IBM Storage Protect 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 *Using the Application Programming Interface* for more details.

Mac OS X | Linux | AIX Running applications that use the API as a non-root user ID

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.
```

5. The root or authorized user must generate the TSM.PWD password file by using either the backup-archive client or the **dapismp** 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*:

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

6. If 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.

Windows Journal Based Backup problem determination

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

Windows Determining if a backup will be journal-based

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 IBM Storage Protect 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 IBM Storage Protect 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:\tsmjournal\tsmH__.jdb' for fs 'H:' has been
deleted and reset to the invalid state.
```

Windows Linux AIX Restarting a valid journal

You can increase performance by restarting a valid journal.

The reasons for restarting a valid journal:

- Error conditions in the journal daemon
  - Buffer overflow errors caused by excessive change activity on the journal file system being monitored for changes
  - Journal database access errors (disk full errors, etc.)
- Request by a backup client
- Clients will issue a journal restart request when it is determined that a journal file system lacks integrity for one of the following reasons:
  - Server 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)

Windows Running the journal daemon in the foreground

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

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

Windows The Journal Database Viewing utility

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 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
Storage Protect
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
Storage Protect
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 Syst em 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
```

Windows Using Windows Volume Shadow Copy Services

The IBM Storage Protect 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.

Windows Defining VSS transient errors

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

Windows Defining Windows VSS test flags

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 IBM Storage Protect:

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

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

Option file example:

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

Windows Volume Shadow Copy Services tuning

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.

Windows Gathering VSS diagnostic information for Microsoft assistance

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 Storage Protect, 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.

AIX Collecting a Windows VSS trace by using the VSS trace tool

In some cases, a Windows Volume Shadow Copy Service (VSS) trace can be useful for diagnosing issues on Windows operating systems.

Before you begin

When error messages indicate issues with Microsoft VSS or when further information is required to help diagnose problems involving Microsoft VSS, you can use the VSS trace tool developed by Microsoft to collect traces and generate additional log files. Before you install the VSS trace tool, you must download a Windows Software Development Kit (SDK).

Restriction: Because Windows SDKs are developed by Microsoft, diagnosing or troubleshooting issues related to Windows SDKs and VSS traces are outside the scope of IBM Software Support.

To download and install the VSS trace tool, complete the following steps:

1. Download the Windows 10 SDK. Go to the Microsoft Developer website and search for "Windows 10 SDK".

   Tip: Microsoft asks customers to download Visual Studio packages along with the SDK. If you do not install either the SDK or Visual Studio package, you must contact Microsoft support to obtain the VSS trace tool.
2. Run the SDK installer. During the installation process, you must, at a minimum, select the following features:
   - Windows SDK Signing Tools for Desktop Apps
   - Windows SDK for UWP Managed Apps
   - Windows SDK for UWP C ++ Apps
   - Windows SDK for Desktop C++ x86 Apps
   - Windows SDK for Desktop C++ amd64 Apps

   Tip: The feature names might vary slightly, depending on the SDK version that you install.
3. Click **Install**.
4. To verify the installation, ensure that the vsstrace.exe executable file is available in the following directory: `C:\Program Files (x86)\Windows Kits\10\bin\10.0.19041.0\x64`. This file path can vary based on the Windows version release that you are using at the time of installation.

Procedure

To use the VSS Trace (vsstrace.exe) tool to create a trace of a VSS action, complete the following steps:

1. Open a Windows command prompt.
2. Create a directory for the trace file. For example, if you want to store a trace file in the `C:\VSS_trace` directory, issue the following command:

   ```
   mkdir C:\VSS_trace
   ```
3. Change to the directory where the vsstrace.exe executable file is located. Typically, the file is located in the default directory that was created during the installation process:

   ```
   C:\Program Files (x86)\Windows Kits\10\bin\10.0.19041.0\x64
   ```
4. Run the vsstrace.exe executable file by issuing a command that is similar to the following example:

   ```
   vsstrace -f 255 -o C:\VSS_trace\vsstrace.log +indent
   ```

   The VSS trace process is visible in the foreground, and the trace log is similar to the following example:

   ```
   C:\Program Files (x86)\Windows Kits\10\bin\10.0.19041.0\x64>vsstrace -f 255 -o C:\
   VSS_trace\vsstrace.log +indent
   Session started. Enabling provider...
   Enabling provider {9138500e-3648-4edb-aa4c-859e9f7b7c38}on level 170 and /la, 0x000000ff...
   Real-time tracing now in progress. Press Ctrl-C to stop tracing...
   ```
5. Recreate the VSS issue.
6. In the vsstrace.exe command prompt, press Ctrl+C. The tracing ends with a message that is similar to the following example.

   Tip: In the following example, the start and end times are the same (19:46:10.001), but the last message is time-stamped 19:46:45.107.

   ```
   [19:46:45.107 P:0ABB T:0288 CORSVCC(0924) COORD]
   ENTER[CVsServiceModule::Unlock]
   [19:46:45.107 P:0ABB T:0288 CORSVCC(0933) COORD] VSSVC: Idle period begins
   [19:46:45.107 P:0ABB T:0288 CORSVCC(0924) COORD]
   EXIT[CVsServiceModule::Unlock] Time spent: 00:00:00-0000; total: 0; HRESULT: 0
   Tracing stopped.

   Total Events Processed 1719
   Start Time 3/22/2021 19:46:10.001 (0x01071F7588F39790)
   ```

```
End Time 3/22/2021 19:46:10.001 (0x01D71F759DE047E5)
Elapsed Time 35 sec
+------------------------------------------------------------------------+
EventCount    Event Name     EventType Guid
+------------------------------------------------------------------------+
1719          Default        {77d8f687-8130-4a14-b8a6-3b922e05b99c}
+------------------------------------------------------------------------+

C:\Program Files (x86)\Windows Kits\10\bin\10.0.19041.0\x64>
```

7. Gather the trace file that was specified on the **VSSTRACE** command.

Example

The following sample trace file shows that, a simulated disk-full condition was created by changing the `C:` drive's shadow storage to the EFI system partition and a manual snapshot of the C drive was taken with Windows shadow copy functionality. This operation failed with error `0x8004231f`, as shown in the trace file:

```
[19:46:24.905 P:0330 T:0DAC SPRQSNPC(0690) SWPRV].
ENTER[CVssQueuedSnapshot::MarkAsProcessingPrepare]
[19:46:24.905 P:0330 T:0DAC SPRQSNPC(0548) SWPRV].
ENTER[CVssQueuedSnapshot::GetSnapshotProperties]
[19:46:24.905 P:0330 T:0DAC SPRQSNPC(0548) SWPRV].
EXIT[CVssQueuedSnapshot::GetSnapshotProperties] Time spent: 00:00:00-0000; total: 0; HRESULT: 0
[19:46:24.905 P:0330 T:0DAC SPRQSNPC(0690) SWPRV].
EXIT[CVssQueuedSnapshot::MarkAsProcessingPrepare] Time spent: 00:00:00-0000; total: 0; HRESULT:
0
[19:46:24. 905 P:0330 T:0DAC SPRQSNPC(0083) SWPRV].
ENTER[CVssQueuedSnapshot::OpenVolumeChannel]
[19:46:24. 905 P:0330 T:0DAC SPRQSNPC(0548) SWPRV].
ENTER[CVssQueuedSnapshot::GetSnapshotProperties]
[19:46:24.905 P:0330 T:0DAC SPRQSNPC(0548) SWPRV].
EXIT[CVssQueuedSnapshot::GetSnapshotProperties] Time spent: 00:00:00-0000; total:0; HRESULT: 0
[19:46:24. 905 P:0330 T:0DAC SPRQSNPC(0083) SWPRV].
EXIT[CVssQueuedSnapshot::OpenVolumeChannel] Time spent: 00:00:00-0000; total: 0; HRESULT: 0
[19:46:24. 905 P:0330 T:0DAC SPRQSNPC(0172) SWPRV].
ENTER[CVssQueuedSnapshot::PrepareForSnapshotloctl]
[19:46:24. 905 P:0330 T:0DAC SPRQSNPC(0548) SWPRV].
ENTER[CVssQueuedSnapshot::GetSnapshotProperties]
[19:46:24. 905 P:0330 T:0DAC SPRQSNPC(0548) SWPRV].
EXIT[CVssQueuedSnapshot::GetSnapshotProperties] Time spent: 00:00:00-0000; total: 0; HRESULT: 0
[19:46:24.908 P:0330 T:0DAC INCICHLH(0521) SWPRV].
THROW[CVssQueuedSnapshot::PrepareForSnapshotloctl] ERROR_DISK_FULL detected.
DeviceioControl(\ \?\ Volume{1/800977-b3c7-4832-bdd0-29c3e99967d8}
-0000000000000210,0x0053c008,000001A035718120,0,000001A03571D140,4096, [OJ)
[19:46:24.908 P:0330 T:0DAC SPRQSNPC(0172) SWPRV].
EXIT[CVssQueuedSnapshot::PrepareForSnapshotloctl] Time spent: 00:00:00-0016; total: 0x10;
HRESULT: 0x8004231f
[19:46:24.908 P:0330 T:0DAC SPRALLOC(0187) SWPRV].
ENTER[CVssDiffAreaAllocator::~CVssDiffAreaAllocator]
[19:46:24. 908 P:0330 T:0DAC SPRALLOC(0187) SWPRV].
EXIT[CVssDiffAreaAllocator::~CVssDiffAreaAllocator] Time spent: 00:00:00-0000; total: 0;
HRESULT: 0
[19:46:24.908 P:0330 T:0DAC SPRPROVC(0654) SWPRV] HRESULTexception caught: hr: 0x8004231/
[19:46:24.908 P:0330 T:0DAC SPRPROVC(0515) SWPRV]
EXIT[CVsSoftwareProvider::EndPrepareSnapshots] Time spent: 00:00:00-0063; total: 0x3f; HRESULT:
0x8004231f
[19:46:24.909 P:0AB8 T:1214 CORSNPSC(1754) COORD]......
THROW[CVssSnapshotSetObject::EndPrepareAllSnapshots] insufficient di// area storage detected
while calling EndPrepareAllSnapshots. Provider ID = {b5946137-7b9/-4925-a/80-51abd60b20d5}
[19:46:24.910 P:0AB8 T:1214 CORSOFTC(0160) COORD]
ENTER[CVssSoftwareProviderWrapper::Release]
```

Windows Troubleshooting errors using a VSS trace

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:ffffffff
"TraceEnterExit"=dword:00000001
"TraceToFile"=dword:00000001
"TraceToDebugger"=dword:00000000
"TraceFileLineInfo"=dword:00000001
"TraceForceFlush"=dword:00000000
```

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

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 IBM Storage Protect encounter the same problem. If **vsreq** can reproduce the same problem as IBM Storage Protect, 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**.

Windows Comparing IBM Storage Protect and Ntbackup.exe interaction with VSS

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

The known difference between `Ntbackup.exe` and IBM Storage Protect 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. IBM Storage Protect 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 is displayed by a **SHOW** command, 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 on page 44 are not all of the available **SHOW** commands.

Table 3. SHOW commands for the backup-archive client

SHOW Command	Description	Information
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.
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.

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

You can use IBM Storage Protect 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 IBM Storage Protect 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" on page 45
You get SOAP messages 'content not available on server' triggered by VMware	"Resolving 'Content not available' error messages during VMware operations" on page 46
You can see only inactive copies of SQL databases when you use the Data Protection for SQL GUI or the **tdpsqlc** command.	"Viewing active copies of Microsoft SQL databases" on page 47
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 47
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 47
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 48
You want to view the status of VSS writers within the guest virtual machine.	"Determining whether a virtual machine backup might fail" on page 48

Windows Resolving database access problems

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:

i) Ensure that the 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.

ii) 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   dsmagent_node   *      \VMFULL-vm_name\*\*
```

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.

Resolving 'Content not available' error messages during VMware operations

During VMware operations, SOAP messages might be issued and logged on the IBM Storage Protect client. Typically, these messages are issued during VMware restore operations. The messages report issues with missing or damaged data, or other errors with the VMware restore operation.

Although the messages are reported in the IBM Storage Protect client, the messages are triggered by the VMware product and are passed to the IBM Storage Protect server. For more information about SOAP messages, see the VMware documentation.

To understand the reported errors, examine the details of the SOAP messages.

To help resolve the reported issues, complete the following steps:

1. Issue the **QUERY ACTLOG** command to display messages that were generated by the server during VMware operations.
2. In the activity log, look for messages that indicate that data or volumes are not available to the server. If you find such messages, check the mounted volumes that were listed in error messages at the time of failure. If necessary, update the volumes to make them available and accessible.
3. If you still get failure messages, check the storage pool or volumes that contain the data for any damaged objects or contents.

   Tip: Damaged data objects in storage pools are the most likely cause of SOAP messages.

4. Repair any damaged data objects to make them available again. You might be able to recover the damaged data objects from a copy storage pool or by repairing the storage pool, or by retrieving the data from a target replication server. For more information about repairing damaged data in an IBM Storage Protect environment, see *Repairing and recovering data* in IBM Documentation.

Windows Viewing active copies of Microsoft SQL databases

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.

Windows Microsoft SQL databases with DBCS names

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.

Windows Responses to messages for virtual machine backups with application protection

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 IBM Storage Protect at the IBM Support Portal for IBM Storage Protect.

Windows Saving VSS XML manifest files

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\tsmvmbackup\fullvm\vmtsmvss\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.

Windows Determining whether a virtual machine backup might fail

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 IBM Storage Protect server problems

When working with IBM Storage Protect, 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
- Examine the job information
- 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:

  Linux AIX *`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 IBM Storage Protect is installed. For example:

Linux AIX `/opt/tivoli/tsm`

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

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

Examining the job information

Use the **QUERY JOB** command to display information about one or more jobs. The information can include job IDs and statuses. You can query jobs after starting a copy data, tiering, or replication storage rule. Additionally, you can query retention set jobs, including jobs for copying a retention set to tape storage. You can filter the list of jobs by specifying a job ID or other job attributes that are displayed in the command output information.

A log is generated after every job run. The log displays the information such as the status of the job, the start and end time for the job, and a message that is associated with the job. You can also use the **QUERY ACTLOG JOB=n** command option to see the messages related to the job. This provides additional information about the symptoms for the problem or might provide information about the actual cause of the problem that the job encountered.

For more information about command options, see *QUERY JOB* in IBM Documentation.

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 Storage Protect 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 Storage Protect 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 IBM Storage Protect 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.
- 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.

Linux AIX Configuring a server instance to use shared memory

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

Before you begin

Tip: Windows For Windows, in addition to the TCP/IP communication protocol, you can set and use the Named Pipes communication method when you backup or restore databases. When you run the server and client on the same Windows machine, the Named Pipes method will help to avoid loopback problems. For more information about setting the Named Pipes communication method, see *Configuring server and client communications* in IBM Documentation.

Procedure

Complete the following steps to update the database backup node configuration for your server to enable shared memory:

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

```
COMMMethod SHAREdmem
SHMPort         1510
```

2. Locate the `dsm.sys` client API system options file, which is stored in the following location by default:

server_bin_directory/dbbkapi/dsm.sys

3. Modify the stanza for the database backup node in the file, dsm.sys.
   a. Remove the following lines from the stanza:

```
COMMMethod TCPip
TCPServeraddress 127.0.0.1
TCPPort 1500
```

   b. Add the following lines to the stanza:

```
COMMMethod SHAREdmem
SHMPort 1510
```

   c. Save and close the file.

Changing the copy frequency

IBM Storage Protect 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 VTLPOOL.
ANR8847E No LTO-type drives are currently available in library IBMVTL.
ANR8801I LABEL LIBVOLUME process 25 for library IBMVTL completed; 0 volume(s)
labeled, 0 volume(s) checked-in.
ANR0985I Process 25 for RELABEL running in the BACKGROUND completed with
completion state SUCCESS at 22:10:36.
```

To resolve a RELABEL error, complete the following steps:

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

   If a RELABEL operation cannot obtain a drive or fails to relabel a volume, IBM Storage Protect 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 IBM Storage Protect 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
ANR3174E Communication error with managed server ADMIN.
ANR0484W Session 2 for node ADMIN 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, IBM Storage Protect 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 or IBM Cloud Object Storage.
   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 IBM Cloud Object Storage.

      Tip: If you are using IBM Cloud Object Storage as your object storage system, log in to IBM Cloud Object Storage and click the **Security** tab. In the **dsNet Fingerprint** section, click **dsNet certificate authority** and copy the certificate information into a certificate file for Part 2.

   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 IBM Cloud Object Storage 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 Storage Protect 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

There are two methods of tracing that are available for the backup-archive client.

Resolving installation and upgrade problems

Resolving installation problems with the IBM Storage Protect 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:

Linux AIX `/var/ibm/InstallationManager/logs`

Windows `C:\ProgramData\IBM\Installation Manager\logs`

Installation wizard fails to start

AIX 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 IBM Storage Protect server, the installation might not start. An error about missing gtk libraries is issued.

Related information

Installing IBM Storage Protect by using the installation wizard

Resolving GSKit installation problems

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

If the IBM Storage Protect 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.

Linux 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:

- Linux `/usr/local/ibm/gsk8_64/lib64`

AIX For AIX, issue the following command:

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

Linux AIX 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
```

- 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
```

- 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
```

- Linux

```
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 Storage Protect 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:

- Linux | AIX 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.

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 Storage Protect 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 IBM Storage Protect.

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`.
   - Linux 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 V*xxxx* 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 on page 60 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.

For instructions about migrating the server to a different operating system, see IBM Storage Protect upgrade and migration process: Frequently asked questions.

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.

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 IBM Storage Protect 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 IBM Storage Protect environment:

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

  You can use the `db2support.zip` file that is generated for troubleshooting.
- In addition to the output from the listed commands, or in the cases of a server that cannot respond to commands, collect a dump. The way that you collect a dump depends on the operating system.
  - Linux AIX 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

If the IBM Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect server.
2. AIX Issue the following commands:
   a) `db2set DB2_ALTERNATE_GROUP_LOOKUP=GETGRSET`
   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

Linux AIX 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`.

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).

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`.

- Linux Copy the `/var/log/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.

Windows Detecting errors after a server service starts and stops

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 64.
- 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 64 to server with the client.
  3. Delete the files that are listed in Table 6 on page 64.

Table 6. Files that can be deleted after these files are archived	
File name	**Description**
`instance_name.nfy` `instance_name.`*n*`.nfy` (where *n* is a number)	Administration notification logs
`db2dasdiag.log`	Db2 administration server (DAS) diagnostic log
`db2eventlog.`*xxx* (where *xxx* is the database partition number)	Db2 event log
nnnnnnn`.`*nnnnn*`.`*nnn*`.dump.bin` (where *n* is a number)	Binary dump files of key in-memory structures
nnnnnnn`.`*n*`.`*nnn*`.trap.txt` (where *n* is a number)	Trap files
nnnnnnn`.`*nnnnn*`.`*nnn*`.apm.bin` (where *n* is a number)	Access the plan-manager binary dump files
nnnnnnn`.`*nnnnn*`.`*nnn*`.stack.txt` (where *n* is a number)	Stack traces
`FODC_`*xxxx*`/core<pid>`	Core files These `FODC_`*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 *n* 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
The `db2diag.log` file contains diagnostic information that can help you to resolve problems that might occur with your database.

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 IBM Storage Protect 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 IBM Storage Protect deploys.

For more information, see the Db2 product information.

Resolving database manager starting problems

The IBM Storage Protect 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, OSSHLibrary::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 _ZN11OSSHLibrary4loadEPKcm + 0x1D3
  [4] 0x00002AEF63F63BDC _Z20secLoadPluginGenericP19SEC_PLUGIN_HANDLE_TPc + 0x68
  [5] 0x00002AEF63F62FBB _Z23secLoadServerAuthPluginP19SEC_PLUGIN_HANDLE + 0x57
  [6] 0x00002AEF63F6C833 _Z25sqlexLoadAllPluginsServerP5sqlca + 0x3B5
  [7] 0x00002AEF6431737C /home/hannigan/sqllib/lib64/libdb2e.so.1 + 0x123637C
  [8] 0x00002AEF643164C5 sqloRunInstance + 0x191
  [9] 0x000000000040D31D 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/sqllib/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 _Z25sqlexLoadAllPluginsServerP5sqlca + 0x3B5
[7] 0x00002AEF6431737C /home/hannigan/sqllib/lib64/libdb2e.so.1 + 0x123637C
[8] 0x00002AEF643164C5 sqloRunInstance + 0x191
[9] 0x000000000040D31D 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:

Linux AIX

1. Ensure that the server has write authority to the ~/sqllib/db2dump/ directory.
2. Add the following text to the ~instance/sqllib/userprofile file:

   ```
   export DB2_DSMDB2PW_TRACEFILE=filename
   ```

   where *filename* is a fully qualified path and file name of the trace file, for example ~/sqllib/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_DSMDB2PW_TRACEFILE=c:\Program Files\Tivoli\TSM\s1\sqllib\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 IBM Storage Protect 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 Storage Protect 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.

Linux AIX
```
> 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).

Linux AIX 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 TSMSERVER1, the **db2set deb2instprof** command shows this path:

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

   The diagnostic log file is in the TSMSERVER1 sub directory:

   ```
   C:\ProgramData\IBM\DB2\DB2TSM1\TSMSERVER1
   ```

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

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

Db2 upgrade log files

When you upgrade the IBM Storage Protect 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:

- Linux AIX /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 Storage Protect 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 the 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 the Tivoli Storage Manager V6.1 server and then restore the database.

If you must restore an 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 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 IBM Storage Protect 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**.

   Linux AIX 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`
- Linux AIX `tsm/client/api/bin64/sample/dsmrc.h`

Linux AIX 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

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.

The DSMI_* variables are set in the instance's `userprofile` file. For example: `/home/tsminst1/sqllib/userprofile`

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

Procedure

Open the `/home/tsminst1/sqllib/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
commmethod 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 IBM Storage Protect 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 an IBM Storage Protect 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 IBM Storage Protect 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 bàckup/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=SQLO_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 Storage Protect 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 IBM Storage Protect client API executable files. To obtain information about the error, open a Db2 command line processor window and issue the following command:

```
db2 => ? sql2071
```

Verify that if any changes were made to the `tsmdbmgr.opt`, `dsm.sys`, or `sqllib/userprofile` files, the Db2 instance is recycled so that it picks up the new values. To recycle the Db2 instance, stop and start the IBM Storage Protect server. Also, verify that the **EXPORT** command precedes the `DSMI_*=` entries in the `sqllib/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
You can configure a server instance to use shared memory to resolve slow database backups that can occur because of Transmission Control Protocol (TCP) loopback problems.

Characteristics of the $$_TSMDBMGR_$$ user ID

The IBM Storage Protect 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 IBM Storage Protect 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. IBM Storage Protect reorganizes tables by default. If automatic reorganization is affecting server performance, you can schedule reorganization manually.

The following suggestions might help when you configure your reorganization:

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

Analyzing the process symptoms to resolve problems

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

You might encounter one of the following process symptoms:

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

Reviewing process messages to determine the state of server operations

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

Processes that run on the server

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

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

Many server processes can be run in the FOREGROUND or synchronously. Processes that run in the FOREGROUND can be initiated by a command using the WAIT=YES parameter. Commands that start server processes that do not allow the WAIT=YES parameter or commands specified with WAIT=NO are run in the BACKGROUND or asynchronously.

Some server processes can initiate multiple processes simultaneously to accomplish the task. See Table 7 on page 75 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.	

Table 7. Server processes (continued)

Process or command	Description	Runs in the foreground or as a multiple process
Expiration	Delete client backup and archive files from the server, based on the policies defined to manage those files. You can run expiration automatically by specifying EXPINTERVAL=*n* in the server options file, where *n* 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.	The **EXPIRATION** command can run as a synchronous process by specifying WAIT=YES.
IMPORT	Import data from sequential media volumes or directly from another server using TCP/IP communication connections between the servers. Import processing can be started by any of the following commands: • **IMPORT ADMIN** • **IMPORT NODE** • **IMPORT POLICY** • **IMPORT SERVER**	
LABEL LIBVOLUME	Label one or more library volumes in a library.	
Migration	Migrate data from one storage pool to the next in the storage hierarchy. 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.	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.

Table 7. Server processes (continued)

Process or command	Description	Runs in the foreground or as a multiple process
MOVE DRMEDIA	Manage the disaster recovery media by moving on-site volumes off-site, or by bringing back off-site volumes, on-site. Disaster recovery media is the database backup and storage pool backup volumes necessary to protect and recover the server.	The **MOVE DRMEDIA** command can run as a synchronous process by specifying WAIT=YES.
MOVE MEDIA	Move volumes from a tape library to the overflow location to prevent a library from becoming full.	
MOVE NODEDATA	Move all the data for the node or nodes specified to other volumes in the same storage pool or to a different storage pool.	The **MOVE NODEDATA** command can run as a synchronous process by specifying WAIT=YES.
PREPARE	Create a recovery plan file.	The **PREPARE** command can run as a synchronous process by specifying WAIT=YES.
Reclamation	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
FOREGROUND or **BACKGROUND**. If the process is running in the foreground, the command was issued with the **WAIT=YES** parameter. Foreground processing causes the administrative session that issued the command to wait until the processing completes. A process running in the background returns immediately to the administrative session that issued the command, indicating that a process was started while the process still runs. Processes running in the background might be monitored with the **QUERY PROCESS** command.

time
The time that the process was started.

Messages issued when processes end

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

Process ended

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

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

The following list defines the variables from this message:

process_id
Numeric process identifier.

process_name
The name of the process.

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

completion_state
SUCCESS or FAILURE.

time
The time that the process was started.

Process ended with items and bytes

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

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

The following list defines the variables from this message:

process_id
Numeric process identifier.

process_name
The name of the process.

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

number_of_items
The number of items processed.

bytes_processed
The number of bytes processed.

completion_state
SUCCESS or FAILURE.

time
The time that the process was started.

Process ended with items

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

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

The following list defines the variables from this message:

process_id
Numeric process identifier.

process_name
The name of the process.

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

completion_state
SUCCESS or FAILURE.

time
The time that the process was started.

Process ended with bytes

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

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

The following list defines the variables from this message:

process_id
Numeric process identifier.

process_name
The name of the process.

process_state
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.

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 Storage Protect 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* **DISCARDDATA=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 an IBM Storage Protect 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, IBM Storage Protect 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

IBM Storage Protect 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 IBM Storage Protect.

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 IBM Storage Protect, 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 IBM Storage Protect 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 Storage Protect 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 IBM Storage Protect to function correctly. Storage pool errors can be related to the following issues:

- Failed transactions
- A storage pool experiencing a high volume usage after increasing the MAXSCRATCH value
- A storage pool having "Collocate?=Yes" but volumes still containing data for many nodes
- Unable to store data in an active data pool by using the simultaneous-write function or by issuing the **COPY ACTIVEDATA** command

"ANR0522W Transaction failed..." message received

The ANR0522W message is displayed when the server is unable to allocate space in the storage pool that is identified to store data for the specified client.

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 FILESPACE), 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 IBM Storage Protect, you can backup 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. For more information, see technote 3241809.

Use the following guidance to resolve issues and handle limitations:

Removing data marked as damaged or orphaned during an audit

A damaged data extent is a file that has references in the server database, but has missing or corrupted data on the cloud. An orphaned data extent is an object stored in a cloud service provider that does not have a reference in the server database.

To remove missing or corrupted data on the cloud-container storage pool, issue the **AUDIT CONTAINER** command with the **VALIDATECLOUDEXTENTS=YES** parameter setting. When you issue this command, the command reads the data from the cloud-container storage pool and validates individual extents in addition to the consolidated metadata in the container.

To delete orphaned extents from the cloud-container storage pool, issue the **AUDIT CONTAINER** command with the **FORCEORPHANDBDELETE** parameter. When you issue this command, the command enables the server to force the deletion of orphaned extents from the server database, even if they are not deleted from the cloud-container storage pool.

For more information, see the **AUDIT CONTAINER** command in IBM Documentation.

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
- 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 IBM Storage Protect 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 Data Protection for VMware control files and Data Protection for SQL metadata files (for legacy SQL backups). For more information, see the following documents:

- Using tape, VTL, or container storage pools with IBM Storage Protect for Virtual Environments, technote 1659833
- IT11763: METADATA CONSIDERATIONS ARE MISSING IN DATA PROTECTION FOR SQL SERVER DOCUMENTATION.

Verifying synchronization between source and target replication server

The purpose of replication is to maintain the same level of files on the source and the target replication servers so that data can be restored from the target server if necessary. When data is replicated from the source replication server to the target replication server, you might want to verify if both the servers are synchronized. It is necessary that you understand the processes that are involved in verifying synchronization between the source and target replication server. You must also understand the ways to troubleshoot for reasons if the servers are not synchronized.

View client replication status

On the **Client Replication** page, you can view client replication status to verify that the client file spaces were successfully replicated to target replication servers and that the replicated data is current.

Note: The **Client Replication** page displays file spaces regardless of whether they were replicated by a replication storage rule or by using the **REPLICATE NODE** command.

To view client replication status, complete the following steps:

1. On the Operations Center menu bar, click **Clients** > **Client Replication**.

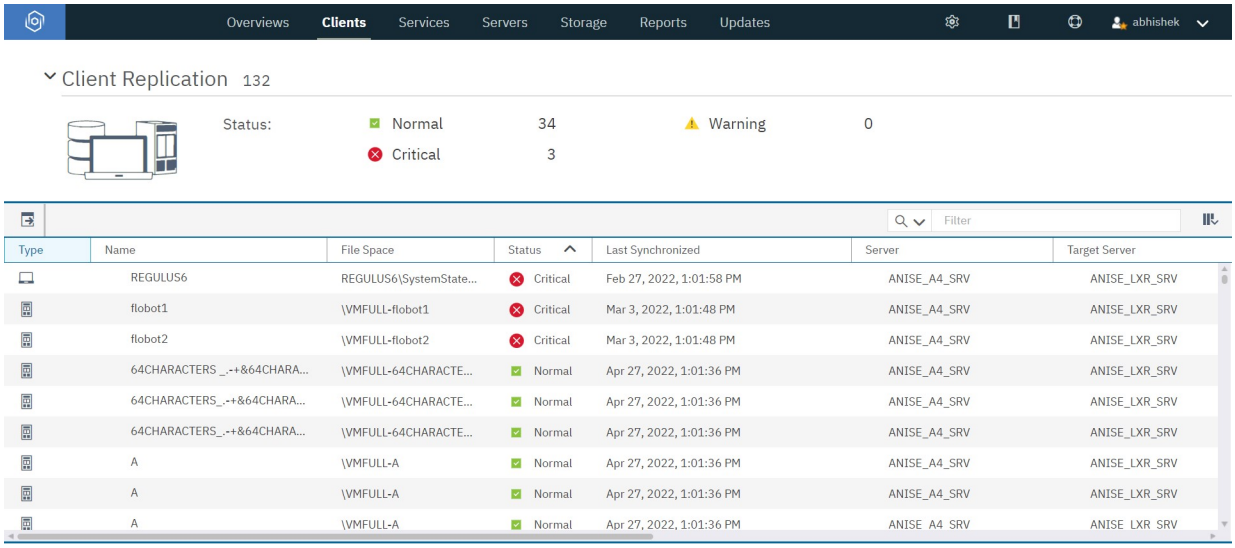

2. Review the information that is shown in the table. Each row of the table shows a separate file space that is replicated to a target replication server. To verify that a file space was successfully replicated and that the replicated data is current, review the information in the Status and Last Synchronized columns.

Verify data synchronization between source and target replication servers

You can verify if both servers are in sync by issuing **QUERY REPLICATION** command. Specify the name of the client node on the source replication server, name of the target replication server, and the parameter **FORMAT=DETAILED**.

To verify synchronization between the client node, NODE1 on the source replication server and the target replication server PHOENIX-DR, issue the following command:

```
query replication node1 server=phoenix-dr format=detailed
```

From the query output, you need to focus on the following fields:

Total Files To Replicate
Displays the total number of files to replicate to the target replication server. The value must be zero.

Total Files Not Replicated Due To Errors
Displays the total number of files that were not replicated because of errors. The value must be zero.

Total files Not Yet Replicated
Displays the total number of files that are not yet replicated to the target replication server. The value must be zero.

If any of the values in the fields that are mentioned is not equal to zero, it would indicate that the source and target replication servers are not in sync. For more information about **QUERY REPLICATION** command, see *Guidelines for reducing database size by pruning the SDRC table* in IBM Documentation

View information about the failed replication

The **QUERY REPLICATION** command output might display information on files that failed to replicate from the source replication server to the target replication server. You can obtain more details on the files that failed to replicate and the reason for the failure by issuing the **QUERY REPLFAILURES** command.

When you issue the command, specify the name of the client node that you want to query, the name of the target replication server, and the parameter **TYPE=OBJECTS**.

To obtain the list of files that failed to replicate from the client node, NODE1 on the source replication server to the target replication server PHOENIX-DR, issue the following command:

```
query replfailures node1 server=phoenix-dr type=objects
```

From the query command output you can view the names and IDs of the files that failed to replicate, the error code, and the explanation of the errors.

Check the server activity log file and other log files

To have a better understanding of the replication error, you might want to obtain detailed information about the recent server activities. Also, you might need log filestoraise a service request with IBM Software Support.

To review the messages in the server activity log, issue the **QUERY ACTLOG** command.

To view additional information about the cause of the problem and to find ways to resolve it, you can refer to 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`

Windows Linux AIX Chapter 4. Resolving Operations Center problems

If a problem occurs with the IBM Storage Protect 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.

Windows Linux AIX Log files overview

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:

  Linux AIX *`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 IBM Storage Protect is installed. For example:

Linux AIX `/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
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.

Windows Linux AIX Viewing the Operations Center log from within the Operations Center

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
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.

Object agent is not shown in the Operations Center

If you have defined an object agent on the server by issuing the **DEFINE SERVER** command or you modified an object agent by issuing the **UPDATE SERVER** command, the command completes without errors. And you can verify that the object agent was defined by issuing the **QUERY SERVER** command. However, the object agent is not shown in the Operations Center.

Symptom

On the **Object Agent** page for the server, a message is displayed that says that an object agent must be configured before you register an object client. If you try to register a client by using the **Add Client** wizard, you also get a message that says that no object agent is configured. If you try to define an object agent by using the Operations Center, the following error message is written to the activity log:

```
ANR4610E Command: An object agent server is already defined.
```

Cause

This happens when you define an object agent by using the Operations Center, the Operations Center defines an endpoint to which object clients can send S3 requests. This endpoint is formed by using the IP address of the server and a port number. In the Operations Center, you can configure only the port number. The IP address is automatically determined to be the same as the server's IP address. When the

object agent was created or modified by issuing the **DEFINE SERVER** or **UPDATE SERVER** command, a high-level address was specified that does not match the IP address of the server. For this reason, the Operations Center incorrectly determines that no object client is defined on the server.

Solution

To resolve this problem, issue the **UPDATE SERVER** command and specify the **HLADDRESS** parameter to use the IP address of the IBM Storage Protect server.

Usage and configuration information is not shown for a container storage pool

In the Operations Center, sometimes you cannot find the detailed information about a container storage pool.

Symptom

To display the detailed information about a container storage pool in the Operations Center, you select the storage pool on the Storage Pools page and click **Details**. On the Summary tab of the Details notebook, no information is shown in the Usage and Configuration area. In the Usage and Configuration area, the status message **Loading...** is displayed. Eventually, the operation times out and the Usage and Configuration area displays the message `Cannot retrieve data`. The following error is displayed in a separate window: `Error occurred while making the Web server request: Timeout exceeded`

Cause

This problem occurs if there are many containers in the storage pools that are defined on the server. This condition is most likely when there are cloud-container storage pools defined on the server.

Solution

Although the usage and configuration information is not available on the Summary tab of the Details notebook, you can view some of the same information from other tabs of the Details notebook. On the Properties tab of the Details notebook, you can view the following information for a storage pool:

- Capacity and utilization measurements
- The space savings that were realized by compression and deduplication
- The overflow pool name
- The protection pool name
- The name of the tiering storage rule

On the Directories tab of the Details notebook, you can view the file system directories that are associated with the container storage pool. For directory-container pools, the directories define the storage pool size. For cloud-container pools, the directories define the size of a disk cache that is used to optimize data transfer to the cloud.

Storage rule summary pages are not shown in the Recent History area

In the Operations Center, sometimes you cannot find the storage rule summary pages in the two-week history chart or the Recent History area.

Symptom

To view the performance history of a storage rule in the Operations Center, you select the storage rule on the Storage Rules page and click **Details**. From the storage rule's details notebook, you can see the

summary page for the storage rule, a subrule, or a source pool. The two-week history chart is empty or is missing the most-recent data. The Recent History area of the page is also missing data.

Cause

In some situations, background processes for identifying eligible file spaces to be copied can degrade the performance of a storage rule. As a workaround to improve performance, the CHECKFORTIERSTATS server option is set to No. A side effect of that workaround is that the Operations Center cannot populate storage-rule summary pages with recent history. The information that is normally collected in the background to populate the two-week history chart and the Recent History area is not collected while the CHECKFORTIERSTATS server option is set to No.

Solution

To display the current setting of the CHECKFORTIERSTATS server option, issue the QUERY OPTION command. From the command line, issue the following command:

```
q opt checkfortierstats
```

If the CHECKFORTIERSTATS server option is set to No, you can resolve this problem by setting the server option to Yes. If you set the CHECKFORTIERSTATS server option to Yes, the recent history is not immediately restored, but begins to be collected again the next time the storage rule runs. If the server option was previously set to No to resolve a performance problem, understand that setting it to Yes might degrade storage rule performance. You might need to choose between improved performance or displaying the two-week historical data.

To set the CHECKFORTIERSTATS server option, issue the following command:

```
setopt checkfortierstats yes
```

Windows Linux AIX Alerts are not updated immediately

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 94 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 IBM Storage Protect
1	6 seconds	7 seconds
10	6 seconds	7 seconds
100	6 seconds	8 seconds
1,000	10 seconds	20 seconds
10,000	45 seconds	1.25 minutes

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.

Windows Linux AIX Active tasks are not canceled immediately

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 on page 95 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.

Windows Linux AIX Further known issues with the Operations Center

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 IBM Storage Protect Operations Center.

Chapter 5. Resolving communication problems

The need for connectivity in IBM Storage Protect 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 TCPPORT; 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 be locked in following scenarios:
  - If server administrator issues the LOCK NODE administrative command.
  - If someone attempts to login by using the wrong password more than the number of times specified by the SETINVALIDPWLIMIT command. For more information about the number of invalid logon attempts, see *SET INVALIDPWLIMIT (Set the number of invalid logon attempts)* in IBM Documentation.

  To view the messages in the server activity log, issue the QUERY ACTLOG command. This might provide additional information about the actual cause to lock the node. If the node is locked due to invalid logon attempts, evaluate to identify that those attempts are made by authorized user or unauthorized user. If the attempts are made by unauthorized user consider to change the password before unlocking the node.

  If it is appropriate to unlock the node, issue the following command: `UNLOCK NODE <node_name>`

  where *<node_name>* 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 or TLS 1.3 but the client is not at a sufficient level. Client program version 6.3 is required for TLS 1.2 and version 8.1.11 is required for TLS 1.3.
- The server is set up for TLS 1.2 or TLS 1.3 but the client imported the `cert.arm` file instead of the `cert256.arm` file.
- The server is not set up for TLS 1.2 or TLS 1.3 but the client imported the `cert256.arm` file instead of the `cert.arm` file.
- The server is set up for TLS 1.3, but the size of the key in `cert256.arm` is less than 2048 bits.

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 Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- -------------
     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 IBM Storage Protect 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 IBM Storage Protect attempts to establish a session with the server. If true, follow the guidance in the IBM Storage Protect 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.

Resolving the connection issues between a client system and the server

The connection between the IBM Storage Protect server and the clients require the use of SSL to secure the connection. The IBM Storage Protect server also communicates with the IBM Storage Protect Plus server by using the SSL encrypted connection. There are a number of reasons why the client's SSL certificates can become invalid. Though in most cases the certificates are renewed automatically, in some cases you must resolve the problems that are related to the certificates.

Before you begin

On the IBM Storage Protect server, update the value of the **SESSIONSECURITY** parameter to **TRANSITIONAL** by issuing the following command:

```
update node node_name sessionsecurity=transitional
```

Procedure

To resolve the connectivity problems, re-create the SSL certificates by completing the following steps:

1. Obtain the IBM Storage Protect server self-signed certificate: `cert256.arm`.

   Tip: Each IBM Storage Protect server stores its certificate in the server's instance directory similar to the following example: `/opt/tivoli/tsm/server/bin/cert256.arm`.
2. Copy this certificate into the IBM Storage Protect client installation directory.
3. Stop any IBM Storage Protect administrator and backup-archive clients that are running on the affected system.
4. From the IBM Storage Protect client installation directory, remove the following files, if they exist:
   - `dsmcert.crl`
   - `dsmcert.idx`
   - `dsmcert.kdb`
   - `dsmcert.rdb`
   - `dsmcert.sth`
5. From the IBM Storage Protect client installation directory, issue the **dsmcert** command to import the `cert256.arm` file with the correct server name. The command is similar to the following example:

   ```
   dsmcert -add-server SERVER1 -file cert256.arm
   ```

   where SERVER1 is the name of the IBM Storage Protect server.

   The **dsmcert** command writes new `dsmcert.*` files.
6. Issue the **dsmc q** session command to check the connectivity.

Renewing an SSL certificate of the IBM Storage Protect server

Self-signed SSL/TLS certificate that is generated by the IBM Storage Protect server expires after 10 years. You can renew the certificate when it is expired or before it expires.

Before you begin

To verify the SSL/TLS certificate expiration details, issue the following command from the server instance directory:

```
gsk8capicmd_64 -cert -details -db cert.kdb -stashed -label "TSM Server SelfSigned SHA Key"
```

The output is similar to the following example and the line starting with `Not After` shows the certificate expiration date.

```
      Label : TSM Server SelfSigned SHA Key
   Key Size : 2048
    Version : X509 V3
     Serial : aaabbbcccddd
     Issuer : "CN=TSM Self-Signed Certificate,OU=TSM Network,O=TSM,C=US"
    Subject : "CN=TSM Self-Signed Certificate,OU=TSM Network,O=TSM,C=US"
Not Before : November 15, 2012 11:16:40 AM GMT+01:00
 Not After : November 14, 2022 11:16:40 AM GMT+01:00
```

Procedure

To create and distribute a new SSL/TLS certificate (`cert256.arm`), complete the following steps:

1. Stop the IBM Storage Protect server.
2. Make a backup copy of the following certificates and key stores present in the IBM Storage Protect server instance directory.

   `cert256.arm`

   `cert.kdb`

   `cert.sth`

   `cert.rdb`

   `cert.crl`
3. Delete only the `cert256.arm` file from the server instance directory.
4. Delete the server's certificate from the key store by issuing the following command:

   ```
   gsk8capicmd_64 -cert -delete -db cert.kdb -stashed -label "TSM Server SelfSigned SHA Key"
   ```
5. Start the IBM Storage Protect server. On startup, the server generates a new certificate and stores it in the key store with label `"TSM Server SelfSigned SHA Key"`. Also, a new `cert256.arm` file is created.
6. Check the `cert.kdb` and `cert256.arm` files for more information, by issuing the following commands:

   ```
   gsk8capicmd_64 -cert -details -file cert256.arm
   ```

   ```
   gsk8capicmd_64 -cert -details -db cert.kdb -stashed -label "TSM Server SelfSigned SHA Key"
   ```

   In the output, expiration date in the line starting with `Not After` must be the same in both the files. Verify that the expiration date is 10 years in the future.
7. To provide the new certificate to other servers that communicate with the server being updated, issue the **UPDATE SERVER** command on each of the other servers by specifying the **FORCESYNC=YES** and **CERTFINGERPRINT=""** parameters. This action forces the server to update its local copy by synchronizing with the updated server's certificate.
8. To provide the new certificate to backup-archive and API clients, issue the **UPDATE NODE** command for each client node by specifying the **SESSIONSECURITY=TRANSITIONAL** parameter to change the node's SESSIONSECURITY state to TRANSITIONAL.

   Important: You must do this step for each node including the IBM Storage Protect server host itself, because the client certificate on IBM Storage Protect server host are used for doing the server database backup operation.

   a) At each node, make a backup copy of `dsmcert.kdb`, `dsmcert.idx`, and `dsmcert.sth`.

      Tip: These files are located in client installation directory. On Unix, Linux and Mac systems, if client sessions were ever started from a non-root user, copies of the certificate can be located in `$HOME/IBM/StorageProtect/certs/` or in the directory determined by the **PASSWORDDIR** client option.

   b) If a node is set to connect multiple servers, it is possible that the redistribution of the other server's certificates to that node. To avoid redistribution of certificates from other servers to that

node, you can delete the certificate for only the affected IBM Storage Protect server by completing the following steps:

i) List the certificates stored in the keystore by issuing the following command:

```
gsk8capicmd_64 -cert -list -db dsmcert.kdb -stashed
```

Note the affected server certificate name and use it in the next step.

ii) To delete the affected certificate, issue the following delete command:

```
gsk8capicmd_64 -cert -delete -db dsmcert.kdb -stashed -label certificate_name
```

If only one certificate is listed for the server that is being updated, you can simply delete all the following files: `dsmcert.kdb`, `dsmcert.sth`, and `dsmcert.idx`

Important: You must delete `dsmcert.idx` in any case, whether you delete the entire key store or you delete a single certificate from the key store.

c) Connect to the IBM Storage Protect server by using the backup-archive or API client to get the new certificate.

9. To provide the new certificate to administrative clients, issue the **UPDATE ADMIN** command for each client administrator by specifying the **SESSIONSECURITY=TRANSITIONAL** parameter to change the administrator's SESSIONSECURITY state to TRANSITIONAL.

   On each system from which an administrator uses the command line administrative client (**dsmadmc**) to connect to the server, complete the following actions:

   a. Make a backup copy of `dsmcert.kdb`, `dsmcert.sth`, and `dsmcert.idx`.

      Tip: These files are located in client installation directory. On Unix, Linux and Mac systems, if client sessions were ever started from a non-root user, copies of the certificate can be located in `$HOME/IBM/StorageProtect/certs/` or in the directory determined by the **PASSWORDDIR** client option.

   b. List the certificates stored in the keystore by issuing the following command:

      ```
      gsk8capicmd_64 -cert -list -db dsmcert.kdb -stashed
      ```

      Note the affected server certificate name and use it in the next step.

   c. If **dsmadmc** is used to connect the multiple servers from this system, issue the following command to delete only the certificate for the server whose certificate is expired.

      ```
      gsk8capicmd_64 -cert -delete -db dsmcert.kdb -stashed -label certificate_name
      ```

      If only one certificate is listed for the server that is being updated, you can simply delete all the following files: `dsmcert.kdb`, `dsmcert.sth`, and `dsmcert.idx`

      Important: You must delete `dsmcert.idx` in any case, whether you delete the entire key store or you delete a single certificate from the key store.

   d. Connect to the IBM Storage Protect server by using **dsmadmc** to get the new certificate.

10. To provide the new certificate to Operations Center, complete the following steps:

    a. Stop the Operations Center service.

    b. On the hub server, cancel any sessions between the Operations Center and the hub server.

    c. Issue the **UPDATE ADMIN** command for the following administrators by specifying the **SESSIONSECURITY=TRANSITIONAL** parameter to change the administrator's SESSIONSECURITY state to TRANSITIONAL:

       i) Any administrator who logs into the Operations Center

       ii) The Operations Center monitoring admin: `IBM-OC-`*hub_server_name*, where *hub_server_name* is the server name of the IBM Storage Protect server hub

    d. Start the Operations Center service.

e. Log in to the Operations Center with an administrator that is specified in step 10.c.i.

Automating the distribution of IBM Storage Protect server certificate to clients

Self-signed SSL/TLS certificate that is generated by the IBM Storage Protect server expires after 10 years. You can renew the certificate when it is expired or before it expires. Here an utility is provided, which helps automate the distribution of new certificate generated on IBM® Storage Protect Server to various IBM Storage Protect Clients (including BA Client, HSM, VE, various types of TDP clients etc.) This automation utility is available in the form of a script that runs on the administrative client. Several schedules, with `action=command` and are defined in this process of distribution and associated with nodes (clients) registered on the IBM Storage Protect server. These schedules are responsible for sending the certificate to the client, when the client scheduler runs those. The new certificate is then added into the client certificate database. Enabling the clients to trust the new certificate from IBM Storage Protect, for further communications, after the IBM Storage Protect switches over to the new certificate.

Before you begin

Ensure that the following prerequisites are met:

- The administrative client (`dsmadmc`) must be configured to connect with the IBM Storage Protect server.
- The client scheduler must be configured for the clients. For the details on client scheduler, refer to IBM Storage Protect scheduler overview.
- You must load and configure the utilities before the execution.

Loading Utility

The process of automation to distribute certificate to clients initiates from the administrative client (`dsmadmc`), which can be on the same machine where server is installed or on a separate machine.

To download the tar package, click here. You must extract the package on the administrative client machine.

After the extraction, you must change the directory to `cert_distribute`.

- After changing the directory, you must set directory to *windows* or *linux-aix* based on the platform of the admin client machine.
- You must use the `cert_distribute.sh` **BASH** script on Linux, AIX or other UNIX kernel based admin clients.
- You must use the `cert_distribute.ps1` Microsoft Windows powershell script on Microsoft Windows based admin client.

Configuring Utility

You must place the configurations file `cert_distribute.ini` in the same directory where the script is present. The configuration file contains parameters needed by utility to execute. You must also set appropriate values to these parameters in the configuration file.

Note: The IBM Storage Protect clients registered with **`type = OBJECTCLIENT`** or **NAS** are not covered with this utility

Table 10. System Requirements

IBM Storage Protect Components	Minimum Version
IBM Storage Protect Server	8.1.19
IBM Storage Protect Client	8.1.2

About this task

Note: This process only distributes the certificate to clients and does not set the distributed certificate to default for server and client communication.

Restriction: In the following cases, clients are not covered by this utility:

- IBM GSkit **(gsk8capicmd_64)** is not installed or located in the default location.

  You must refer to the following default location for GSkit binary on different platforms:

Platform	Default Installation Directory
AIX	`/usr/opt/ibm/gsk8_64/bin/gsk8capicmd_64`
Linux	`/usr/local/ibm/gsk8_64/bin/gsk8capicmd_64`
MAC	`/Library/ibm/gsk8/bin/gsk8capicmd`
Windows	`C:\Program Files\IBM\gsk8\bin\gsk8capicmd_64`

- Client certificate keystore file `dsmcert.kdb` path is not found in default set of locations, including client installation path.

  You must refer to the following default paths for BA client installation or binaries is:

Platform	Default BA Client Installation Directory
AIX	`/usr/tivoli/tsm/client/ba/bin`
Linux	`/opt/tivoli/tsm/client/ba/bin`
MAC	`/Library/Application\ Support/tivoli/tsm/client/ba/bin`
Windows	`C:\Program Files\Tivoli\TSM\baclient`

  If the client is not installed in above default paths, then the script looks for the client's certificate key store file under the path defined with environment variable *DSM_DIR*.

  Other than above mentioned paths, you can also find the client's certificate key store file in the path mentioned by environment variables such as PASSWORDDIR, or in `~/IBM/StorageProtect/certs` on Unix based systems and `C:\Users\user\IBM\SpectrumProtect\certs` on Windows system.

In addition to the above cases, the automated certificate renewal utility does not provide the new certificate to other IBM Storage Protect servers that communicate with the server being updated, or the administrative clients configured to connect with the server being updated, or to the Operations Center. In such cases, you must refer to the manual procedure documented under "Renewing an SSL certificate of the IBM Storage Protect server" on page 100.

Procedure

To automate the distribution of the IBM Storage Protect server certificate to clients, complete the following steps:

1. You must generate a new certificate on the IBM Storage Protect server for distributing certificate to all clients by using **CREATE CERTIFICATE** command.

   Log in to server administrative client (DSMADMC), and create the certificate by using the following command:

   ```
   CREate CERTificate "certificate_label"
   ```

   To know more details about **CREATE CERTIFICATE** command, refer to CREATE CERTIFICATE (Create a new TLS certificate).

   Important: Certificate file *certificate_label.arm* is created in the server's instance directory, where *certificate_label* is the label used while generating the certificate.

2. Copy the generated certificate file from server's instance directory to the admin client. For example, copy *certificate_label.arm* from server's instance directory to the admin client.

   Important: If server and admin client are on the same machine, you must provide the full path of the certificate file in the configuration file instead of copying the file path.
3. Place the script and configuration file in the same directory on the admin client machine.
4. Set appropriate values for different parameters in configuration file.
5. Ensure that the script has the **execute** permissions to the logged in user. You can also give permissions to the script for execution.

   For example, use the following command on Linux admin client :

   ```
   chmod +x cert_distribute.sh
   ```
6. Execute this script with dsmadmc *user_id* and *password*, along with the *action* argument as shown in the following examples:

   For Windows admin client :

   ```
   powershell .\cert_distribute.ps1 -id <dsmadmc-user-id>
                                    -pass <dsmadmc-user-password>
                                    -action <action-to-perform>
   ```

   For Linux /UNIX /Aix admin client:

   ```
   ./cert_distribute.sh -id <dsmadmc-user-id>
                        -pass <dsmadmc-user-password>
                        -action <action-to-perform>
   ```

   Note: Please provide dsmadmc user ID and password with single quotes in case they contain special characters.

   For example:

   ```
   ./cert_distribute.sh -id 'test&44admin'
                        -pass 'Colt!44lifelate'
                        -action <action-to-perform>
   ```
7. To start a certificate distribution job, run the script with action **distribute**. Set of client schedules will be defined on server to distribute certificate to all clients.

   Example of executing the utility with **distribute** action on Linux admin client:

   ```
   # ./cert_distribute.sh -id admin -password passw0rd -action distribute

   ANS8002I Highest return code was 0.

   >> Schedules are defined to distribute the certificate to clients.
   >> A return code other than 0 indicates an error. For more information, take a look at the
   log file.
   >>The log file is located at /home/cert_distribute/cert_distribute.log
   >> Client nodes will begin receiving a new certificate after 5 hours from the start time of
   the scheduler window.
   >> To monitor the progress of certificate distribution, it is recommended to run this script
   in Report mode regularly for a couple of days.

   #
   ```
8. To monitor the progress of the certificate distribution job, run the script with action **report**. The distribute schedules will be run on a daily basis. You must leave the schedules to run for certain days if some clients are not reported as **Completed**.

   Example of executing the utility with **report** action on Linux Admin Client:

```
# ./cert_distribute.sh -id admin -password passw0rd -action report

The detailed status report is : /opt/tivoli/tsm/client/ba/bin/
certdistribute.report.2023-11-02-09_49_19

>> Any status other than 'Completed' indicates a failure.
>> Client nodes with status ‘Missed’ specifies that the scheduled startup window is already
passed.
>> Client nodes with status ‘Future’ specifies that the beginning of the startup window is in
the future.
>> For ‘Missed’ and ‘Future’ status there is a chance that they will catch up as these
schedules will be re-run on a daily basis.
>> Please regularly monitor the progress of certificate distribution status, by executing
this script in Report mode for a couple of days.
>> Client nodes that are still in failed state, it is recommended to follow the manual steps
for adding the certificate.
>> 1. Copy the new certificate file over to the client box.
>> 2. Execute below command to add the certificate.
>> gsk8capicmd_64 -cert -add -label "<<New Certificate Label>>" -file "<<New Certificate File
Path>>" -db dsmcert.kdb -stashed
>> Tip: The dsmcert.kdb file is located in client installation directory. On Unix, Linux
and Mac systems, if client sessions were ever started from a non-root user, copies of the
certificate can be located in $HOME/IBM/StorageProtect/certs/ or in the directory determined
by the PASSWORDDIR client option.

#
```

9. After completing the certificate distribute job, run the script with action **cleanup** to remove the defined schedules from the IBM Storage Protect server.

   Note: You may not see some clients with **Completed** status after the retries as client scheduler might not run or there are some other issues. Such clients are considered failed. To distribute certification for the failed clients, refer to “Renewing an SSL certificate of the IBM Storage Protect server” on page 100 .

   Example of executing the utility with **cleanup** action on Linux admin client:

```
# ./cert_distribute.sh -id admin -password passw0rd -action cleanup

ANS8002I Highest return code was 0.

>> Cleanup Completed.
>> A return code other than 0 indicates an error. For more information, take a look at the
log file.
>> The log file is located at /cert_distribute/cert_distribute.log

#
```

10. You must confirm the distribution of certificate to all clients by executing script in **report** action.

    Note: The new and detailed status report is generated regularly. This report file contains the status of the schedule responsible for adding the certificate in client certificate keystore. For details, check possible values for status information at QUERY EVENT (Display client schedules)

11. Set the distributed certificate to **default** value after the report file shows the **completed** status for all clients.

    To get the new certificate in effect, the admin must execute the following command from the admin client:

```
SET DEFAULTTLSCert <certificate_label>
```

    Where, *certificate_label* is the label mentioned in configurations file while executing the script. For details on **SET DEFAULTTLSCert** command, refer to SET DEFAULTTLSCERT (Mark a TLS certificate as the default)

12. Restart the server to get the new certificate in effect for further communication with all BA and TDP clients.

    Note: This results in aborting all the ongoing sessions (like backup or restore or replication) currently running on the server.

Recovering the key database file password

If you forgot the current key database file password, IBM Storage Protect 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 certificate key database that is the `cert.kdb` file ensures that Transport Layer Security (TLS) starts when you restore the IBM Storage Protect server. If you do not have a backup copy of certificate key database, create a new certificate key database.

Procedure

To create a new certificate key database, complete the following steps:

1. Delete all `cert*` files from the server instance directory.
2. Shut down the server.
3. Start the server. The server automatically creates a new `cert.kdb` file.

   Important: The server 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.
4. Redistribute the new `.arm` file to all backup-archive clients that are using TLS. If you are using TLS 1.2 or TLS1.3, use the `cert256.arm` file. 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

Back up all the following `cert*` files from the server instance directory, which helps you to avoid losing connectivity in case certificate and password files are lost or corrupted:

cert.kdb
This file is the actual certificate key database.

cert.sth

This is the password stash file for the certificate key database. The stash file contains an obfuscated copy of the password that is required to access the contents of the certificate key database. This is the only copy of the password and if this password is lost, the existing certificate key database can no longer be used.

cert256.arm

This is an exported copy of the server's self-signed TLS certificate. Distribute this file to new clients so that they can connect to the server by using TLS.

Chapter 6. Resolving storage agent problems

IBM Storage Protect 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.

To find the latest online product information for the storage agent, see the IBM Tivoli Storage Manager for Storage Area Networks documentation: https://www.ibm.com/docs/en/tsmfsan.

The documentation for IBM Tivoli Storage Managerfor SAN version 7.1 is applicable for use with the IBM Storage Protect version 8.1 product family.

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 Storage Protect 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 Storage Protect 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 Storage Protect can, at times, experience problems that you can resolve through trace.

Windows Linux AIX Starting an extended trace of the Operations Center

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

If you contact IBM Software Support about a problem with the Operations Center, they might request that you send them log files.

Related tasks

Viewing the Operations Center log from within the Operations Center

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.

Windows Linux AIX Tracing the Operations Center by enabling logging functions from within the Operations Center

From 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

By performing the following procedure, you can 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 then click **Enable**.

   Tip: Generally, if you are not sure which logging groups to select from the list of logging groups, select the row with the `OC.LEVEL_NORMAL` logging group, and then click **Enable**.
5. Confirm that you want to enable the logging groups, and click **Close**.
6. Re-create the problem that you want to troubleshoot.

   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 then click **Disable**.
9. Confirm that you want to disable the logging groups, and then click **Close**.

What to do next

For the location and names of the Operations Center log files, see "Log files overview" on page 91.

Related tasks

Viewing the Operations Center log from within the Operations Center

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.

Tracing the Operations Center by enabling functions in the logging configuration file

If the problem that you want to troubleshoot , prevents you from opening the Operations Center, you can open and modify the logging configuration file, and then start an extended trace that adds data to the Operations Center log.

Windows Linux AIX Tracing the Operations Center by enabling functions in the logging configuration file

If the problem that you want to troubleshoot , prevents you from opening the Operations Center, you can open and modify the logging configuration file, and then 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:

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

   Windows *installation_dir*\ui\Liberty\usr\servers\guiServer

   where *installation_dir* represents the directory in which IBM Storage Protect 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.

   Tip: Generally, if you are not sure which logging groups to enable in the text editor, enable the logging group `OC.LEVEL_NORMAL` by replacing the word `OFF` with the word `ON`.
6. Save and close the file.
7. Start the Operations Center web server.
8. Re-create the problem that you want to troubleshoot.

   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 91.

Related tasks

Tracing the Operations Center by enabling logging functions from within the Operations Center
From the Operations Center, you can enable logging functions and start an extended trace that adds troubleshooting data to the Operations Center log.

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.

Procedure

Perform the following steps to enable trace classes for the server or storage agent:

1. Determine the trace classes that you want to enable. To issue trace messages for a given trace class, enable the trace class 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 the following command: **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 the following command: **TRACE BEGIN**
   - For tracing to a file with no size limitation, issue the following command: **TRACE BEGIN** *fileName*
   - For tracing to a file with a size limitation, issue the following command: **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
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.

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 11 on page 117 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 11. 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.
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.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
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.
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.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
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.
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**, **BFRTRV**, **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.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
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.
BFRTRV	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.
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.
CLOUD	This trace class displays information related to cloud input/output operations for database backup to cloud, retention cloud pools, and cloud-container storage pools with IBM Storage Protect 8.1.13 and later.	Use this trace class to diagnose cloud storage related problems.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
CLOUDJ	This trace class displays information related to cloud input/output operations in the Java GUI components. The displayed information is about all the cloud operations and other code that uses the Java GUI components. This trace class is similar to the trace class SDCLOUDJ.	Use this trace class to diagnose problems related to cloud storage and the other code that uses the Java GUI component.
COLLOCATE	This trace class displays information about collocation processing on storage pools. COLLOCATEDETAIL trace class can also be used to get more detailed information about the collocation processing. 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.
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.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
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 Storage Protect 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.
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.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
DEDUP2	Traces the fingerprinting and digital signatures path.	Use DEDUP2 to trace fingerprinting and digital signature paths.
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.
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.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
DSDEALLOC	This trace class displays information about releasing and de-allocating space on disk volumes. Typical deallocation operations on the server are **EXPIRATION**, **MIGRATION**, **MOVE DATA**, **MOVE NODEDATA**, **AUDIT VOLUME**, **DELETE VOLUME**, and **DELETE FILESPACE**.	Use this trace class to diagnose during the deletion of data.
DSRTRV	This trace class displays information about reading data from disk volumes.	Use this trace class to diagnose problems 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.
GROUP	Trace class for logical group functions.	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.
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**).

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
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**).
JVM	This trace class displays information related to the startup and basic function of the Java Virtual Machine (JVM) used by IBM Storage Protect.	Use this trace class to diagnose problems related to the startup of the JVM. Generally, this trace class should be enabled before you start the IBM Storage Protect server.
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.
MONITOR	This trace class displays information about alert monitoring.	Use this trace class to determine why an alert might not be generated.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
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.
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.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
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 the other trace classes that are 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 replication server to the target replication 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.
REPLSTATS	This trace class displays tracing related to updating the statistics as replication runs. Also includes insertion or update of history records in the replication history table.	Use this trace class to diagnose replication problems with statistical updates.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
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.
RETS	This is an aggregate trace class that enables the other trace classes that are RETSIM, RETSBF, RETSAF, RETSSD, RETSCS, RETSREPL, and RETSSC.	Use this trace class to diagnose problems with defining, creating, and scheduling retention rules and retention sets. The RETSSCD trace class provides more detailed tracing related to copying of retention set data to tape or cloud.
ROWMGR	Traces activities for row-based operations. Row-based operations are the following operations: • Abbrev • Delete • Fetch • FetchNext • FetchPrev • Insert • SearchBounds • Update	Use this trace class to trace the activities for row-based operations.
SC	This trace class displays tracing related to storage rules defined for the operations such as data retention, data replication, tiering data, and copying data.	Use this trace class to diagnose the problems related to the storage rule operations and the storage pool defined with the storage rule.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
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.
SD	This trace class displays tracing related to data processing with container storage pool, which includes both directory-container and cloud-container storage pools.	Use this trace class to diagnose the problems related to data processing with directory-container and cloud-container storage pools.
SDCLOUD	This trace class displays information related to cloud input/output operations for cloud-container storage pools.	Use this trace class to diagnose problems related to cloud-container storage pools.
SDCLOUDJ	This trace class displays information related to cloud input/output operations in the Java components related to cloud-container storage pools and other Java related operations. In general, use this trace class in conjunction with SDCLOUD.	Use this trace class to diagnose problems related to cloud-container storage pools or other code that uses the Java GUI component.
SDCRYPT	This trace class displays tracing related to encryption and decryption of data stored in directory-container and cloud-container storage pools.	Use this trace class to diagnose problems related to encrypting or decrypting data in container storage pools.
SDREPL	This trace class is used for chunk processing with any of the **PROTECT STGPOOL** operations or replication-based operations. To troubleshoot **PROTECT STGPOOL** operations, the trace class must be enabled on both the source and the target server.	Use this trace class to diagnose the problems with **PROTECT STGPOOL** of type **REPLSERVER**.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
SDREPL AS	This trace class is used for chunk processing with any of the **PROTECT STGPOOL** operations or replication-based operations. Additionally, **AS** tracing is required for **PROTECT STGPOOL** of type **LOCAL**, which helps to diagnose issues with tape-based operations.	Use this trace class to diagnose the problems with **PROTECT STGPOOL** of type **LOCAL**.
SEC	This trace class displays tracing related to the server security component, which is responsible for securely storing passwords and other sensitive information.	Use this trace class to diagnose problems with storing or fetching passwords.
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.
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 IBM Storage Protect 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.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
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.
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.
S3	This trace class displays information related to the front-end operations of an Amazon S3 API, which is also known as the S3 Object Agent that is provided by IBM Storage Protect.	Use this trace class to diagnose problems related to the spObjectAgent service and cold-data-cache storage pool.
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.

Table 11. Server or storage agent trace classes (continued)

Trace classes	Description	Usage
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.
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.

Table 11. Server or storage agent trace classes (continued)		
Trace classes	**Description**	**Usage**
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 12. 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 aggr***ID-high* aggr*ID-low*. aggr*ID-high* and aggr*ID-low* 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 ASVol**.	As sequential media volumes are assigned for use by a session or a process, they are tracked in an in-memory list. You can view this list to determine the state of in-use volumes, 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 12. Server or storage agent SHOW commands (continued)

SHOW command	Description	Recommendation
BFOBJECT	Shows the following information in the server storage hierarchy data: • The active/inactive state of logical files within an aggregate • The offset/length of logical files within an aggregate • The active state or owner bitfile ID of logical files within an aggregate • The link bitfile ID if the deduplicated extent is linked to another extent The syntax is **SHOW BFObject**.	This command helps you determine the existence and attributes of a bitfile object in the server's storage hierarchy. If you have problems restoring, retrieving, expiring, or auditing the object, you might issue this command.
DEDUPDELETEINFO	Shows the status of background deletion threads for dereferenced deduplicated objects. The syntax is **SHOW DEDUPDeleteinfo**.	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 12. 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: • 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*. *objID-high* and *objID-low* 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*. *objID-high* and *objID-low* 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 12. 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** *objID-high objID-low*. *objID-high* and *objID-low* 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: • 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** *libraryName* where *libraryName* 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** *libraryName* where *libraryName* 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 12. 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 Storage Protect. This command requires a NAS node and datamover definition only.

Table 12. 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 IBM Storage Protect 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 12. 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** *libraryName*.	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 IBM Storage Protect 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 12. 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 12. 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 IBM Storage Protect device driver

Tracing is available for the IBM Storage Protect device driver. The IBM Storage Protect 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 IBM Storage Protect device driver on all platforms where the device driver is supported. For devices that use device drivers other than the IBM Storage Protect 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
To trace the driver from the server, you must first issue the proper commands.

Tracing data from a command shell for AIX and Windows

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

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 IBM Storage Protect 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.

Windows AIX Tracing data from a command shell for AIX and 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 Storage Protect 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 Storage Protect 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.	

Trace Class Name	Description	When to use	Additional Notes
VERBINFO	Collect information regarding the client-server protocol used by IBM Storage Protect.	To debug protocol violations, transaction processing errors, or in cases where the client is stopped and not responding.	
VERBDETAIL	Detailed information regarding the client-server protocol used by IBM Storage Protect. 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 13 on page 145 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 13. 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.
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.

Table 13. Trace classes (continued)

Trace Class	Description	Recommendation
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.
AUDIT	Displays auditing information for backup and restore processing. Part of the SERVICE trace aggregate.	Use this trace class to keep record of files processed, committed, and restored in a file.
CLIENTTYPE	Displays client type on each trace output line.	Use this trace class for tracing situations when more than one client component is involved, such as the client acceptor and the file system agent.
COMPRESS	Displays compression information. Part of the SERVICE trace aggregate.	Use this trace class to determine how much data is compressed on a per-file basis.
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.

Table 13. Trace classes (continued)

Trace Class	Description	Recommendation
FIOATTRIBS	Displays comparisons of file attributes between the local client version and the active version on the server. Part of the SERVICE 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.
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 involve multiple processes or multi-threaded processes. When you use PID, also use TID.
POLICY	Displays policy information available to the backup-archive client. Part of the SERVICE and ALL_BACK trace aggregates.	Use this trace class to see which policies are available during a backup or archive operation.

Table 13. Trace classes (continued)

Trace Class	Description	Recommendation
SCHEDULER	Displays general processing information for the scheduler. An aggregate that includes most of the client trace classes 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.
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.
SM	Displays detailed processing information of backup and restore operations on file systems that are managed by the IBM Storage Protect for Space Management client. Aggregate of several trace classes related to IBM Storage Protect for Space Management when the backup-archive client is used to protect migrated files. Part of the SERVICE trace aggregate.	Use this trace class to diagnose problems that are related to backup and restore operations on a file system that is managed by the IBM Storage Protect for Space Management client.
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.

Table 13. Trace classes (continued)

Trace Class	Description	Recommendation
THREAD_STATUS	Displays thread status. Part of the SERVICE trace aggregate.	Use this trace class when you are diagnosing problems related to threading.
TID	Displays the thread ID on each trace statement. Part of the SERVICE trace aggregate.	Use this trace class to diagnose problems that involve multiple processes or multi-threaded processes. When you use TID, also use PID.
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.
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.
VM	Displays detailed information related to virtual machine's data processing. Aggregate of several trace classes related to virtual machines when the backup-archive client is used to send data to IBM Storage Protect for Virtual Environments. Part of the SERVICE trace aggregate.	Use this trace class to diagnose problems that are related to virtual machine's backup and restore operations.
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`
   Linux AIX `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

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.

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.
  - Linux AIX 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 14. Availability of the dsmtrace utility

Client Component	AIX or Linux 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 14 on page 152 includes Macintosh OS X.
- Tracing for the Data Protection components is for the IBM Storage Protect application programming interface (API) only.
- The IBM Storage Protect API tracing is available with any multithreaded application that uses the IBM Storage Protect 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
Storage Protect
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:

```
fvtlinuxppc:/opt/tivoli/tsm/client/ba/bin # dsmtrace q p

IBM
Storage Protect
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

fvtlinuxppc:/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
Storage Protect
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.

- Linux AIX 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
Storage Protect
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 dsmtrace is used only with the -TRACEFLAGS option, for example, **dsmtrace enable 2346 -traceflags=service**, then you still see the following message:

  ```
  ANS2805I Tracing has been enabled.
  ```

  In this case, the trace flags were enabled, but tracing is not active until a trace file is specified by using the -TRACEFILE option.
- Do not use the dsmtrace enable command to start tracing the application programming interface (API) for Data Protection applications if the Data Protection application is run in a manner that does not cause it to connect to the IBM Storage Protect server. For example, the Data Protection for HCL Domino command line interface has several such commands:
  - domdsmc help
  - domdsmc set
  - domdsmc query domino
  - domdsmc query pendingdbs
  - domdsmc query preferences

  If you use dsmtrace to enable tracing for such commands, the result can be a stoppage of the dsmtrace process and (AIX and Linux only) a residual named pipe in the /tmp directory.
- Windows You must be logged in with a local administrative account to use dsmtrace.
- You must be logged in as root to use dsmtrace. If a client process stops or is stopped, it might leave a named pipe (UNIX FIFO) in the /tmp directory. These FIFOs have names 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 IBM Storage Protect 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).

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: Linux AIX 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:

- Linux AIX 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 Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect API trace flags. For the list of other available trace classes, see “Client trace classes” on page 145. For example:

- `TRACEFILE /log/trace.out`
- TRACEFLAGS api api_detail pid tid verbinfo verbdetail

Important: If you do not have write permission for the file specified by the TRACEFILE option, the API calls such as `dsmSetup`, `dsmInitEx`, or `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 `dsmtrace` utility. The `dsmtrace` utility lets you turn on tracing while the problem is occurring, without having trace constantly enabled. Refer to the *dsmtrace* section.

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 Storage Protect.

The IBM Storage Protect 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 Storage Protect 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.

Linux Supporting 32-bit or 64-bit Linux kernel modules for 32-bit or 64-bit applications

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.

Linux Running an IBM Storage Protect Linux server on x86_64 architecture

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 Storage Protect Linux server and storage agent can run only on a AMD64/EM64T system with a 64-bit Linux operating system. Likewise, a 32-bit IBM Storage Protect Linux server and storage agent can run only on an AMD64/EM64T system with 32-bit Linux operating system.

A 64-bit IBM Storage Protect 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.

Linux Enabling multiple LUN support on Linux kernels

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.

Linux Using IBM Storage Protect to perform a ddtrace on 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:

```
trace enable lpdd <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 Storage Protect passthru device driver trace cannot be enabled through the ddtrace utility.

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 **CFGMGR** command to force the operating system to reconfigure itself. Then, run SMIT to reconfigure your IBM Storage Protect devices.

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 IBM Storage Protect 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 15. 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 Storage Protect 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 IBM Storage Protect 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 an IBM Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect is installed. Typically, a hard disk drive or disk subsystem is defined to the computer where IBM Storage Protect is installed as a drive or file system. After the hard disk drive or disk subsystem is defined to the operating system, IBM Storage Protect might use this space by allocating a database, recovery log, or storage pool volume on the device. The IBM Storage Protect 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 Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect 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 an IBM Storage Protect 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 Storage Protect 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 IBM Storage Protect 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 an IBM Storage Protect 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 Storage Protect 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 or Linux.

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 16. 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 Storage Protect 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 Storage Protect functions with a SAN device.

Related reference
HBA configuration issues
Host bus adapters (HBAs) have many different configuration settings and options.

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 Storage Protect 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 IBM Storage Protect SAN Device Mapping, you can complete SAN discovery on most of the systems and the persistent binding of the devices are not required. An IBM Storage Protect 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 IBM Storage Protect 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?

  IBM Storage Protect requires that an HBA is configured with tape support. IBM Storage Protect typically uses SANs for access to tape drives and libraries. As such, the HBA setting to support tapes must be enabled.

Linux AIX 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 IBM Storage Protect server or the IBM Storage Protect 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 Storage Protect 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 17 on page 174 for errors that are generated for SAN devices.

Table 17. Common SAN device errors

Error	Explanation
ANR8302E I/O error on drive *TSMDRIVE01* (*/dev/mt9*) (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	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. 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.
ANR8957E: *command*: Autodetect is OFF and the serial number reported by the library did not match the serial number in the library definition	The IBM Storage Protect SAN Device Mapping encountered a path for the library that reports a different serial number than the current IBM Storage Protect 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. Determine the new path and issue the **UPDATE PATH** command to correct this error.
ANR8958E: *command*: Autodetect is OFF and the serial number reported by the drive did not match the serial number in the drive definition	IBM Storage Protect SAN Device Mapping encountered a path for a drive that reports a different serial number than the current IBM Storage Protect 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. Determine the new path and issue the **UPDATE PATH** command to correct this error.

Table 17. Common SAN device errors (continued)

Error	Explanation
ANR8963E: Unable to find path to match the serial number that is defined for drive *driveName* in library *libraryName*	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: • Device Removed If the device was removed from the SAN, delete the server definitions that refer to this device. Issue the `QUERY PATH F=D` command to determine any paths that reference the device. Then issue the **DELETE PATH** command to remove these paths. • Device Replace A SAN Device was replaced with a new device as a result of maintenance or an upgrade. Perform the following procedures: – Try not to delete the drive or drive path definition after you replace the drive. – Issue one of the following server commands: - `UPDate DRive <libraryName> <driveName> SERIAL=AUTODetect` This command force-records the new serial number into the server database. Because the drive is replaced, the element number stays the same. - `UPDate PATH <sourceName> <driveName> SRCT=SERVER DESTT=DRIVE LIBRary=<libraryName> DEVIce=xxxxx AUTODetect=Yes` This command force-records the new serial number into the database. Because the drive is replaced, the element number stays the same. – If the drive or drive path is deleted, redefine this new, replaced drive. You must restart the IBM Storage Protect server so that the element number/serial number map for the library is refreshed. This mapping occurs only at initialization. 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: `UPDATE PATH AUTODetect=Yes`
ANR8972E: Unable to find element number for drive *driveName* in library *libraryName*	If the **ELEMent** parameter is set to AUTODetect when defining the drive, IBM Storage Protect automatically gets the drive's element number. However, if the library does not provide an element number/serial number map, this message is issued. Perform the following steps to correct this error: 1. Determine the element number for this tape drive. 2. Issue the **UPDATE DRIVE** command to update the device element number.

Linux AIX 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 IBM Storage Protect server or the IBM Storage Protect storage agent is installed.

Related concepts

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.

SAN device mapping hints and tips

SAN device discovery and device mapping are supported on Windows, AIX, and Linux.

The following items illustrate the advantages of IBM Storage Protect SAN device discovery and device mapping:

IBM Storage Protect 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.

IBM Storage Protect can update the device path automatically when a device's path changes

IBM Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect 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`

If any of these files are missing, the "ANR1791W HBAAPI wrapper library xxxxxxxxx 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, 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 HBAAPI library is not working in general:

- ANR1786W HBAAPI 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 HBAAPIn error on the SAN.

"Get HBA target mapping failed" means that the HBA encountered an error while gathering device mapping information by sending various SCSI commands.

Verify that all SAN devices are working properly (for example, a SAN Data Gateway 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 IBM Storage Protect 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 IBM Storage Protect. Verify that the wrapper file is in the same directory as the IBM Storage Protect 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`

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 or Linux 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` (QlLogic'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)

The result is that the server is not able to perform SAN discovery.

ANR1793W IBM Storage Protect SAN discovery is not supported on this platform or this version of OS

Error message ANR1793W is only displayed if the IBM Storage Protect 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:

- 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 IBM Storage Protect 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.

Windows | Linux | AIX **`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.

Linux | AIX 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.

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

AIX

Due to problems in AIX HBAAPI 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 IBM Storage Protect 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
SAN device discovery and device mapping are supported on Windows, AIX, and Linux.

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 Storage Protect devices) is /dev/rmtX for tape drives and /dev/smcX for medium changers.

Tip: Special file name: `/dev/rmt0, /dev/smc0`

For the IBM Storage Protect devices:

To update the special files, use **smitty** > **devices** > **IBM Storage Protect Devices** > **remove all defined devices**, then **discover devices supported by IBM Storage Protect**. 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. IBM Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect database. The special file name is `/dev/mtA.B.C.D` for both IBM Storage Protect devices and IBM devices, and `/dev/lbA.B.C.C` for both IBM Storage Protect 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 IBM Storage Protect devices:

Users can issue autoconf, the IBM Storage Protect device driver auto configure script. This script resides in the `/opt/tivoli/tsm/devices/bin` directory (or in the same directory as the `tsmscsi` file) to be able to configure devices and get all the current special file names and device information. The device special file name is `/dev/mtX` for tape devices and `/dev/lbX` for medium changers.

Tip: Special file name: `dev/tsmscsi/mt0, /dev/tsmscsi/lb0`

Alternatively, you can reinstall the IBM device driver. IBM Storage Protect device driver updates all the current special file names.

With the Linux pass-thru device driver for the IBM Storage Protect 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 IBM Storage Protect device driver can create configuration files (`/dev/tsmscsi/lbinfo` and `/dev/tsmscsi/mtinfo`). These files are used by the IBM Storage Protect server to create the special file names after each SAN discovery.

32 bits (Linux xSeries)

Ensure that the HBAAPI wrapper library libhbaapi32.so is in the same directory as `dsmserv` or in the `/opt/tivoli/tsm/server/bin` directory.

64 bits (Linux pSeries)

Ensure that the HBAAPI wrapper library `libhbaapi64.so` is in the same directory as `dsmserv` or in the `/opt/tivoli/tsm/server/bin` directory.

64 bits (Linux zSeries)

Ensure that the pseudo-HBAAPI wrapper library `libhbaapi64.so` is in the same directory as `dsmserv` or in the `/opt/tivoli/tsm/server/bin` directory. The wrapper library, `libhbaapi64.so`, is a link to the `/usr/lib64/libzfcphbaapi.so` file.

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 Storage Protect 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. Linux AIX
   Use the `dsmsanlist` module to obtain information about devices on a storage area network (SAN). The `dsmsanlist` module is installed by default when the IBM Storage Protect server or the IBM Storage Protect 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-IBM Storage Protect server operation hints and tips

IBM Storage Protect 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 IBM Storage Protect 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, IBM Storage Protect uses the following items:

- Up to two extra TCP/IP ports.
- A control port that is used internally by IBM Storage Protect during both backup and restore operations.
- A data port during NDMP backup operations to an IBM Storage Protect native storage pool.

The data port is an ephemeral port that is acquired at the beginning of NDMP backup operations to an IBM Storage Protect native storage pool. If a port is not available, an error message is issued and backup of NAS devices to IBM Storage Protect 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 IBM Storage Protect server because NAS restores from IBM Storage Protect native pools.

Resolving firewall issues with NDMP filer-to-IBM Storage Protect server backup and restore

A firewall might prevent the network-attached storage (NAS) file server from contacting the IBM Storage Protect 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 IBM Storage Protect server, use the NDMPPORTRANGELOW and NDMPPORTRANGEHIGH server options.

A firewall might prevent the IBM Storage Protect 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 IBM Storage Protect from accessing the NAS file server, the outbound connection from IBM Storage Protect fails.

Resolving connectivity issue with NDMP filer-to-IBM Storage Protect server

If you have more than one network adapter on the IBM Storage Protect server, the NDMP filer might connect to the server with an incorrect network adapter. If this happens, the connection fails and the server logs this information in the filer. If the NDMP filer cannot make the connection, the backup operations also fail.

To resolve the connectivity issue, use the **NDMPPREFDATAINTERFACE** option on the IBM Storage Protect server. This option specifies the IP address for the server network adapter that the filer should use for all NDMP backup data. You can update this server option without stopping and restarting the server by issuing the **SETOPT** command, similar to the following command:

```
setopt ndmpprefdatainterface ip_address
```

For more information, see the section*NDMPPREFDATAINTERFACE* section in IBM Documentation.

Resolving SCSI device problems

Tape drives and libraries might report information back to IBM Storage Protect 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 IBM Storage Protect 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 IBM Storage Protect.

Using the information reported in IBM Storage Protect message ANR8300, ANR8301, ANR8302, ANR8303, ANR8943, or ANR8944, refer to the IBM Storage Protect 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 Storage Protect 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. Collecting statistics from the servermon component for problem resolution

The servermon component automatically collects statistics about the IBM Storage Protect storage environment and archives the data. The statistics provide wide-ranging information that can facilitate problem resolution. If you cannot resolve issues in your storage environment by using statistics from the servermon component, you can provide the statistics to IBM Software Support.

About this task

Starting with IBM Storage Protect 8.1.10, the servermon component is automatically installed and enabled when the IBM Storage Protect server is installed. The servermon component collects data at regular intervals by taking snapshots of the storage environment and archiving the statistics daily.

Procedure

To share data that is collected by the servermon component with IBM Software Support, complete the following steps:

1. Obtain a list of available archive files by issuing the following command from the instance directory:

   ```
   servermon –list
   ```

   The output of this command provides the index identifier to use for extraction, the date and time that the archive was created, and the size of the compressed file. The compressed file includes data from a single day of collection and is available at the end of each 24-hour cycle of data collection.
2. Select a data archive file by issuing the following command and specifying the ID that you want to extract for analysis:

   ```
   servermon –extract –id id_to_extract
   ```

   where *id_to_extract* specifies the ID of the data archive that you identified in step 1.

   Optionally, to extract the latest data archive file, you can issue the following command without specifying the ID:

   ```
   servermon -extract
   ```
3. Go to the `srvmon` subdirectory and locate the archive `.zip` file. The file name includes the timestamp that is associated with the extracted archive ID, for example, `servermonFile-SERVER1-20210931.zip`.
4. Send the archive file to IBM Software Support.

Appendix B. Getting call stack information from a core file

You can use debugging tools or a script to get the call stack information for each running thread from a core file. The call stack provides wide-ranging information that can facilitate problem resolution.

Linux Getting call stack information on Linux

You can use the getcoreinfo script to get the call stack information for each running thread from a core file. The getcoreinfo script gets the function traceback for the failing thread and registers values and function traceback for all other threads.

When the system fails, IBM Software Support requires you to run getcoreinfo to collect call stack information. To get the call stack information, issue the following command from a Unix shell prompt:

```
/opt/tivoli/tsm/server/bin/getcoreinfo path_to_dsmserv path_to_corefile
```

where

path_to_dsmserv
Specifies the path to `dsmserv.exe`, which is a server executable file.

path_to_corefile
Specifies the path to the core file that is produced by the system failure.

This command produces the following two files as output:

- `getcoreinfo.txt` contains the call stacks of all the threads in the core file.
- `getcoreinfo-shlibs.tar.gz` is a compressed file that contains a copy of dsmserv, the core file, and all system libraries that were loaded at the time of the crash.

If you cannot resolve the problem in your storage environment by using call stack, you can provide both the `getcoreinfo.txt` and `getcoreinfo-shlibs.tar.gz` files to IBM® Software Support.

AIX Getting call stack information on AIX

You can use the dbx debugging tool to get the call stack information for the failing thread from a core dump file. The dbx tool reads the information in the core dump file and provides call stack trace information of the processing thread at the time the error occurred.

To get the call stack information from the core dump file, complete the following steps:

1. Issue a command similar to the following command:

   ```
   dbx path_to_dsmserv path_to_core
   ```

   where

   path_to_dsmserv
   Specifies the path to `dsmserv.exe`, which is a server executable file.

   path_to_core
   Specifies the path to the core dump file that is produced by the system failure.
2. Issue the `dbx where` command to display the call stack tracing for the main thread or the thread that caused the core dump.

For more information about dbx debugging tool, see dbx - Use the debugger and about `dbx where` command, see where subcommand for dbx.

If you cannot resolve the problem in your storage environment by using call stack tracing, IBM Support requires that you to run the `snapcore` command and provide the collected documentation package. The

collected package includes the server executable file, the core dump file, and all libraries loaded by the server at the time of the system failure.

To collect the documentation package for IBM Support, issue one of the following commands:

```
snapcore core_file_name program_name
```

or

```
snapcore core_file_name
```

The `snapcore` command creates a `.pax.Z` package file, and by default saves the file in the `/tmp/snapcore` directory.

Windows Getting call stack information on Windows

When the system fails, the call stack information is appended to the `dsmserv.err` file, which is in the same directory as the server. If the server is running as a service, the file is named `dsmsvc.err`. You can use the information in the error file for debugging the problem.

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 IBM Storage Protect 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
```

Attention: 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 18 on page 189 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 19 on page 195 for the key management return codes.

Table 18. 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 `open()` function call.
0x00000002	2	GSK_API_NOT_AVAILABLE	The dynamic link library (DLL) was unloaded and is not available (occurs on Microsoft Windows systems only).
0x00000003	3	GSK_INTERNAL_ERROR	Internal error. Report this error to IBM Software Support.
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 `init()` operation on a handle twice.

Table 18. 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.

Table 18. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
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.
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 gsk_secure_socket*() command that is attempted after a gsk_close_environment() 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.

Table 18. IBM Global Security Kit SSL general return codes (continued)

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

Table 18. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
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.
0x000001b6	438	GSK_ERROR_INCORRECT_SESSION_TYPE	Incorrect value is returned from the reset session type callback function. Only GSKit `gsk_sever_session`, `gsk_sever_session_with_cl_auth`, or `gsk_sever_session_with_cl_auth_crit` 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 `reset_cipher()`, and the connection uses SSLv2.

Table 18. IBM Global Security Kit SSL general return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
0x0000025a	602	GSK_MISC_INVALID_ID	A valid ID was not specified for the `gsk_secure_soc_misc()` 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 `gsk_start_trace()` must be a valid full path file name.

Table 19. 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.

Table 19. IBM Global Security Kit key management return codes (continued)

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

Table 19. IBM Global Security Kit key management return codes (continued)

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

Table 19. IBM Global Security Kit key management return codes (continued)

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

Table 19. IBM Global Security Kit key management return codes (continued)

Return code (hex)	Return code (decimal)	Constant	Explanation
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_ LABELMISMATH	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.
0x0000025a	602	GSK_MISC_INVALID_ID	An ID that is not valid was specified for the gsk_secure_soc_misc() function call.

Table 19. IBM Global Security Kit key management return codes (continued)

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

Table 19. IBM Global Security Kit key management return codes (continued)			
Return code (hex)	**Return code (decimal)**	**Constant**	**Explanation**
0x00000645	1605	GSK_TRACE_OPEN_FAILED	Trace file cannot be opened. The first parameter of gsk_start_trace() must be a valid, full-path file name.

Appendix D. Accessibility features for the IBM Storage Protect product family

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

Overview

The IBM Storage Protect family of products includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Storage Protect family of products uses the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM Documentation is enabled for accessibility.

Keyboard navigation

This product uses standard navigation keys.

Interface information

User interfaces do not have content that flashes 2 - 55 times per second.

Web user interfaces rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

Web user interfaces include WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Storage Protect 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 accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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
US

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

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

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
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.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

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

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

COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

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.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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.

Red Hat®, OpenShift®, Ansible®, and Ceph® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

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

VMware, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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 Storage Protect family of products.

See the IBM Storage Protect glossary.

Index

Special Characters

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

Product Number: 5608-E01
5608-E02
5608-E03