IBM Tivoli Storage Manager: Problem Determination Guide

Document Number: SC32-9103-02

Before getting started...
Using this guide...
Feedback

Information about...

- Help facilities
- Other sources for information
- Problem determination
- Reporting a problem
- SHOW commands
- Tracing
- Stack trace on messages
- Tivoli Storage Manager ODBC driver
- Understanding Tivoli Storage Manager messages

Hints and tips...

- Device drivers
- Hard disk drives and disk subsystems
- Storage area network (SAN)
- SAN Device Mapping
- Tape drives and libraries
- Miscellaneous server information

I have a problem with...

Client

- Diagnostic Tips
- File Include-Exclude
- Passwords and Authentication
- Scheduling
- Journal Based Backups (JBB)
- Application Programming Interface (API)
- Client Options Sets
- Windows Volume Shadowcopy Services (VSS)
- Open File Support (OFS) and the Logical Volume Snapshot Agent (LVSA)

Data Protection for

- Domino
- Enterprise Storage Server (ESS)
- Enterprise Storage Server (ESS) for mySAP.com
- Exchange
- Informix
- mySAP.com
- Oracle
- SQL

Administration Center

- Install
- Unable to access from a Web browser
- Integrated Solutions Console Server crash
- Integrated Solutions Console has excessive memory consumption
- Authority problems for an ISC User
- Unable to establish a connection with a Tivoli Storage Manager server
- Blank Pages in Administration Center
- Red Error Messages: Portlet is unavailable/Application error
- Administration Center task fails or returns unexpected results
- Administration Center task fails with a message
- Message refers to the Problem Determination Guide for more information
- Health Monitor
- Help does not display
- Problems with the Demos on the Getting Started page
- Performance Tuning Suggestions

Server

- Diagnostic Tips
- Crash
- Hang (not responding)
- Database
- Recovery Log
- Storage Pool
- Processes

Storage Agent

- Diagnostic Tips
- LAN-free setup

Data Storage

- Diagnostic Tips
- SAN devices
- SCSI devices
- Sequential media volume (tape)

Before getting started

Who should read this guide

This guide is intended for anyone administering or managing Tivoli Storage Manager. Similarly, information provided by this guide may be useful to business partners and anyone with responsibility for supporting Tivoli Storage Manager.

What you should know before reading this guide

You should be familiar with:

- Tivoli Storage Manager
- The operating systems used for the configured Tivoli Storage Manager environment

This document references error logs, trace facilities and other diagnostic information for the product. These trace facilities and diagnostic tools are not a programming interface for the product. Tivoli Storage Manager product development and support use these tools for diagnosing and debugging problems. With regard to this guide, these are provided only to aid in the diagnosing and debugging of problems. Trace facilities are subject to change without notice and may vary depending upon the version and release of the product or the platform on which the product is being run.

Information referenced within this guide may not be supported or applicable to all versions or releases of the product. Changes are periodically made to the information herein. IBM may make improvements and changes in the product(s) and the program(s) described in this publication at any time without notice.

We are very interested in hearing about your experience with Tivoli products and documentation. We also welcome your suggestions for improvements. If you have comments or suggestions about our documentation, please complete our customer feedback survey at Feedback Survey.

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make them available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that IBM product, program or service may be used. Subject to IBM's valid intellectual property or other legally protectible rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

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

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
USA

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:

Site Counsel
IBM Corporation
P.O. Box 12195
3039 Cornwallis
Research Triangle Park, NC 27709-2195
USA

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.

Trademarks

IBM, the IBM logo and the following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX

AS/400

DB2

DFS

DFSMS/MVS

DFSMShsm

DFSMSrmm

DPI

Enterprise Storage Server

ESCON

eServer

FICON

FlashCopy

HACMP

Informix

iSeries

Lotus

Lotus 1–2–3

Lotus Approach

Lotus Domino

Lotus Notes

Magstar

MVS

NetView

OpenEdition

OS/2

OS/390

OS/400

Passport Advantage

pSeries

RACF

Rational Redbooks

RS/6000

S/390

SANergy

SecureWay

StorageSmart

SystemView

Tivoli

Tivoli Enterprise Console

Tivoli Management Environment

TotalStorage

TME

VTAM

WebSphere

z/OS

zSeries

Intel, Intel Inside (logos), MMX and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

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

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

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

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

Application Programming Interface

Click on the symptom in the table or skip the table and scroll through the information below:

Application Program Interface Problem Determination

- I. How do I trace the Tivoli Storage Manager API?
- II. Where do I go to locate solutions for the Tivoli Storage Manager API?
- III. What information should I gather before calling IBM?
- IV. What files should I gather before calling IBM?
- V. What option files are used by the Tivoli Storage Manager API?
- VI. Determining if the data is encrypted and/or compressed during backup/archive
- VII. Determining if the data is being sent to the Tivoli Storage Manager Storage Agent rather than Server
- VIII. Where can I find out information about API return codes?
- IX. Running application that use the Tivoli Storage Manager API as a non-root user
- X. How do I get instrumentation information with the API?

How do I trace the Tivoli Storage Manager API?

To enable tracing for the Tivoli Storage Manager API, you need to 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. Trace flags are separated by a space.

The trace flags specific to the Tivoli Storage Manager API:

- api - information about the API function calls
- api_detail - detailed information about the API function calls

You can also specify other Tivoli Storage Manager backup/archive client and Tivoli Storage Manager API trace flags. Refer to backup/archive client documentation for a list of available trace classes.

For example:

- TRACEFILE /log/trace.out
- TRACEFLAGS api api_detail verbinfo verbdetail timestamp

Note that if the user does not have write permission for the file pointed by the TRACFILE option, dsmSetup or dsmInitEx/dsmInit will fail with return code DSM_RC_CANNOT_OPEN_TRACEFILE (426).

Also, to enable tracing for the multi-threaded Tivoli Storage Manager API after an application has been started you can use the dsmtrace utility. It provides the ability to turn on tracing while the problem is occurring without having trace enabled all the time. Refer to the Dsmtrace section of the *Problem Determination Guide*.

Return to Steps to Diagnose Client Application Programming Interface (API) Problems

Where do I go to locate solutions for the Tivoli Storage Manager API?

A number of resources are available to learn about or to diagnose the Tivoli Storage Manager API. The IBM Support Solutions database contains a knowledge base of articles and information on issues related to backup-restore problems. Refer to Support for

IBM Tivoli Storage Manager to review this information. Click the "Hints and Tips, Solutions, and Support Flashes" links in the self-help table to access the search tool. Enter a term in order to search through solutions to previously encountered issues.

Return to Steps to Diagnose Client Application Programming Interface (API) Problems

What information should I gather before calling IBM?

Collecting all the necessary information about the environment can significantly assist in determining the problem. Gather as much of the following information as possible before contacting IBM Support. This information will assist IBM Support with resolving your problem.

- What operating system is the problem being experienced on.
- Exact level of the operating system, including all service packs and hot fixes that have been applied.
- Exact level of the Tivoli Storage Manager API.
- Exact level of the Tivoli Storage Manager server.
- Tivoli Storage Manager Server platform and operating system level.
- Exact level of the Tivoli Storage Manager storage agent (if LAN-free environment).
- Tivoli Storage Manager Storage Agent platform and operating system level (if LAN-free environment).
- Output from the Tivoli Storage Manager server query system command.
- List of other applications running on the system.
- List of the steps needed to re-create the problem (if the problem is re-creatable).
- If the problem is not re-creatable, list the steps that caused the problem.

Return to Steps to Diagnose Client Application Programming Interface (API) Problems

What files should I gather before calling IBM?

A number of log files and other data may be created by the Tivoli Storage Manager API. Gather as many of the following files as possible before contacting IBM Support. The contents of these files will assist IBM Support with resolving your problem.

- Tivoli Storage Manager API options file. For Windows(R), the default options file dsm.opt or the file referenced by the DSMI_CONFIG environment variable. For UNIX(R), the default options file is dsm.sys in the directory referenced by the DSMI_DIR environment variable.
- Tivoli Storage Manager API error log file. The default API error log is dsierror.log.
- Any trace files created for the API (the recommended trace flags are api api_detail verbdetail)
- Output from failed command or operation. This may be either the console output redirected to a file or an actual screen image of the failure.
- Tivoli Storage Manager server activity log. A Tivoli Storage Manager administrator can view this log for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- If the API client is configured for LAN-free data movement, also collect the options file for the Tivoli Storage Manager storage agent. The default name for this file is dsmsta.opt.
- A short program or sections of the application source code invoking the Tivoli Storage Manager API function calls that are suspected to cause the problem.

Return to Steps to Diagnose Client Application Programming Interface (API) Problems

What option files are used by the Tivoli Storage Manager API?

On UNIX and OS/400 operating systems, the options reside in two options files:

- The client options file (dsm.opt).
- The system options file (dsm.sys).

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

DSMI_CONFIG - The fully-qualified name for the client options file.

DSMI_DIR - This points to the API installation directory and also is used to find the dsm.sys file (on Unix). Wherever the DSMI_DIR is set to, ensure that a dsm.sys file exists in this directory.

DSMI_LOG - Points to the path for the dsierror.log file.
Note that if this variable points to a directory for which the user does not have write permission, dsmSetup and dsmInitEx will fail with return code DSM_RC_ACCESS_DENIED (106).

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

You can use the following technique to verify that the API uses the correct option file and/or server stanza in dsm.sys.

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

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

SERVERNAME srvr1.cmpron
COMPRESSION YES
ERRORNEOUS_OPTION 12345
TCPSERVERADDRESS machine.company.com

SERVERNAME srvr1.pwdfl
PASSWORDACCESS GENERATE
PASSWORDDIR .
TCPSERVERADDRESS machine.company.com
...
```

2. Verify that the API detects this error. The sample API program dapismp can be used 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.
```

3. If no error is reported, you have updated the wrong options file. Check the environment variable values previously mentioned or repeat Steps 1 and 2 with another options file/server stanza.

4. Remove the option inserted in Step 1.
Return to Steps to Diagnose Client Application Programming Interface (API) Problems

Determining if the data is encrypted and/or compressed during backup/archive

To verify that the data is encrypted and/or compressed during backup/archive through trace, follow these steps:

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

TRACEFILE <trace file name>
TRACEFLAGS api api_detail

2. Examine the trace file after the operation.

3. Locate a statement that looks similar to this one.
dsmSendObj ENTRY:... objNameP: '<the file name>'

4. The line in Step 3 is followed by this trace message that indicates whether the object is compressed, encrypted, compressed and encrypted.

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

```
   +-------------------------------------------------------------------------------
+
   | encryptType/compress |             0                 |
1              |
   +-------------------------------------------------------------------------------
+
   | NO                   | not compressed, not encrypted | compressed, not
encrypted |
   | CLIENTENCRKEY        | not compressed, encrypted     | compressed,
encrypted     |
   | USER                 | not compressed, encrypted     | compressed,
encrypted     |
   +-------------------------------------------------------------------------------
+
```

Alternatively, your application itself can determine encryption type/strength and compression of your data by using the dsmEndSendObjEx function call and the dsmEndSendObjExOut_t data structure.

```
/*-----------------------------------------------------------------------+
|  Type definition for dsmEndSendObjExOut_t
+-----------------------------------------------------------------------*/
typedef struct dsmEndSendObjExOut_t
{
   dsUint16_t         stVersion;           /* structure version */
   dsStruct64_t       totalBytesSent;      /* total bytes read from app */
   dsmBool_t          objCompressed;       /* was object compressed */
   dsStruct64_t       totalCompressSize;   /* total size after compress */
   dsStruct64_t       totalLFBytesSent;    /* total bytes sent LAN Free */
   dsUint8_t          encryptionType;      /* type of encryption used   */
}dsmEndSendObjExOut_t;

objCompressed  - A flag that displays if the object was compressed.
encryptionType - A flag that displays the encryption type.

For example:

   ...
   rc = dsmEndSendObjEx(&endSendObjExIn, &endSendObjExOut);
   if (rc)
   {
```

```
            printf("*** dsmEndSendObjEx failed: ");
            rcApiOut(dsmHandle, rc);
        }
        else
        {
            printf("Compression:          %s\n",
                endSendObjExOut.objCompressed == bTrue ? "YES" : "NO");

            printf("Encryption:           %s\n",
                   endSendObjExOut.encryptionType & DSM_ENCRYPT_CLIENTENCRKEY ?
     "CLIENTENCRKEY" :
                   endSendObjExOut.encryptionType & DSM_ENCRYPT_USER ? "USER" : "NO");
            printf("Encryption Strength:  %s\n\n",
                   endSendObjExOut.encryptionType & DSM_ENCRYPT_AES_128BIT ?
     "AES_128BIT" :
                   endSendObjExOut.encryptionType & DSM_ENCRYPT_DES_56BIT ? "DES_56BIT" :
     "NONE");
        }
        ...
```

See Chapter 6. API Function Calls in Using the Application Programming Interface for details.

Return to Steps to Diagnose Client Application Programming Interface (API) Problems

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

To verify that the data is being sent to the Tivoli Storage Manager storage agent rather than server, follow these steps:

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

   TRACEFILE <trace file name>
   TRACEFLAGS api api_detail verbdetail

2. Examine the trace file after the operation.

3. Locate a statement that looks similar to this one.

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

4. The line in Step 3 is followed by this trace statement that indicates whether the object totalLFBytesSent was sent to the Tivoli Storage Manager storage agent. If totalLFBytesSent is 0 0, the data was sent directly to the Tivoli Storage Manager server.

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

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

```
/*-------------------------------------------------------------------------+
|  Type definition for dsmEndSendObjExOut_t
+-------------------------------------------------------------------------*/
typedef struct dsmEndSendObjExOut_t
{
   dsUint16_t         stVersion;           /* structure version */
   dsStruct64_t       totalBytesSent;      /* total bytes read from app */
```

```
   dsmBool_t          objCompressed;       /* was object compressed */
   dsStruct64_t       totalCompressSize;   /* total size after compress */
   dsStruct64_t       totalLFBytesSent;    /* total bytes sent LAN Free */
   dsUint8_t          encryptionType;      /* type of encryption used   */
}dsmEndSendObjExOut_t;

totalLFBytesSent  - The total LAN-free bytes that were sent.

For example:

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

See Chapter 6. API Function Calls in Using the Application Programming Interface for details.

Return to Steps to Diagnose Client Application Programming Interface (API) Problems

Where can I find out information about API return codes?

The complete list of return codes along with their explanations is located in Appendix D in *Using the Application Programming Interface*. The return codes are also listed in the dsmrc.h include file.

Return to Steps to Diagnose Client Application Programming Interface (API) Problems

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

Here is a check-list of things that need to be done in order for a non-root user to run an application that uses the API successfully on Unix:

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

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

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

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

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

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

4. If PASSWORDACCESS GENERATE is set in system options file dsm.sys, perform the following steps.

4.1 Check ownership and permissions of the Trusted Communication Agent (TCA) that should reside in the directory pointed by DSMI_DIR environment variable. For example, the following TCA has the correct ownership and permissions:

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

Wrong permissions or ownership will result in DSM_RC_AUTH_FAILURE (137) returned from dsmInit. Additionally, it is imperative to use the same version of the API library and dsmtca. Mixing versions will result in an error.

```
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.
```

For details about the TCA, see Chapter 3 in Using the Application Programming Interface. 4.2 Generate password file TSM.PWD using either the Tivoli Storage Manager Backup/Archive client or sample API application dapismp. Location of the password file is determined by PASSWORDDIR option in dsm.sys. This step must be performed by the root user. In the following example, the sample API application is used to generate TSM.PWD for a node whose password is oddesy.

```
# dapismp

************************************************************************
* Welcome to the sample application for the Tivoli Storage Manager API. *
* API Library Version = 5.3.0.0                                        *
************************************************************************

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.
```

5. If tracing is enabled, verify that the non-root user has write permission for the file pointed by the TRACEFILE option.

Return to Steps to Diagnose Client Application Programming Interface (API) Problems

How do I get instrumentation information with the API?

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

Return to Steps to Diagnose Client Application Programming Interface (API) Problems

Main Menu | Client diagnostic tips

Client diagnostic tips

For a client problem, review these steps to try to isolate or resolve the problem:

Diagnosing a client problem...

1. My dsmc/dsmadmc/dsmj/... won't start.
2. Examine any error messages that were issued.
3. Examine the server activity log messages for this session.
4. Is this an error connecting (communicating) to the server?
5. Were client options changed?
6. Were policy settings changed on the server?
7. Is the client being run with the QUIET option?
8. Verify INCLUDE/EXCLUDE syntax and ordering.
9. Was this the correct Tivoli Storage Manager server?
10. Identify when and where the problem can occur.
11. If the problem can be reproduced, try to minimize the circumstances under which it can occur.
12. Documentation to collect.

My dsmc/dsmadmc/dsmj/... won't start.

This message is displayed and the process stops:

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

Beginning with Tivoli Storage Manager Version 5.3, client applications will not run without a writable log. The system has denied write access to the log file named in the message. The following rules apply. (The dsmerror.log is used as an example, but all logs are treated equally.)

If the log does not exist, it will be created with default permissions.

- The name and the directory specified by the ERRORLOGNAME option will be used.
- If the option is absent, the name dsmerror.log in the directory specified in the DSM_LOG environment variable, if present, will be used.
- Otherwise, the name dsmerror.log in the current working directory will be used.

Because default permissions are used, a log created by the root user may not be writable by any other user. If this is the case, the root user must set the proper permissions or ACLs to allow free use of the server application by all users who need to use it.

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

The client does not attempt to create logs in the root directory. ANS1398E is displayed when the method in rule 1 directs the log file to be created in the root directory.

If a log file exists and can be located using the method in rule 1, it will be used. It can also be in the root directory if you so chose. Furthermore, whatever permissions you give that log file will be preserved by Tivoli Storage Manager code.

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

Additional info:
Certain background Tivoli Storage Manager applications may fail to start due to errors writing to dsmerror.log. When these errors occur, a number of errors are recorded in the Windows system event log ('system log' on other platforms). One of these errors indicates that the dsmerror.log cannot be written to.

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

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

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

Return to diagnosing a client problem...

Examine any error messages that were issued.

Check for ANSnnnnx messages issued to the console, dsmsched.log, or dsmerror.log.

Additional information for ANSnnnnx messages is available in either the *Tivoli Storage Manager Messages* or from the client HELP facility.

Return to diagnosing a client problem...

Examine the server activity log messages for this session.

Check for server activity log using QUERY ACTLOG for messages issued for this client session. The messages from the server activity log may provide additional information about the symptoms for the problem or may provide information about the actual cause of the problem the client encountered.

Additional information for ANRnnnnx messages in the server activity log is available in either the *Tivoli Storage Manager Messages* or from the server HELP facility.

Return to diagnosing a client problem...

Is this an error connecting (communicating) to the server?

If the client is unable to connect to the server, it is likely either a configuration error with the client network options or a network problem.

For problems connecting to the server, check the following:

- Whether or not the client communication options in the client option file has been changed. If so, review these changes and try reverting back to the previous values and retrying the connection.
- Whether or not the server communication settings have been changed. If so, either update the client communication options to reflect the changed server values or else revert the server back to its original values.
- Whether or not any network settings have been changed. For example, has the TCP/IP address for the client or server been changed? If there have been network changes, work with the network administrator to understand these changes and update the client, server, or both for these network changes.

Return to diagnosing a client problem...

Were client options changed?

Changes to client options are not recognized by the client scheduler until the scheduler is stopped and restarted.

Stop and restart the client scheduler.

Return to diagnosing a client problem...

Is the client being run with the QUIET option?

The QUIET processing option for the client suppresses messages.

Restart the client without the QUIET option. This will allow all the messages to be issued and will assist with a more complete understanding of the problem.

Return to diagnosing a client problem...

Verify INCLUDE/EXCLUDE syntax and ordering.

The include/exclude processing can impact which files are sent to the server for a backup or archive operation.

Refer to the Client file include-exclude section of this document for more information on client include/exclude syntax and ordering.

Return to diagnosing a client problem...

Was this the correct Tivoli Storage Manager server?

In an environment with multiple Tivoli Storage Manager servers, a client may use different servers for different operations.

If multiple servers are available for this client to use, ensure that the client connected to the appropriate server for the operation that was attempted.

Return to diagnosing a client problem...

Identify when and where the problem can occur.

Client processing problems often only occur when performing specific operations, at certain times, or only on certain client machines.

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

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

Return to diagnosing a client problem...

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

By minimizing the complexity of the environment for recreating a problem, Tivoli Storage Manager support will be better able to assist in diagnosing or recreating the problem.

Consider the following options to minimize the complexity of the environment for recreating a problem:

- Use a minimal options file consisting of only TCPSERVERADDRESS 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 running the command manually.

Return to diagnosing a client problem...

Documentation to collect.

The Tivoli Storage Manager client provides information in a number of different sources. These should be inspected and if they have information relating to this, they should be provided to support.

Documentation for client problems and configuration information may be found in one or more of the following:

- Error log. The client error log file is dsmerror.log.
- Scheduler log. The error log for the client scheduler is dsmsched.log.
- Web client log. The error log for the web client is dsmwebcl.log.
- Options files. The client may use a combination of files for its configuration. These files are dsm.opt, dsm.sys for UNIX systems, and the include/exclude file.
- Trace data. If tracing was active, the file containing the trace data should be provided to support.
- Application dump. If the client crashed, many platforms will generate an application dump. The application dump is provided by the operating system.
- Memory dump. If the client hung, a memory dump can be generated that can be used to help with diagnosis. The ability to create a memory dump varies by system and is provided by the operating system.
- List all the software installed on the client system. The client may experience problems due to interactions with other software on the machine or because of the maintenance levels of software that the client uses.
- Starting with Tivoli Storage Manager Version 5.2, the command *dsmc query systemInfo* is available and will collect most of this information in the file *dsminfo.txt*.
- Client option sets defined on the server that apply to this client node.
- Server options. There are a number of server options that are used to manage the interaction between the client and server. An example of one such server option is TXNGROUPMAX.
- Information about this node as it is defined to the server. This can be collected by issuing QUERY NODE *nodeName* F=D using an administrative client connected to the server.
- Schedule definitions for the schedules that apply to this node. These can be queried from the server using the QUERY SCHEDULE command.
- The policy information configured for this node on the Tivoli Storage Manager server. This information can be queried from the server using QUERY DOMAIN, QUERY POLICYSET, QUERY MANAGEMENTCLASS, and QUERY COPYGROUP.

Return to diagnosing a client problem...

Main Menu | File include-exclude

File include-exclude

Click on the symptom in the table or skip the table and scroll through the information below:

List of file include-exclude problems
• A file is not being included during backup processing • A file is not being excluded during backup processing

A file is not being included during backup processing

If you have implicitly or explicitly indicated that a file should be included during backup processing, but it is being excluded or not being processed, there are several things to check:

- Some files are automatically excluded from backup processing
- Windows(R) system files are automatically excluded from backup of the system drive
- EXCLUDE.DIR will exclude all files in the parent directory
- The server client options set has an exclude statement that matches the file name
- The server has a policy that dictates a certain number of days must pass between backups
- Include statements for compression, encryption and subfile backup do not imply inclusion for backup
- Volume delimiters and directory delimiters are not specified correctly
- The include/exclude list is coded incorrectly

Return to list of file include-exclude problems

A file is not being excluded during backup processing

If you have implicitly or explicitly indicated that a file should be excluded during backup processing, but it is still being processed, there are several things to check:

- Windows system files are automatically backed up as part of the SYSTEMOBJECT or SYSTEMSTATE domain
- EXCLUDE.DIR statements should not be terminated with a directory delimiter
- Selective backup of a single file does not honor EXCLUDE.DIR
- The server client options set has an include statement that matches the file name
- Exclude statements for compression, encryption and subfile backup do not imply exclusion from backup
- Volume delimiters and directory delimiters are not specified correctly
- The include/exclude list is coded incorrectly

Return to list of file include-exclude problems

Some files are automatically excluded from backup processing

There are some files that should not be backed up by the backup application. Reasons could include files that have been identified by the operating system as not necessary for backup and files that Tivoli Storage Manager uses for internal processing.

If there is a need to have these files included in the backup processing, Tivoli Storage Manager can include these files by coding include statements in the client options set on the server. Because these files have been explicitly identified as files not being backed up, including them in the server client options set is not recommended.

You can issue the Backup-Archive client command **DSMC QUERY INCLEXCL** to identify these files; the output from this command will show "Operating System" as the source file for files that have been automatically excluded from backup processing:

```
tsm> q inclexcl
*** FILE INCLUDE/EXCLUDE ***
Mode Function  Pattern (match from top down)  Source File
---- --------- ------------------------------ -----------------

Excl All       C:\WINDOWS\Registration\*.clb  Operating System
Excl All       C:\WINDOWS\netlogon.chg        Operating System
```

The following files are automatically excluded on various backup-archive client platforms. This list is current for the backup-archive clients.

Windows

1. Files enumerated in the registry key HKLM\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup
2. The client staging directory C:\ADSM.SYS
3. RSM database files (these files are processed in the system object or system state backup)
4. IIS metafiles (these files are processed in the system object or system state backup)
5. Registry files (these files are processed in the system object or system state backup)
6. Client trace file

UNIX

1. Client trace file

NetWare

1. Client trace file

Macintosh

1. Volatile, temporary, and device files used by the operating system
2. Client trace file

Return to list of file include-exclude problems

Windows system files

Windows system files are silently excluded from the system drive backup processing.

These files cannot be included in the system drive backup processing. To process these files the user must issue a **DSMC BACKUP SYSTEMOBJECT** (Windows 2000 and Windows XP) or a **DSMC BACKUP SYSTEMSTATE** (Windows 2003) command.

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 protected by Windows system file protection (sfp). These files will not be processed during backup of the system drive, for example, during an incremental backup of the C: drive. These files are excluded from the system drive processing internally instead of relying on explicit exclude statements (due to the sheer number of exclude statements that would be needed to represent these files which would adversely affect backup performance).

You can issue the Backup-Archive client command **DSMC QUERY SYSTEMINFO** to identify these files. The output of this command is written to the file dsminfo.txt

```
(partial contents of file dsmfino.txt)
=====================================================================
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)
```

Return to list of file include-exclude problems

EXCLUDE.DIR

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

If the intent is to include all files based on a file pattern, regardless of their location within a directory structure, the EXCLUDE.DIR statements should not be used.

For example, consider this set of UNIX include/exclude statements:

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

This Include statement in this example indicates that all files with a ".o" extension should be included, but the preceding **exclude.dir** statement will exclude all files in the /usr directory, even if they have a ".o" extension. This would be true regardless of the order of these two statements.

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

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

When using wildcards in INCLUDE/EXCLUDE, use * if you want all the files rather than *.*. *.* means all files containing at least one dot (.) character, while * means all files. If you use *.*, files containing no dot characters (such as C:\MYDIR\MYFILE) will not be filtered. Restart the client without the QUIET option. This will allow all the messages to be issued and will assist with a more complete understanding of the problem.

Return to list of file include-exclude problems

EXCLUDE.DIR syntax

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

These are examples of invalid EXCLUDE.DIR statements due to a terminating directory delimiter:

```
exclude.dir /usr/             (UNIX)
exclude.dir c:\directory\     (Windows)
exclude.dir SYS:\PUBLIC\      (NetWare)
exclude.dir Panther:User:     (Macintosh)
```

These examples show the correct coding of EXCLUDE.DIR:

```
exclude.dir /usr              (UNIX)
exclude.dir c:\directory      (Windows)
exclude.dir SYS:\PUBLIC       (NetWare)
exclude.dir Panther:User      (Macintosh)
```

Return to list of file include-exclude problems

Selective backup of a single file does not honor EXCLUDE.DIR

A selective backup of a single file from the command-line client does not honor the EXCLUDE.DIR option.

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

For example, consider the UNIX include/exclude statement:

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

If the user issues a selective backup against a single file in this directory, the file will be processed:

```
dsmc selective /home/spike/my.file  (this file will be processed)
```

If the user issues a selective backup with a wildcard, no files will be processed because the directory is excluded:

```
dsmc selective "/home/spike/my.*"   (nothing will be processed because /home/spike is
excluded)
```

Note that a subsequent incremental backup of the /home file system will inactivate the file "/home/spike/my.file".

Return to list of file include-exclude problems

Include and Exclude statements in the Tivoli Storage Manager server client option set

The Tivoli Storage Manager administrator has the ability to include or exclude files on behalf of the client. Include or exclude statements that come from the server will override include and exclude statements coded in the local client's option file.

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

You can issue the Backup-Archive client command **DSMC QUERY INCLEXCL** to identify files that are included or excluded by the server client options set. The output from this command will show "Operating System" as the source file for files that have been automatically excluded from backup processing:

```
tsm> q inclexcl
*** FILE INCLUDE/EXCLUDE ***
Mode Function  Pattern (match from top down)  Source File
---- --------- ------------------------------ -----------------

Excl All       /.../*.o                        Server
Incl All       /.../*.o                        dsm.sys
```

In this example, the users indicated that they wanted all files that end with a ".o" extension to be included in the local options file, but the server sent the client an option to exclude all files that end with a ".o" extension. The server-provided option prevails.

Return to list of file include-exclude problems

Tivoli Storage Manager server policy dictates an incremental copy frequency that is non-zero

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

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

- Contact the Tivoli Storage Manager server administrator to change the copy frequency attribute

- Execute a selective backup of the file, for example, **DSMC SELECTIVE C:\FILE.TXT**

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

```
tsm: WINBETA>q copygroup standard active f=d

                  Policy Domain Name: STANDARD
                          ...
                      Copy Frequency: 1
                          ...
```

Return to list of file include-exclude problems

Include and exclude statements for compression, encryption, subfile backup do not imply the files will be included for backup processing

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

Use include and exclude statements in combination with these other statements to produce desired results.

Consider the UNIX example:

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

This statement indicates that the file /usr/file.o will be excluded from backup processing. The **include.compress** statement indicates that "if a file is a candidate for backup processing and matches the pattern /usr/*.o; then compress the file". The **include.compress** statement should not be interpreted as "backup all files that match the pattern /usr/*.o and compress them". If the user wants to back up file /usr/file.o in this example the exclude statement must be removed.

Return to list of file include-exclude problems

Include and exclude syntax for "everything" and "all files under a specific directory" is platform specific

If the volume delimiters or directory delimiters are not correct, it may cause **include** and **exclude** statements not to function properly.

If you want to code an **include** statement for "all files under a specific directory", ensure that the slashes are correct and volume delimiters are correct. For example, suppose you want to exclude all of the files under a directory called "home", or simply all files:

- Windows uses the backwards slash "\" and the volume delimiter ":"

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

- UNIX uses the forward slash "/"

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

- NetWare can use either slash will work but the volume name is required with the volume delimiter ":"

```
*include everything in the SYS:\HOME directory
include sys:\home\...\*  or
include sys:/home/.../*
*include everything
include *:\...\*         or
include *:/.../*
```

- Macintosh uses the ":" as the directory and volume delimiter

```
*include everything in the Users directory
include Panther:Users:...:*
*include everything
include *:...:*
```

Return to list of file include-exclude problems

The include/exclude list is coded incorrectly

Due to the complexity or number of include/exclude statements, there might be the unintentional inclusion or exclusion of a file.

A new trace statement was added to the Tivoli Storage Manager backup-archive client Version 5.2.2 that can help determine why a file was included or excluded. This can be done by configuring the client with the **INCLEXCL** traceflag.

For example, the user believes that file "c:\home\file.txt" 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'
```

Further investigation using the backup-archive client command **DSMC QUERY INCLEXCL** shows that this statement is coming from the Tivoli Storage Manager server client options set:

```
tsm> q inclexcl
*** FILE INCLUDE/EXCLUDE ***
Mode Function  Pattern (match from top down)  Source File
---- --------- ------------------------------ -----------------

Excl All       c:\home\*.txt                  Server
```

Return to list of file include-exclude problems

Main Menu | File Journal Based Backup

Journal Based Backups (JBB) Problem Determination

Click on the symptom in the table or skip the table and scroll through the information below:

Journal Based Backup (JBB) Problem Determination
• Determining if a backup will be journal based • Client and Journal Daemon Traceflags • Running the journal daemon in the foreground • Journal Database Viewing Utility

Determining if a backup will be journal based.

The following conditions must be met in order for a backup to be journal based:

- The journal daemon must be configured to journal the file system being backed up.
- The journal for the file system being backed up must be in the valid state.
- The Tivoli Storage Manager node and server the backup is using must match the node and server for which the journal is valid.

The journal daemon is configured to journalize a file system by adding the file systems to the list of journalized file systems in the journal daemon configuration file tsmjbbd.ini as follows:

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

In order for a journal to be valid, a full incremental backup must be performed on the corresponding file system while the file system is actively being journalized. This full incremental backup must set the "Last Backup Completed" date on the Tivoli Storage Manager server file space in order for the journal to be set to valid. The "Last Backup Completed" date may be viewed by issuing the QUERY FILESPACE server command. After the journal is set to the valid state, subsequent backups by the same node to the same Tivoli Storage Manager server will be journal-based. If a backup uses a different node and/or 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.

The Journal Database Viewing Utility may also be used to determine the current state of a journal. If a valid journal is restarted backups will be non-journal based until the journal is revalidated. 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.

Specific messages detailing the specific cause of the journal restart are also written to both the eventlog and the journal errorlog (jbberror.log).

Reasons for restarting a valid journal are:

1. 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.)
2. Request by a backup client
   - Clients will issue a journal restart request when it is determined that a journal file system lacks integrity for one of the following reasons:
     - the server filespace no longer exists
     - the server filespace was deleted after the last backup
     - the node policy set was updated after the last backup
     - the Last Backup Completed or Last Backup Started dates aren't valid (not set).

Return to Journal Based Backups (JBB) Problem Determination

Client and Journal Daemon Traceflags

Trace Settings for the Journal Daemon

Trace settings for the journal daemon are specified in the journal daemon 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

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
FILEOPS - internal database activity
ALL_JBB - aggregate traceflag that includes all of the above settings

Trace Settings for the B/A client specified in dsm.opt:

JOURNAL - journal based backup tracing

Return to Journal Based Backups (JBB) Problem Determination

Running the journal daemon in the foreground

For diagnostic and testing purposes, it is sometimes useful to run the journal daemon in the foreground as opposed to running it as a Windows service.

The journal daemon may be started from a Windows command prompt as follows:

tsmjbbd.exe i

Return to Journal Based Backups (JBB) Problem Determination

Journal Database Viewing Utility

The Journal Database Viewing Utility provides the following information:

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

This utility also allows searching, inserting, or deleting specific entries in a journal database.
The syntax of this utility is :

dbviewb <fully qualified journal database basefile name> dbviewb <fully qualified journal database basefile name> <i>

```
Example 1:
D:\tsm530c\debug\bin\winnt_unicode>dbviewb c:\tsmjournal\tsmh__.jdb
```

```
Tivoli Storage Manager
Journal Database Viewing Utility
Version 5, Release 3, Level 0.0
Last Update: Oct 11 2004


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 Mon Oct 11 11:49:05 2004
     Journal Validation Date Tue Oct 12 16:41:11 2004
     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:\tsm530c\debug\bin\winnt_unicode>


Example 2:

D:\tsm530c\debug\bin\winnt_unicode>dbviewb c:\tsmjournal\tsmh__.jdb i
Tivoli Storage Manager
Journal Database Viewing Utility
Version 5, Release 3, Level 0.0
Last Update: Oct 11 2004


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 Mon Oct 11 11:49:05 2004
     Journal Validation Date Tue Oct 12 16:41:11 2004
     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 may 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     : Thu Oct 14 10:15:28 2004
     Object Time  : Thu Oct 14 14:15:28 2004
     Hit Count    : -2110169276
  ----------------------------------------

  Please enter your request: quit
```

Return to Journal Based Backups (JBB) Problem Determination

Main Menu | File Volume Shadowcopy Services

Open File Support (OFS) and the Logical Volume Snapshot Agent (LVSA) Problem Determination

Click on the symptom in the table or skip the table and scroll through the information below:

Open File Support (OFS) and the Logical Volume Snapshot Agent (LVSA) Problem Determination
• Examine the Windows system event log • Obtaining debug print output • Configuring the system for a full dump • Forcing a dump for a system hang when a LVSA problem is suspected

Examine the Windows system event log

In Tivoli Storage Manager 5.3.0, support was added to the tsmlvsa.sys driver to write critical information for problem determination to the Windows system event log. Examining the event log is the first step in isolating potential problems with LVSA in the context of online image or open file backup.

WOpen File Support (OFS) and the Logical Volume Snapshot Agent (LVSA) Problem Determination

Obtaining debug print output

Many problems require trace data to be obtained from the tsmlvsa.sys driver to supplement what is obtainable from a client service trace. The LVSA debug print output can be obtained by using the DebugView tool.

1. Download the latest version of the DebugView tool. Make sure you use the version that is labeled "you plan on using DebugView on WinNT/2K/XP".
2. Install the DebugView tool. This is simply extracting the files from the dbgvnt.zip file.
3. Place the following line in the dsm.opt file. Please ensure that you direct the trace file to a location with many GB of free space.

TRACEFLAG SERVICE
TRACEFILE trace.txt

4. Run the executable dbgview.exe prior to running the Tivoli Storage Manager client.
5. Configure dbgview.exe to log the output to a file via the File->Log to File option.
6. Perform the failing operation.

Open File Support (OFS) and the Logical Volume Snapshot Agent (LVSA) Problem Determination

Configuring the system for a full dump

When a system bug check occurs, it is necessary to obtain a full memory dump to assist in the diagnosis of possible LVSA problems. The failing system must be configured to take a full memory dump. The configuration steps are:

Setup system to take a full memory dump [if not already setup] (changing these settings requires reboot):
1. open the control panel
2. open the system icon
3. select the advanced tab
4. select the startup and recovery... button
5. In the "System Failure" section:
   a. Insure the following checkboxes are selected:
      i. Write an event to the system log
      ii. Send an administrative alert
   b. Insure the following checkbox is not selected:
      i. Automatically reboot
6. In the "write debugging information" section
   a. Select "Complete Memory Dump"
      i. make a note of where the file will be written (%SystemRoot%\MEMORY.DMP)
      ii. Insure the checkbox "Overwrite any existing file" checkbox is selected
7. Windows will ask to reboot so that new setting can take effect. Please reboot at this time
8. After reboot if the "Overwrite any existing file" was NOT previously checked then rename the previous dump file (if any) to a new name.
9. When a bug check occurs please note the contents of the bug check screen (commonly referred to as the Blue Screen Of Death or BSOD). Collect the memory.dmp file upon reboot for examination.

Open File Support (OFS) and the Logical Volume Snapshot Agent (LVSA) Problem Determination

Forcing a dump for a system hang when a LVSA problem is suspected

If you have ensured proper configuration and a dump is not taken, a dump may need to be forced when the system hangs. There are two methods to employ.

1. If you have the opportunity to reboot and recreate the hang then see Microsoft Knowledge Base article 244138 "Windows feature allows a Memory.dmp file to be generated with the keyboard". This method requires a registry change and reboot to enable an on-demand dump when the right CTRL key is held and SCROLL LOCK key is pressed twice. This method may also be required if the BANG! Tool mentioned below is unable to cause a bug check and memory dump.

2. If the system is hung and you cannot afford a reboot and recreate:

a. Download and install BANG! from www.osronline.com. Follow instructions provided in the BANG! package/website.

b. Run BANG! and press the "Crash Now" button.
c. The system should get a blue screen and generate a full memory dump.

NOTE: IBM does not support the BANG! utility; any questions or problems regarding the BANG! utility should be reported to osronline.

Open File Support (OFS) and the Logical Volume Snapshot Agent (LVSA) Problem Determination Main Menu | Client Option Sets

Guide to Client Option Sets

Click on the symptom in the table or skip the table and scroll through the information below:

Guide to Client Option Sets
• I. What are Tivoli Storage Manager Client Option Sets? • II. Where can I find detailed information on client option sets? • III. Why must I use INCLEXCL option instead of INCLUDE or EXCLUDE? • IV. What order will the options be processed? • V. Tivoli Storage Manager client and client option sets Traceflags. • VI. Examples of when to use client options sets.

What are Tivoli Storage Manager Client Option Sets?

A Tivoli Storage Manager administrator can create a set of client options to be used by a client node that is on Tivoli Storage Manager Version 3.0 or later. These client options are defined on the Tivoli Storage Manager server. The client options specified in the client option set are used in conjunction with the client options file. Client option sets allow the administrator to specify additional options that may not be included in the client's option file. The client can use these defined options during a backup, archive, restore, or retrieve process.

Return to Guide to Client Option Sets

Where can I find detailed information on Client Option Sets?

General overview: *IBM Tivoli Storage Manager Administrator's Guide* > "Managing Client Nodes" > "Managing Client Option Files"

Tivoli Storage Manager server commands: *IBM Tivoli Storage Manager Administrator's Reference* > "Administrative Commands"

DEFINE CLOPTSET
DELETE CLOPTSET
COPY CLOPTSET
QUERY CLOPTSET
UPDATE CLOPTSET
REGISTER NODE
QUERY NODE
UPDATE NODE
DEFINE CLIENTOPT
DELETE CLIENTOPT
UPDATE CLIENTOPT

Tivoli Storage Manager client options: *IBM Tivoli Storage Manager Backup-Archive Client Installation and User's Guide*.

NOTE: A list of Tivoli Storage Manager client options which are supported can be found in the *IBM Tivoli Storage Manager*

Administrator's Reference > the DEFINE CLIENTOPT command.

Technical Support Article #1178114 http://www-306.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManagerExtendedEdition.html

Return to Guide to Client Option Sets

Why must I use INCLEXCL option instead of INCLUDE or EXCLUDE?

Options passed to the Tivoli Storage Manager client from the Tivoli Storage Manager server are provided in groups. This means if INCLUDE and EXCLUDE options were supported on the Tivoli Storage Manager server that all INCLUDE options would be sent in a group and all EXCLUDE options would be sent in a group. You could not intermix these options to get desired results of including some files from excluded directories. By using the INCLEXCL option this allows you to intermix and order the include and exclude options.

Return to Guide to Client Option Sets

What order will the options be processed?

Multiple options can be defined and assigned a sequence number. These options are then processed from low to high sequence. As an example if you defined the following INCLEXCL options:

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

The above INCLEXCL statement would result in all files on the Novell NetWare SYS: volume being excluded, directories would still be backed up. With exception of the files in the following path 'SYS:\SYSTEM*' and 'SYS:\TMP*' which will be backed up.

Return to Guide to Client Option Sets

Tivoli Storage Manager client and client option sets Traceflags.

Trace settings for the client option sets are specified in the Tivoli Storage Manager option file for all clients. The exception is MAC which uses the Tivoli Storage Manager User Preference file.

TRACEFILE <trace file name>
TRACEFLAGS config

Return to Guide to Client Option Sets

Examples of when to use Client Options Sets

Following are customer scenarios which can take advantage of the client option set and examples of what commands would be required.

a. The customer has multiple NetWare servers and wants to control specific options for backup.
   i. Wants SYS: volume to be backed up.
   ii. Wants the Novell NetWare server which has the root level NDS to be backed up.
   iii. Wants all Tivoli Storage Manager NetWare clients to use standard display on client commands with 'scrollines and scrollprompt'. This will control the display so things won't scroll off the screen.
   iv. Wants to exclude all license objects from NDS backup. License objects can not be backed up or restored by Tivoli Storage

Manager. License must be reinstalled via Novell NetWare.

Note: Domain and inclexcl options can not use the force parameter. As a result, the client can override the client option set via the Tivoli Storage Manager option file.

Tivoli Storage Manager server command:

```
Define cloptset netware description="NetWare Server Option Sets"
Define clientopt netware domain sys:
Define clientopt netware domain nds:
Define clientopt netware scrollprompt yes
Define clientopt netware scrollline 20
Define clientopt netware inclexcl " exclude 'NDS:*License ID*' "
Update node mock cloptset=netware
```

b. The customer has a critical environment and restore is high priority.

i. Wishes to use collocatebyfilespec so all filespec data is stored on as few tapes as possible. This will enhance restore processing because fewer tape mounts will be needed.

ii. Does not want the client to be able to override this option.

Tivoli Storage Manager server command:

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

c. The customer has machines on a slow network and limited space on the Tivoli Storage Manager server for data.

i. Use compression option to limit the amount of data sent and stored.

Tivoli Storage Manager server command:

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

d. The customer has a 24/7 database which cannot be stopped. Since the files are open, the server cannot back them up. The customer wants to exclude all files and subdirectories from Tivoli Storage Manager backups and add this to the existing client option set "space_rest".

i. 'exclude.dir', specify the directory path to be excluded.

Tivoli Storage Manager server command:

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

e. The customer has a fast network and wants to make the best possible use of Tivoli Storage Manager client resources to complete backups.

i. Wants 'resourceutilization' set to maximum.

Tivoli Storage Manager server command:

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

Return to Guide to Client Option Sets

Main Menu | Client passwords and authentication

Client passwords and authentication

Click the symptom in the table or skip the table and scroll through the information below:

List of client password and authentication problems
• ANS1025E Session rejected: Authentication failure • ANS1874E Login denied to NetWare Target Service Agent 'server-name' • ANS2025E Login failed to NetWare file server 'server-name'

ANS1025E Session rejected: Authentication failure

Generally, this error occurs when the password has expired. It can also occur, however, if either the server or the client has been renamed.

If you receive **ANS1025E** during an interactive session, it is probable that the password given is incorrect.

- Have the Tivoli Storage Manager administrator reset the node's password by using the **UPDATE NODE** command.
- Issue a **DSMC QUERY SESSION** command and when prompted give the new password.

If you are receiving **ANS1025E** during a noninteractive session such as central scheduling, ensure that the client option is PASSWORDACCESS GENERATE. This option causes the client to store the password locally. The password is encrypted and stored either in the registry for Windows(R) clients or in a file named TSM.PWD for Macintosh, UNIX(R) and NetWare clients. Editing the registry or the TSM.PWD file is not recommended. Instead, follow these recommendations:

- Make sure that **PASSWORDACCESS GENERATE** is set in the option file.
- Issue a **DSMC QUERY SESSION** command. This command will force the locally stored password to be set.
- If this does not resolve the problem, update the node's password using the **UPDATE NODE** administrative command.
- Reissue the **DSMC QUERY SESSION** command, providing the new password.

To see the password expiration setting for a particular node, issue the administrative command **QUERY NODE F=D**. Look for the *Password Expiration Period* field. Note: If this field is blank, the default password expiration period of 90 days is in effect.

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

Return to list of client password and authentication problems

ANS1874E Login denied to NetWare Target Service Agent 'server-name'

ANS2025E Login failed to NetWare file server 'server-name'

These messages indicate an authentication problem between the Tivoli Storage Manager NetWare client and the NetWare server.

The most probable cause of this error is either not ensuring that the proper Target Service Agent is loaded, or not using the Novell distinguished name.

The password for the NetWare Target Service Agent is encrypted and stored into the same file that is used to store the Tivoli Storage Manager password. The process of storing the password in this file is not related to the client option **PASSWORDACCESS GENERATE**. If you do not want the password stored locally, code the option **NWPWFILE NO** into the dsm.opt file.

Some things to verify:

- **LOAD TSA500/TSA600** (depending on NetWare OS version).

- For NDS backups, **LOAD TSANDS.**
- Use the Novell typeful name. For example, instead of *Admin* use *.CN=Admin.O=IBM.*
- Ensure that **NWPWFILE YES** (the default setting) in dsm.opt.
- Check if Tivoli Storage Manager can connect to the file system Target Service Agent (TSA) by issuing the command **DSMC QUERY TSA.**
- Check if the server can connect to the NDS Target Service Agent (TSA) by issuing the command **DSMC QUERY TSA NDS.**
- **Note:** The **DSMC QUERY TSA** command can be used to store the password in the local password file and also to test if the stored password is valid.

Other things to check for when experiencing NetWare login failures include:

- The NetWare user-id has been disabled.
- The NetWare user-id/password is invalid or expired.
- The NetWare user-id has inadequate security access.
- The NetWare user-id has insufficient rights to files and directories.
- The NetWare user-id specified has a login restriction based on time-of-day.
- The NetWare user-id specified has a Network address restriction.
- The NetWare user-id specified has a login restriction based on number of concurrent connections.
- NetWare is not allowing logins (**DISABLE LOGIN** was issued at the console).
- The temporary files in "SYS:\SYSTEM\TSA" are corrupt which can prevent logins. Shut down all Tivoli Storage Manager processes and the SMS modules, and then either move or delete these temporary files.
  The following Novell TID discuss this issue: error: FFFDFFD7 when the tape software tries to login in order to backup nds
- The SMDR configuration has become corrupt. Reset the SMDR with the following command: 'LOAD SMDR /NEW' at the NetWare console.
- If the message is displayed intermittently during a multiple session backup or restore, the probable cause is that there are not enough Novell licenses available. Each Tivoli Storage Manager session requires at least one licensed connection to the file server. Either add more NetWare licenses, or reduce the **RESOURCEUTILIZATION** setting. The NetWare utility **Nwadmn32** can be used to determine the current number of licenses.
- Check the SYS volume for free space. A lack of free space can cause the authentication failure.
- If the above items have checked out, then please contact Novell for additional support to resolve this issue.

Return to list of client password and authentication problems

Main Menu | Client scheduling

Client scheduling

Click the symptom in the table or skip the table and scroll through the information below:

List of client scheduler problems
• Troubleshooting the Tivoli Storage Manager client scheduler • Using QUERY EVENT to query scheduled events • Checking the server activity log • Inspecting the client schedule log • Restarting the scheduler process on a remote machine

Troubleshooting the client scheduler

If you are experiencing problems with a client scheduler, there are a number of diagnostic steps to help determine the cause of the

problem:

- Use the **QUERY EVENT** command to determine the status of a scheduled event
- If a scheduled event is missed but other consecutive scheduled events for that node show a result of *Completed*, then check for errors in the server activity log and the client schedule log for more information.
- Use the diagnostic tool SHOW PENDING to display schedules, nodes, and when they should execute next.
- Additional helpful information when diagnosing why a node missed a scheduled event includes the dsm.sys stanza for the node and the MANAGEDSERVICES, PRESCHEDCMD, and POSTSCHEDCMD option values from the client options file.

Return to list of client scheduler problems

Using QUERY EVENT to query scheduled events

The server maintains records of scheduled events which can be helpful when managing Tivoli Storage Manager schedules on numerous client machines. The **query event** command allows an administrator to view the event records on the server. A useful query that shows all of the event results for the previous day is:

```
query event * *  begind=today-1 begint=00:00:00 endd=today-1 endt=23:59:59
```

Or, the query results can be limited to exception cases:

```
query event * *  begind=today-1 begint=00:00:00 endd=today-1 endt=23:59:59
exceptionsonly=yes
```

The query results include a status field that gives a summary of the result for a specific event. By using the **format=detailed** option you can also see the result of an event that is the overall return code passed back by the Tivoli Storage Manager client. The following table summarizes the meaning of the event status codes that are likely to exist for a scheduled event that has already taken place:

```
Status      Meaning
Completed   The scheduled client event ran to completion without a critical failure.
            There is a possibility that the event completed with some errors or warnings.
            Query the event with detailed format to inspect the event result for more
information.
            The result can either be 0, 4, or 8.
Missed      The schedule start window has elapsed without action from the TSM client.
            Common explanations for this result include the schedule service not running
on the client,
            or a previously scheduled event not completing for the same or a different
schedule.
Started     Typically this indicates that a scheduled event has begun processing.
            However, if an event showing a status of Started is followed by one or more
Missed events,
            it is possible that the client scheduler encountered a hang while processing
that event.
            One common cause for a hanging client schedule is the occurrence of a user
interaction prompt
            such as a prompt for an encryption key that is not responded to.
Failed      The client event ran to completion; however, a critical failure occurred.
```

Return to list of client scheduler problems

Checking the server activity log

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

For example, if investigating the following suspect event:

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

You could use 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
```

Return to list of client scheduler problems

Inspecting the client schedule log

The Tivoli Storage Manager client keeps a detailed log of all scheduled activities. If queries of the server's activity log are not able to explain a failed scheduled event, the next place to check is the client's local schedule log.

Access to the client machine is required for inspecting the schedule log. The schedule log file is typically stored in the same directory as the dsmerror.log in a file named dsmsched.log. The location of the log file can be specified using client options, so you may need to refer to the options file to see if the **schedlogname** option has been used to relocate the log file. On Windows(R), the schedule log can also be relocated by an option setting which is part of the schedule service definition. The **dsmcutil query** command can be used to check if this option has been set. When you have located the schedule log, it is easy to search through the file to find the time period corresponding with the start date and time of the scheduled event in question. Here are some tips on what to look for:

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

Return to list of client scheduler problems

Restarting the scheduler process on a remote machine

When managing a large number of clients running scheduler processes, it is helpful to be able to start and stop the client service from a remote machine.

The Tivoli Storage Manager client for Windows provides a utility to assist with remote management of the scheduler service. For other platforms, standard operating system utilities are required.

Windows
In order to remotely manage the client scheduler service using **dsmcutil** with the **/machine:** option, you must have administrative rights in the domain of the target machine. To determine whether the scheduler service is running on a remote machine, check the "Current Status" field from a query similar to:

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

To restart a scheduler service that is missing schedules:

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

Or, if you are using the CAD to manage the scheduler, you may have to restart the CAD service, or stop the scheduler service and restart the CAD service:

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

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

UNIX

A shell script can be written to search for and stop running Tivoli Storage Manager schedulers or client acceptor processes, and then restart the processes. Software products such as Symark's Power Broker can be used to allow Tivoli Storage Manager administrators limited access to UNIX(R) servers for the purpose of managing the scheduler processes, and copying off the schedule log file.

The following shell script is an example of how to recycle the Tivoli Storage Manager scheduler process:

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

# If necessary, the following variables can be customized to allow an
# alternate options file to be used.
# export DSM_DIR=
# export DSM_CONFIG=
# export PATH=$PATH:$DSM_DIR

# Extract the PID for the running TSM Scheduler
PID=$(ps -ef | grep "dsmc sched" | grep -v "grep" | awk {'print $2'});
print "Original TSM scheduler process using PID=$PID"

# Kill the scheduler
kill -9 $PID

# Restart the scheduler with nohup, redirecting all output to NULL
# Output will still be logged in the dsmsched.log
nohup dsmc sched 2>&1 > /dev/null &

# Extract the PID for the running TSM Scheduler
PID=$(ps -ef | grep "dsmc sched" | grep -v "grep" | awk {'print $2'});
print "New TSM scheduler process using PID=$PID"
```

Return to list of client scheduler problems

Main Menu | Show Commands | Show commands for the backup/archive client

Show commands for the backup/archive client

To find information about a show command, click the command in the table below or scroll through the list of commands below the table. The information about show commands is in the form:

COMMAND

Description of command

Recommendation for when to use

Other information

List of show commands

- CACHE
- CLUSTER
- DOMAIN
- OPTIONS
- OPTTABLE
- PLUGINS
- SESSION
- SYSTEMOBJECT
- SYSTEMSERVICES
- SYSTEMSTATE
- TRACEFLAGS
- VERSION

CACHE

Displays information about the subfile cache.

For Microsoft(R) Windows(R) clients configured to use subfile backup, this is useful for displaying information about the configured subfile cache.

None.

Return to list of show commands

CLUSTER

Displays information about the disk mappings in a Microsoft Cluster.

Useful for displaying information about the disk mapping (configuration) in a Microsoft Cluster environment.

None.

Return to list of show commands

DOMAIN

Displays information about the configured domains to use for incremental backup processing.

Useful for displaying and summarizing the `DOMAIN, DOMAIN.IMAGE, and DOMAIN.NAS` client options.

None.

Return to list of show commands

OPTIONS

Displays the client options.

Useful to determine the settings of client options.

None.

Return to list of show commands

OPTTABLE

Displays information about which options are administered by the server versus those that are managed by the client option file.

The client may 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 have been defined using `DEFINECLOPTSET`. This command is useful in determining whether the client is using an option configured from the option file or an option configured from a client option set defined on the server.

None.

Return to list of show commands

PLUGINS

Displays information about installed plug-ins for this client.

The client uses plug-ins to provide additional capabilities such as image backup. This show command displays which plug-ins are installed for this client as well as attributes of the various plug-ins such as their version, type, and location.

None.

Return to list of show commands

SESSION

Displays capabilities that this client has for this connection to the server.

The client and server report and negotiate the capabilities that each has when a session is started by a client to a server. This show command reports the capabilities available by this server and client.

None.

Return to list of show commands

SYSTEMOBJECT

For Windows 2000 and XP clients, displays the SYSTEM OBJECT data that is available on this client.

Useful to determine which SYSTEM OBJECT files are installed on this Windows client and those that could be backed up.

None.

Return to list of show commands

SYSTEMSERVICES

For Windows 2003 clients, displays the SYSTEM SERVICES data that is available on this client.

Useful to determine which SYSTEM SERVICES files are installed on this Windows client and those that could be backed up.

None.

Return to list of show commands

SYSTEMSTATE

For Windows 2003 clients, displays the SYSTEM STATE data that is available on this client.

Useful to determine which SYSTEM STATE files are installed on this Windows machine and those that could be backed up.

None.

Return to list of show commands

TRACEFLAGS

Displays information about trace classes and aggregate trace classes for this client.

Useful to determine which trace classes and aggregate trace classes could be used for this client.

None.

Return to list of show commands

VERSION

Displays the version and build date for this client.

Useful to determine which client is being run and when it was built.

None.

Return to list of show commands Main Menu | File Volume Shadowcopy Services

Windows Volume Shadowcopy Services (VSS) Problem Determination

Click on the symptom in the table or skip the table and scroll through the information below:

Windows Volume Shadowcopy Services (VSS) Problem Determination
• VSS and Tivoli Storage Manager • Transient Errors • Test Flags • Windows 2003 hot fixes • Microsoft VSS tuning recommendations • Diagnostic information for Microsoft assistance • Using the vsreq.exe sample program • Comparing Tivoli Storage Manager and Ntbackup.exe interaction with VSS

VSS and Tivoli Storage Manager

The Tivoli Storage Manager Windows client uses the Volume Shadowcopy Services (VSS) of Windows 2003 to perform system state and system services backup. There are a number of improvements to the Tivoli Storage Manager client and to the Windows 2003 operating system to improve reliability of backups performed with VSS.

Windows Volume Shadowcopy Services (VSS) Problem Determination

Transient Errors

The Tivoli Storage Manager client considers the following VSS errors to be transient. When one of these errors occurs, the client will, by default, retry the VSS backup process three times at thirty second intervals. The number of retries and retry interval are configurable via two test flags.

VSS_E_MAXIMUM_NUMBER_OF_VOLUMES_REACHED
VSS_E_SNAPSHOT_SET_IN_PROGRESS
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 Volume Shadowcopy Services (VSS) Problem Determination

Test Flags

SETVSSMAXRETRY - specifies the number of times the VSS backup process will be 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 to wait sixty seconds.

Option file example: retry 10 times at 300 second intervals

TESTFLAG SETVSSMAXRETRY:10
TESTFLAG SETVSSDELAY:300

Windows Volume Shadowcopy Services (VSS) Problem Determination

Windows 2003 Hot Fixes

Microsoft recommends the installation of hotfixes (QFEs) 833167 and 887827 for Windows 2003 VSS problems. Customers must contact Microsoft directly and request these fixes.

Windows Volume Shadowcopy Services (VSS) Problem Determination

Microsoft VSS Tuning Recommendations

Controlling the VSS diff area size:

If after applying these fixes, one of the following events occurs:

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

Consider reducing the I/O load on this system to avoid this problem in the future. If these events still occur, then the following registry key can be used 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, you can increase it up to 3000).

Recommended event log maximum size: Microsoft indicates 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. Microsoft recommends that logs are kept below 64MB in size.

Windows Volume Shadowcopy Services (VSS) Problem Determination

Diagnostic Information for Microsoft Assistance

When it is determined that the VSS failure is outside the scope of Tivoli Storage Manager, then the following information should be gathered for Microsoft support to examine:
1. Windows application event log
2. Windows system event log
3. VSS trace (see instructions below)

Event logs:
Examine the application and system event logs focusing on the error events created by the VolSnap and VSS sources at the time of failure. It is helpful to extract the germane events from the log to isolate the problem and have a more productive interaction with MS support.

How to perform a VSS trace:
1. Create a file tracefile.reg using the contents shown below and change the TraceFile entry to point to a volume that is not going to be shadow copied. Note the double backslash delimiter usage-you need to enter "\\" as the delimiter for each backslash in the path you wish to specify.
2. Install tracing.reg.
3. Reproduce the problem.
4. Turn off tracing by deleting the key "HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\VSS\Debug\Tracing"

Here are the contents of 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 Volume Shadowcopy Services (VSS) Problem Determination

Using the vsreq.exe Sample Program

The VSS SDK contains a sample program vsreq (VSS requestor) which performs a sequence of VSS API calls very similar to the server's. It is useful to compile and run vsreq.exe on the failing system to determine if vsreq and Tivoli Storage Manager encounter the same problem. If vsreq can reproduce the same problem as Tivoli Storage Manager then the output of vsreq can be supplied to MS support to help in diagnosing the VSS problem.

I/O performance diagnosis:
In some cases Microsoft has provided an I/O subsystem analysis tool "yapt" to gather I/O performance data for analysis.

Windows Volume Shadowcopy Services (VSS) Problem Determination

Comparing Tivoli Storage Manager and Ntbackup.exe interaction with VSS

Ntbackup.exe does not fully exploit VSS and cannot always be considered as a benchmark for Tivoli Storage Manager interaction with VSS. This is a list of known differences between Ntbackup.exe and Tivoli Storage Manager in the context of VSS:

- Ntbackup.exe does not use VSS to backup the Active Directory (NTDS). Although Ntbackup.exe uses VSS to take a snapshot, it still uses the legacy NTDS backup API to read data from the disk. Tivoli Storage Manager uses the VSS interface to read NTDS data from the disk. If there is a problem with the VSS writer responsible for NTDS, it will not evidence itself with Ntbackup.exe. Use the Microsoft utility vssadmin.exe to ensure that the NTDS writer is available on the system if a problem is suspected.

Windows Volume Shadowcopy Services (VSS) Problem Determination

Main Menu | Unable to establish a connection with a Tivoli Storage Manager server

Unable to establish a connection with a Tivoli Storage Manager server

For a problem establishing a connection to a Tivoli Storage Manager server, using the Administration Center, review these steps to try to isolate or resolve the problem:

Steps to establish a connection between the Administration Center and a Tivoli Storage Manager server

1. Verify the machine is accessible over the network
2. Verify the Tivoli Storage Manager server running
3. Try using the administrative client to connect to the server
4. Verify the Tivoli Storage Manager port is correct
5. Check firewall settings
6. Verify the TCP/IP address is correct for the server connection
7. Verify the TCP/IP port is correct for the server connection
8. Verify the administrative id is correct for the server connection
9. Verify the administrative password is correct for the server connection
10. Check the activity log of the server
11. Check the server's database file (tsmservers.xml)

Verify that the machine is accessible over the network.

Check to see if the machine is accessible over the network. Use the ping command in a command prompt or shell. Or, try connecting using telnet or ftp.

If the machine is not accessible over the network, then this indicates the machine is down or there is a network communication error. Use the same TCP/IP address entered in the Administration Center.

Return to steps list

Verify that the Tivoli Storage Manager server is running.

Check to see if the Tivoli Storage Manager server is running.

If the server is not running, then the Administration Center won't be able to establish a connection to it.

Return to steps list

Try using the administrative client to connect to the server.

Try connecting to the Tivoli Storage Manager server using the administrative client. Use the same connection settings as in the Administration Center (TCP/IP address, TCP/IP port, admin id, admin password).

This is the quickest way to determine if the problem is on the Tivoli Storage Manager server machine or on the machine running the Administration Center. If both the Tivoli Storage Manager server and the Administration Center are running on the same machine, then skip this step. Go to the directory where the administrative client is installed. Then, issue `dsmadmc` with parameters matching those in the connection settings for the Administration Center.

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

Return to steps list

Verify the Tivoli Storage Manager port is correct.

Check to see what port the Tivoli Storage Manager server is using. You can issue the `QUERY OPT TCPPORT` command on the server to determine the server port.

Make sure the port used in the Administration Center is the same port the server is running on. Go to the directory where the administrative client is installed. Then issue `dsmadmc` with parameters matching those in the connection settings for the Administration Center. This command can also be issued from the server console.

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

tsm: SERVER1>QUERY OPT TCPPORT

Server Option         Option Setting
-----------------     --------------------
TCPPort               1500

tsm: SERVER1>
```

Return to steps list

Check firewall settings.

Check to see if the machine you are attempting to access is behind a firewall.

If the machine running the Tivoli Storage Manager server is behind a firewall, then the machine running the Administration Center must be authenticated to the firewall. Try using the machine running the Administration Center to connect to the machine running the Tivoli Storage Manager server. Use ping, telnet or ftp. If this doesn't work, try connecting to another machine under the same firewall as the server machine.

Return to steps list

Verify the TCP/IP address correct for the server connection.

Check to see if the TCP/IP address for the Tivoli Storage Manager server is entered correctly in the Administration Center.

Use the **Modify Server Connection** table action to view the current connection settings for the Tivoli Storage Manager server connection. Verify that the server connection address is correct. Use the ping command in a command prompt or shell. Verify that the machine indicated by the address is accessible.

Return to steps list

Verify the TCP/IP port is correct for the server connection.

Check to see if the TCP/IP port for the Tivoli Storage Manager server is entered correctly in the Administration Center.

Use the **Modify Server Connection** table action to view the current connection settings for the Tivoli Storage Manager server connection. Verify the server connection port is correct. Use the `QUERY OPT TCPPORT` command in a server console to determine which port the Tivoli Storage Manager server is running on. To check which port the Tivoli Storage Manager server is using, go to the directory where the administrative client is installed. Then issue `dsmadmc` with parameters matching those in the connection settings for the Administration Center. This command can also be issued from the server console.

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

tsm: SERVER1>QUERY OPT TCPPORT

Server Option         Option Setting
-----------------     --------------------
TCPPort               1500

tsm: SERVER1>
```

Return to steps list

Verify the administrative id is correct for the server connection.

Check to see if the administrator id for the Tivoli Storage Manager server is entered correctly in the Administration Center.

Use the **Modify Server Connection** table action to view the current connection settings for the Tivoli Storage Manager server connection. Verify that the administrative id is correct.

Return to steps list

Verify the administrative password is correct for the server connection.

Check to see if the administrator password for the Tivoli Storage Manager server is entered correctly in the Administration Center.

Use the **Modify Server Connection** table action to view the current connection settings for the Tivoli Storage Manager server connection. Enter the correct administrative password to connect to the Tivoli Storage Manager server. Retry the operation to see if changing the password solved the problem.

Return to steps list

Check the activity log of the server.

Check to see if any messages are issued on the server when trying to connect using the Administration Center.

The activity log will have useful information if there is trouble connecting to the server. You will be able to see sessions starting and stopping when the Administration Center is able to communicate with the server successfully. If the network is down or the TCP/IP address or port is incorrect, then no information will show up in the activity log when connection attempts are made. To check the server activity log, go to the directory where the administrative client is installed. Then issue `dsmadmc` with parameters matching those in the connection settings for the Administration Center. This command can also be issued from the server console.

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

tsm: SERVER1>QUERY ACTLOG
```

Return to steps list

Check the server's database file (tsmservers.xml)

Check the server's database file for corruption. The server's database file is located in the ISC install directory under: `PortalServer\shared\app\dsm\tsmservers.xml`.

If this file is corrupted it may cause problems when connecting to servers. You will, however, typically see other error messages indicating this file is corrupt. Open this file in an editor (preferably an XML editor). Verify that the end tag has a closing end tag.

Return to steps list

Problem with Tutorials

Topics

Verify the following:

1. The help system is running.
2. Pop-ups are allowed for the site.
3. MacroMedia Flash player is working.
4. Tutorials are installed correctly.

Verified all of the above but the tutorials are still not working

The help system is running.

The tutorials are part of the help system. When the help system is not running, you will not be able to view the tutorials. The first step is to verify the help server process is running.

On Windows:
Open the Task Manager and select the process tab.

On UNIX:
Run a command like *ps -ef | grep eclipse*.

Look for either eclipse.exe or EclipseSrv.exe in the list of processes. When you do not see either process name in the list, the help system is not running. You will need to start the help system.

When there are multiple processes with the same names in the list, or continued attempts to start the help server fail, further investigation is necessary. This situation can arise if you are using the publications CD on the same system, or some other application is using the port. If by chance the random port number the publications CD chooses is 8423, then the Integrated Solutions Console help server will not start.

Try pointing your Web browser to http://<machine name>:8423. For example, http://wscheid.tucson.ibm.com:8423 would show the help system on wscheid.tucson.ibm.com. You should see the title, Integrated Solutions Console, on the top left hand corner of the screen. If you do not see that title, then some other help system or program is running on port 8423. Stop that application and start the Integrated Solutions Console help server again. See Starting and Stopping the help server for more information.

Pop-ups are allowed for the site.

After you verify that the help server is running, verify that pop-ups are allowed for the site.

Every modern web browser now contains some pop-up blocker. Third-party additions can also perform this function. The Administration Center uses JavaScript to open the tutorial windows. Many pop-up blockers consider this action to be something that should be blocked. Consult the help for the pop-up blocker software to determine how to allow a Web site to display pop-ups. For example, Mozilla will place a little icon (looks like an exclamation point running from an explosion) in the status bar indicating that the pop-up is blocked. To allow the site to open pop-ups, you click on the icon and press the add button in the dialog. Internet Explorer has a similar feature, although the browser version and the operating system version can change how to configure the setting or whether there is a setting.

Another way to test whether a pop-up blocker is causing a problem is to run the tutorials directly in the help system. Open a browser window to http://<machine name>:8423. Expand the IBM Tivoli Storage Manager branch of the contents. Expand the Tutorials branch of the contents. Select a tutorial to view. In the work area to the right of the Contents tree the tutorial should start running. If the tutorial starts using this method, then the pop-up blocker is probably causing the problem. If the tutorial does not start, you probably have a problem with the MacroMedia Flash plug-in for the browser. See the next topic for details.

MacroMedia Flash player is working.

To view the tutorials the browser must have MacroMedia Flash plug-in at version 6.0 or later. To obtain a copy of the plug-in see http://www.macromedia.com/go/getflashplayer. Follow the instructions on the Web site to install the MacroMedia Flash player. After installation is complete, the site provides a way to test the installation. Consult the support section of the MacroMedia Web site if you have problems during installation or are unable to get the plug-in working.

Tutorials are installed correctly.

The Integrated Solutions Console uses Eclipse to provide a help system. The help system contents are stored in a number of plug-ins, which are directories with a defined structure. The Administration Center uses two plug-ins for its help topics: com.tivoli.dsm.admincenter.help.doc.5.3.0.0 and com.tivoli.dsm.admincenter.demo. The second of these plug-ins contains the tutorials.

Verify that the tutorials plug-in exists and contains a doc.zip file. Plug-ins are located in <isc root>/PortalServer/ISCEclipse/plugins where <isc root> is the Integrated Solutions Console installation directory. For example, on Windows look in c:\program files\IBM\ISC\PortalServer\ISCEclipse\plugins for the directory com.tivoli.dsm.admincenter.demo. On UNIX look in /opt/IBM/ISC/PortalServer/ISCEclipse/plugins for that directory. If you see the directory and it contains the doc.zip file, the plug-in is installed correctly. Otherwise, you will need to run the Administration Center installation again to install the help topics.

Verified all of the above items, but the tutorials are still not working.

If none of the previous actions identified the problem, collect the following information before calling support:

- View the source of the help pop-up window. On most browsers a right mouse click will provide a menu with a View Source option. A window should then appear with the HTML source code for that window. Write down the title of that window, which will be the URL or the name of the file the help system is attempting to display.
- Turn on tracing for the SNAPIN trace class and try running the tutorial again. Stop tracing and run the collection option of the service tool.

Main Menu | Red Error Messages: Portlet is unavailable/Application error

Red Error Messages: Portlet is unavailable/Application error

Click the error text in the table or skip the table and scroll through the information below:

Diagnosing an Administration Center red error message
• This portlet is unavailable. If the problem persists, please contact the portal administrator. • There has been an application error! Please contact your system administrator to report this error.

This portlet is unavailable. If the problem persists, please contact the portal administrator.

Typically, this is caused when the Java Server Page (JSP) page has code to initialize a table model or a list model. If there is a connection problem to the Tivoli Storage Manager server, it is possible that the bean that holds all the initial values for the fields on this page was not initialized properly. For example, the reference to a table model in the bean for the Client Nodes Notebook is not initialized and has a null value. This causes an exception when the page is rendered. The message is displayed by the portal server and not the Tivoli Storage Manager Administration Center. Refer to tracing instructions for the Administration Center to turn on the trace and look for any exceptions in the trace. You may see messages like this:

```
  [12/7/04 13:43:37:859 PST] 1a33b798 snapin        e com.tivoli.dsm.snapin.
DsmLaunchedPortlet doView ERROR: view class call failed
  [12/7/04 13:43:37:859 PST] 1a33b798 snapin        e com.tivoli.dsm.snapin.
DsmLaunchedPortlet doView The following exception was logged
                    com.ibm.wps.pe.pc.legacy.impl.TransferPortletException: Cannot set a null
Table Model
    at com.ibm.wps.pe.pc.legacy.impl.PortletContextImpl.include(PortletContextImpl.java:270)
    at com.tivoli.dsm.snapin.views.PolicyDomainsView.performView(PolicyDomainsView.java:184)
    at com.tivoli.dsm.snapin.DsmLaunchedPortlet.doView(DsmLaunchedPortlet.java:727)
    at org.apache.jetspeed.portlet.PortletAdapter.service(PortletAdapter.java:154)  ...
  ...
  ...
  ...
  Nested Exception is java.lang.IllegalArgumentException: Cannot set a null Table Model
    at com.ibm.psw.wcl.components.table.WTable.setModel(WTable.java:2376)
    at com.ibm.psw.wcl.components.table.WTable.setModel(WTable.java:2362)
    at org.apache.jsp._node_5F_prop_5F_notebook._jspService(_node_5F_prop_5F_notebook.
java:2098)
    at com.ibm.ws.webcontainer.jsp.runtime.HttpJspBase.service(HttpJspBase.java:89)
```

In this case, check that you can connect to the Tivoli Storage Manager server. Close the portlet and try the same action again.

Another cause might be a compile error of the Java Server Page (JSP) page which should have not happened unless the page has been modified. Look for any exceptions in the trace. You may see messages like this:

```
  [10/31/04 23:11:32:484 PST] 4c5b22a4 snapin        e com.tivoli.dsm.snapin.DsmMainPortlet
doView The following exception was logged
                    com.ibm.wps.pe.pc.legacy.impl.TransferPortletException: /jsp/console/
tpvw_tbl_libraries.jsp(155,4) Unable to load class com.tivoli.dsm.taglib.DsmRefreshTag
    at com.ibm.wps.pe.pc.legacy.impl.PortletContextImpl.include(PortletContextImpl.java
(Compiled Code))
    at com.tivoli.dsm.snapin.views.TpvwView.performView(TpvwView.java:348)
  ...
  ...
  ...
  Nested Exception is org.apache.jasper.compiler.CompileException: /jsp/console/
tpvw_tbl_libraries.jsp(155,4) Unable to load class com.tivoli.dsm.taglib.DsmRefreshTag
    at com.ibm.ws.webcontainer.jsp.compiler.BasicTagBeginGenerator.init
(BasicTagBeginGenerator.java:81)
    at org.apache.jasper.compiler.JspParseEventListener$GeneratorWrapper.init
(JspParseEventListener.java(Compiled Code))
    at org.apache.jasper.compiler.JspParseEventListener.addGenerator(JspParseEventListener.
java(Compiled Code))
```

This error cannot be recovered unless the JSP page is fixed. Make sure the JSP was not modified. If the problem persists, consider reinstalling the Tivoli Storage Manager Administration Center.

Return to Red Error Messages: Portlet is unavailable/Application error

There has been an application error! Please contact to your system administrator to report this error.

Typically, this error is caused by clicking a button or link more than once before the action is completed. For example, pressing the "Cancel" button on a notebook. If the request is not completed, and the "Cancel" button is pressed again, the application might fail and issue this error message.

Close the application and log on to ISC again.

Return to Red Error Messages: Portlet is unavailable/Application error

Return to Main Menu
Main Menu | Data Protection for Domino

Data Protection for Domino

Data Protection for Domino(R) provides support for protecting Lotus Domino databases and transaction logs.

Follow these steps for diagnosing or reporting problems with Data Protection for Domino:

Steps to diagnose Data Protection for Domino

1. How do I trace the Data Protection client?
2. Where do I go to locate solutions for the Data Protection client?
3. What information should I gather before calling IBM?
4. What files should I gather before calling IBM?

How do I trace the Data Protection client?

The Data Protection client uses the Tivoli Storage Manager API for communicating with the Tivoli Storage Manager server and providing data management functions. Refer to the Backup/Archive client section for a list of available trace classes.

To enable tracing, specify both of the following options on the command line when launching Data Protection for Domino or GUI executable:

/TRACEFILE=*trace file name*
/TRACEFLAG=*trace flags*

where

trace file name=
The name of the file that the trace data will be written to. Note that on UNIX(R) environments, the domdsmc executable is launched by Domino startup script. This script changes the current working directory to the Domino Data directory; therefore, if a full path name is not specified the trace-file-name will be created in Domino Data directory.

trace flags=
The list of *trace flags* to enable. Trace flags are separated by a comma.
- Specify "ALL" to turn on all possible tracing
- Specify "SERVICE" to turn on a subset of trace flags that the service group normally requires
- Specify "API" to trace calls to both the Tivoli Storage Manager API and the Notes C API

For example, issue the following for the command line client:

DOMDSMC SELECTIVE database.nsf /TRACEFILE=trace.log /TRACEFLAG=SERVICE,API

Issue the following for the GUI client:

DOMDSM /TRACEFILE=trace.log /TRACEFLAG=ALL

Return to steps to diagnose Data Protection for Domino

Where do I go to locate solutions for the Data Protection client?

A number of resources are available to learn about or to diagnose the Data Protection client.

The IBM Support Solutions database contains a knowledge base of articles and information on issues related to backup and restore problems. Refer to Support for Data Protection for Domino to review this information.

Return to steps to diagnose Data Protection for Domino

What information should I gather before calling IBM?

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

Gather as much of the following information as possible before contacting IBM Support. This information will assist IBM Support with resolving your problem.

- Exact level of the operating system, including all patches that have been applied.
- Exact level of the Domino server.
- Exact level of Data Protection for Domino.
- Exact level of the Tivoli Storage Manager API.
- Exact level of the Tivoli Storage Manager server.
- Exact level of the Tivoli Storage Manager backup-archive client.
- Exact level of the Tivoli Storage Manager storage agent (if LAN-free environment).
- Tivoli Storage Manager server platform and operating system level.
- Output from the Tivoli Storage Manager server query system command.
- Output from the Data Protection for Domino "DOMDSMC QUERY ADSM" command.
- Output from the Data Protection for Domino "DOMDSMC QUERY DOMINO" command.
- Permissions and name of the user ID being used to run backup and restore operations.
- List of the steps needed to re-create the problem (if the problem is re-creatable).
- If the problem is not re-creatable, list the steps that caused the problem.
- Is the problem occurring on other Domino servers?

Return to steps to diagnose Data Protection for Domino

What files should I gather before calling IBM?

A number of log files and other data may be collected by the Data Protection client.

Gather as many of the following files as possible before contacting IBM Support. The contents of these files will assist IBM Support with resolving your problem.

- Data Protection for Domino configuration file (default: domdsm.cfg).

- Data Protection for Domino log file (default: domdsm.log). This file indicates the date and time of a backup, data backed up, and any error messages or completion codes. This file is important and should be monitored daily. This log file is commonly renamed in scheduled backup scripts by the /LOGFILE parameter specified with the domdsmc backup command.
- Tivoli Storage Manager API options file (default: dsm.opt).
- Tivoli Storage Manager API error log file (default: dsierror.log).
- Output from failed command or operation (redirected console or screen image).
- Tivoli Storage Manager Server activity log. The Data Protection client logs information to the server activity log. A server administrator can view this log for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- If the Data Protection client is configured for LAN-free data movement, also collect the options file for the Tivoli Storage Manager storage agent. The default name for this file is: dsmsta.opt.
- Any command files and scripts used to run Data Protection for Domino.
- If the Tivoli Storage Manager client scheduler is being used, also collect the client schedule log.

Return to steps to diagnose Data Protection for Domino

Main Menu | Data Protection for ESS

Data Protection for ESS

Data Protection for ESS provides support for protecting Oracle or DB2(R) using Enterprise Storage Server(R) (ESS).

If you encounter a problem during Data Protection for ESS processing, follow these steps as an initial attempt to resolve the problem:

1. Retry the operation that failed.
2. If the problem occurred during an incremental FlashCopy(R) backup, run a generic FlashCopy backup. If the generic FlashCopy backup completes successfully, retry the operation that failed. If the problem occurred during FlashCopy restore, try a restore from the Tivoli Storage Manager server.
3. If the problem still exists:
    1. Shut down the Oracle or DB2 server.
    2. Start the Oracle or DB2 server again.
    3. Run the operation that failed.
4. If the problem still exists:
    1. Shut down the entire machine.
    2. Start the machine again.
    3. Run the operation that failed.
5. If the problem still exists, determine if it is occurring on other DB2 or Oracle servers.

Follow these steps for diagnosing or reporting problems with Data Protection for ESS:

Steps to diagnose Data Protection for ESS

1. How do I trace the Data Protection client?
2. Where do I go to locate solutions for the Data Protection client?
3. What information should I gather before calling IBM?
4. What files should I gather before calling IBM?

How do I trace the Data Protection client?

The Data Protection client uses the Tivoli Storage Manager API for communicating to the Tivoli Storage Manager server and providing data management functions. Refer to the Backup/Archive client section for a list of available trace classes.

To enable tracing, specify both of the following options in the user setup file when invoking the Data Protection for ESS (Oracle and DB2) command-line executables:

TRACEFILE = *trace file name*
TRACEFLAGS = *trace flags*

where

trace file name =
The name of the file to write the trace data.
trace flags =
The list of *trace flags* to enable. Trace flags are separated by a space. The trace flags specific to Data Protection for ESS are tdph and tdph_detail. You can also specify Tivoli Storage Manager backup-archive client and Tivoli Storage Manager API trace flags.

For example:
tracefile /home/Data Protectioness/log/trace.out
traceflags tdph tdph_detail api api_detail appl verbinfo timestamp

Make sure to remove the corresponding tracefile and traceflags options from the files pointed to by the DSM_CONFIG and DSMI_CONFIG environment variables.

Return to steps to diagnose Data Protection for ESS

Where do I go to locate solutions for the Data Protection client?

A number of resources are available to learn about or to diagnose the Data Protection client.

The IBM Support Solutions database contains a knowledge base of articles and information on issues related to backup and restore problems. Refer to Support for Data Protection for ESS to review this information. Click the **Hints and Tips, Solutions, and Support Flashes** links in the self-help table to access the search tool. Enter a term in order to search through solutions to previously encountered issues.

Return to steps to diagnose Data Protection for ESS

What information should I gather before calling IBM?

The Data Protection client is dependent upon the operating system as well as the Oracle or DB2 application. Collecting all the necessary information about the environment can significantly assist in determining the problem.

Gather as much of the following information as possible before contacting IBM Support. This information will assist IBM Support with resolving your problem.

- Exact level of the AIX operating system, including all the maintenance levels and patches that have been applied.
- Are you running in an HACMP environment? If so, determine the exact level of the HACMP software including all the maintenance levels and patches that have been applied.
- Exact level of the Oracle or DB2 server, including all the maintenance levels and patches that have been applied.
- If Oracle, are you running in a single server or an OPS or RAC environment? If so, how many hosts are participating in this clustered environment?
- If DB2, are you running in a single partition (EE) or a multipartition (EEE) environment with multiple physical and logical partitions? If so, how many hosts are participating in this clustered environment?
- Exact level of the ESS Server and the ESS Copy Services CLI. Is your ESS FlashCopy feature code enabled?
- Are you using SDD? If so, the exact level of the SDD software.
- Is your database server attached to ESS using SCSI or fibre channel?
- Exact level of Data Protection for ESS (Oracle or DB2).
- Exact level of Data Protection for Oracle.
- Exact level of the Tivoli Storage Manager API.

- Exact level of the Tivoli Storage Manager server.
- Exact level of the Tivoli Storage Manager backup-archive client.
- Exact level of the Tivoli Storage Manager storage agent (if LAN-free environment).
- Tivoli Storage Manager server platform and operating system level.
- Output from the Tivoli Storage Manager server query system command.
- Output from the failing Data Protection for ESS commands on the production and backup systems.
- List of the steps needed to re-create the problem. Note that if the problem is not easily recreatable, then list the steps that caused the instance of the problem that was encountered.
- Does the problem occur on other DB2 or Oracle database servers?

Return to steps to diagnose Data Protection for ESS

What files should I gather before calling IBM?

A number of log files and other data may be collected by the Data Protection client.

Gather as many of the following files as possible before contacting IBM Support. The contents of these files will assist IBM Support with resolving your problem.

- Error log files in the directories pointed to by DSM_LOG and DSMI_LOG on the production and backup systems. Data Protection for ESS error log file (default: tdpess.log) indicates the date and time of the commands, and any error messages or completion codes. This file is important and should be monitored daily.
- Trace files on the production system and the backup system.
- The user setup file used for the Data Protection for ESS commands.
- The tempfile created by Data Protection for ESS (this file is used on the backup system for Tivoli Storage Manager backup).
- The metadata file created by Data Protection for ESS (this file is used for FlashCopy restore).
- Options files pointed to by DSM_DIR, DSMI_DIR, DSM_CONFIG, DSMI_CONFIG on both the production and backup systems.
- Tivoli Storage Manager Server activity log. The Data Protection client logs information to the server activity log. A Tivoli Storage Manager administrator can view this log for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- If the Data Protection client is configured for LAN-free data movement, also collect the options file for the Tivoli Storage Manager storage agent. The default name for this file is: dsmsta.opt.
- Any command files and scripts used to run Data Protection for ESS commands.
- If the Tivoli Storage Manager client scheduler is being used, also collect the client schedule log.

Return to steps to diagnose Data Protection for ESS

Main Menu | Data Protection for ESS for mySAP.com

Data Protection for ESS for mySAP.com

Data Protection for ESS for mySAP.com provides support for protecting Oracle or DB2(R) databases with mySAP.com using Enterprise Storage Server(R) (ESS).

If you encounter a problem during Data Protection for ESS for mySAP.com processing, follow these steps as an initial attempt to resolve the problem:

Follow these steps for diagnosing or reporting problems with Data Protection for ESS for mySAP.com:

Steps to diagnose Data Protection for ESS for mySAP.com

1. How do I trace the Data Protection client?
2. Where do I go to locate solutions for the Data Protection client?
3. What information should I gather before calling IBM?
4. What files should I gather before calling IBM?

How do I trace the Data Protection client?

The Data Protection client uses the Tivoli Storage Manager API for communicating with the Tivoli Storage Manager server and providing data management functions. Refer to the Backup/Archive client section for a list of available trace classes.

To enable tracing, specify the following options in the user setup file (init<SID>.fcs) when invoking the Data Protection for ESS for mySAP.com (Oracle and DB2) command-line executables:

LOG_TRACE_DIR *complete path*
: The *complete path* specifies a directory where all the run logs and trace files are placed.

TRACE YES
: Activate the trace for production and backup.

Both parameters are also described with more details in the *Data Protection for IBM ESS for mySAP.com Technology Installation and User's Guide*.

You can find the logs and traces in the directories specified in parameter LOG_TRACE_DIR of the Data Protection for ESS profile. If no parameter is specified, the logs and traces will be placed in the directory as specified in the parameter WORK_DIR of the Data Protection for ESS profile. The file-naming convention for logs and traces is as follows:

- splitint_b_<splitint function>_<date time stamp>.log
- splitint_p_<splitint function>_<date time stamp>.log
- splitint_b_<splitint function>_<date time stamp>.trace
- splitint_p_<splitint function>_<date time stamp>.trace

where

_b_
: The function is running on the backup system

_p_
: The function is running on the production system

Also see the *Data Protection for IBM ESS for mySAP.com Technology Installation and User's Guide*, Appendix C, for more information.

Return to steps to diagnose Data Protection for ESS for mySAP.com

Where do I go to locate solutions for the Data Protection client?

A number of resources are available to learn about or to diagnose the Data Protection client.

The first reference to consider is the *Data Protection for IBM ESS for mySAP.com Technology Installation and User's Guide*. This is available from the IBM Tivoli Publications Web site.

For the latest news, the IBM Support Solutions database contains a knowledge base of articles and information on issues related to backup-and-restore problems. Refer to Support for Data Protection for Hardware to review this information. Click the H**ints and Tips, Solutions, and Support Flashes** links in the self-help table to access the search tool. Enter a term in order to search through

solutions to previously encountered issues.

Additional information can be found at these Web sites:

- For SAP, you can find more information at SAP Service Marketplace.
- For DB2 Universal Database(TM), you can visit the DB2 Product Family Web site.
- For Oracle, you can search Oracle Technology Network, or visit the Oracle Web site.

Return to steps to diagnose Data Protection for ESS for mySAP.com

What information should I gather before calling IBM?

The Data Protection client is dependent upon the operating system as well as the Oracle or DB2 application. Collecting all the necessary information about the environment can significantly assist in determining the problem.

Gather as much of the following information as possible before contacting IBM Support. This information will assist IBM Support with resolving your problem.

- Exact level of the AIX(R) operating system, including all the maintenance levels and patches that have been applied.
- Are you running in an HACMP environment? If so, determine the exact level of the HACMP software including all the maintenance levels and patches that have been applied.
- Exact level of the Oracle or DB2 server, including all the maintenance levels and patches that have been applied.
- If Oracle, are you running in a single server or an OPS or RAC environment? If so, how many hosts are participating in this clustered environment?
- If DB2, are you running in a single partition (EE) or a multipartition (EEE) environment with multiple physical and logical partitions? If so, how many hosts are participating in this clustered environment?
- Exact level of the ESS Server and the ESS Copy Services CLI. Is your ESS FlashCopy(R) feature code enabled?
- Are you using SDD? If so, the exact level of the SDD software.
- Is your database server attached to ESS using SCSI or fibre channel?
- Exact level of Data Protection for mySAP.com.
- Exact level of the Tivoli Storage Manager API.
- Exact level of the Tivoli Storage Manager server.
- Exact level of the Tivoli Storage Manager backup-archive client.
- Exact level of the Tivoli Storage Manager storage agent (if LAN-free environment).
- Tivoli Storage Manager server platform and operating system level.
- Output from the Tivoli Storage Manager server query system command.
- Output from the failing Data Protection for ESS commands on the production and backup systems.
- List of the steps needed to re-create the problem. Note that if the problem cannot be easily re-created, then list the steps that caused the instance of the problem that was encountered.
- Does the problem occur on other DB2 or Oracle database servers?
- Type of file system (JFS, JFS2)
- Layout of the file system of the mySAP server in relation with the layout of the ESS (file systems, logical volumes, volume groups, LUNs, LSSs)
- Does the problem occur on other mySAP.com database servers?

Return to steps to diagnose Data Protection for ESS for mySAP.com

What files should I gather before calling IBM?

A number of log files and other data may be collected by the Data Protection client.

Gather as many of the following files as possible before contacting IBM Support. The contents of these files will assist IBM Support with resolving your problem.

- Error log files in the directories pointed to by DSM_LOG and DSMI_LOG on the production and backup systems. Data Protection for ESS error log file (default: tdpess.log) indicates the date and time of the commands, and any error messages or completion codes. This file is important and should be monitored daily.
- Log files on SAPs BR*Tools.
- Trace files on the production system and the backup system.
- The Data Protection for Enterprise Storage Server (ESS) profile. This file ends in ".fcs".
- The Data Protection for Enterprise Storage Server (ESS) target volumes file as specified by the parameter SHARK_VOLUMES_FILES. This file ends in ".fct".
  Options files pointed to by DSM_DIR, DSMI_DIR, DSM_CONFIG, DSMI_CONFIG on both the production and backup systems.
- Tivoli Storage Manager Server activity log. The Data Protection client logs information to the server activity log. A Tivoli Storage Manager administrator can view this log for you if you do not have a Tivoli Storage Manager administrator ID and password.
- If the Data Protection client is configured for LAN-free data movement, also collect the options file for the Tivoli Storage Manager storage agent. The default name for this file is: dsmsta.opt.
- Any command files and scripts used to run Data Protection for ESS commands.
- Operating System error info if any (run the command: errpt -a).
- If the Tivoli Storage Manager client scheduler is being used, also collect the client schedule.

Return to steps to diagnose Data Protection for ESS for mySAP.com

Main Menu | Data Protection for Exchange

Data Protection for Exchange

Data Protection for Exchange provides support for protecting Microsoft(R) Exchange databases.

If you encounter a problem during Data Protection for Exchange processing, follow these steps as your first attempt to resolve the problem:

1. Retry the operation that failed.
2. If the problem occurred during an incremental, differential, or database copy backup, run a full backup. If the full backup completes successfully, retry the operation that failed.
3. If the problem still exists, close other applications, especially those applications that interact with Exchange (anti-virus applications, for example). Retry the operation that failed.
4. If the problem still exists:
   1. Shut down the Exchange server.
   2. Start the Exchange server again.
   3. Run the operation that failed.
5. If the problem still exists:
   1. Shut down the entire machine.
   2. Start the machine again.
   3. Run the operation that failed.
6. If the problem still exists, determine if it is occurring on other Exchange servers.

Follow these steps for diagnosing or reporting problems with Data Protection for Exchange:

Steps to diagnose Data Protection for Exchange
1. Determine if the problem is a Data Protection for Exchange issue or an Exchange issue. 2. How do I trace the Data Protection client? 3. Where do I go to locate solutions for the Data Protection client? 4. What information should I gather before calling IBM? 5. What files should I gather before calling IBM? 6. What should I gather if the silent install failed?

Determine if the problem is a Data Protection for Exchange issue or an Exchange issue.

The Data Protection client interacts closely with Microsoft Exchange. Because of this close interaction, it is necessary to first determine if the problem is with Microsoft Exchange or the Data Protection client.

Follow these steps to try to isolate the source of the error:

1. Try re-creating the problem with the Microsoft NTBACKUP utility. This utility uses a call sequence similar to Data Protection for Exchange to run an online backup. If the problem is re-creatable with NTBACKUP then the problem most likely exists within the Exchange server.
2. Try re-creating the problem with the Microsoft BACKTEST (Exchange Server 5.5) or BACKREST (Exchange 2000 Server, Exchange Server 2003) application. This application can run backups using the Microsoft Exchange APIs. If the problem is re-creatable with BACKTEST or BACKREST then the problem most likely exists within the Exchange server. Microsoft ships BACKTEST or BACKREST with the Exchange Software Developer's Kit (SDK). IBM Service can provide a copy of BACKTEST or BACKREST if you encounter problems obtaining or building this application.
3. If the error message "ACN5350E An unknown Exchange API error has occurred" is displayed, the Exchange server encountered an unexpected situation. Microsoft assistance may be needed if the problem continues.
4. Data Protection for Exchange error messages occasionally contain an HRESULT code. Use this code to search Microsoft documentation and the Microsoft Knowledge Base for resolution information. Some Exchange SDK files contain these messages EDBMSG.H (Exchange Server 5.5) or ESEBKMSG.H (Exchange 2000 Server, Exchange Server 2003).

Return to steps to diagnose Data Protection for Exchange

How do I trace the Data Protection client?

The Data Protection client uses the Tivoli Storage Manager API for communicating with the server and providing data management functions. Refer to the Backup/Archive client section for a list of available trace classes.

To enable tracing, specify both of the following options on the command line when launching the Data Protection for Exchange command-line or GUI executable:

/TRACEFILE=*trace file name*
/TRACEFLAG=*trace flags*

where

trace file name=
The name of the file to write the trace data.
trace flags=
The list of *trace flags* to enable. Trace flags are separated by a comma. Use the SERVICE and API trace classes to turn on a subset of trace flags that the service group normally requires.

The following examples show how to enable tracing, depending on your environment:

Command line client:
TDPEXCC BACKUP SG1 FULL /TRACEFILE=trace.log /TRACEFLAG=SERVICE,API

GUI client:
TDPEXC /TRACEFILE=trace.log /TRACEFLAG=SERVICE,API

Return to steps to diagnose Data Protection for Exchange

Where do I go to locate solutions for the Data Protection client?

A number of resources are available to learn about or to diagnose issues with the Data Protection client.

The IBM Support Solutions database contains a knowledge base of articles and information on issues related to backup-restore problems. Refer to Support for Data Protection for Exchange to review this information. Click the **Hints and Tips, Solutions, and Support Flashes** links in the self-help table to access the search tool. Enter a term in order to search through solutions to previously encountered issues.

The Microsoft Knowledge Base contains articles related to backup-restore problems. To review the Microsoft Knowledge Base, visit Microsoft Support.

Return to steps to diagnose Data Protection for Exchange

What information should I gather before calling IBM?

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

Gather as much of the following information as possible before contacting IBM Support. This information will assist IBM Support with resolving your problem.

- Exact level of the Windows operating system, including all service packs and hotfixes that have been applied.
- Exact level of the Exchange Server, including all service packs and hotfixes that have been applied.
- Exact level of Data Protection for Exchange.
- Exact level of the Tivoli Storage Manager API.
- Exact level of the Tivoli Storage Manager server.
- Exact level of the Tivoli Storage Manager backup-archive client.
- Exact level of the Tivoli Storage Manager storage agent (if LAN-free environment).
- Tivoli Storage Manager server platform and operating system level.
- Output from the Tivoli Storage Manager server query system command.
- Output from the Data Protection for Exchange **tdpexcc query exchange** command.
- Device type (and connectivity path) of the Exchange databases and logs.
- Permissions and name of the user ID being used to run backup and restore operations.
- Name and version of anti-virus software.
- List of third-party Exchange applications running on the system.
- List of other applications running on the system.
- List of the steps needed to re-create the problem (if the problem is re-creatable).
- If the problem is not re-creatable, list the steps that caused the problem.
- Is Data Protection for Exchange running in a Microsoft Cluster Server (MSCS) environment?
- Is the problem occurring on other Exchange servers?

Return to steps to diagnose Data Protection for Exchange

What files should I gather before calling IBM?

A number of log files and other data may be collected by the Data Protection client.

Gather as many of the following files as possible before contacting IBM Support. The contents of these files will assist IBM Support with resolving your problem.

- Data Protection for Exchange configuration file. The default configuration file is tdpexc.cfg.
- Data Protection for Exchange log file. The default log file is tdpexc.log. This file indicates the date and time of a backup, data backed up, and any error messages or completion codes. This file is important and should be monitored daily. This log file is commonly renamed in scheduled backup scripts by the /LOGFILE parameter specified with the tdpexcc backup command.
- Data Protection for Exchange Tivoli Storage Manager API options file. The default options file is dsm.opt.
- Tivoli Storage Manager API error log file. The default error log file is dsierror.log.
- Windows Event Log for Application and System. The Exchange Server logs information to the Windows Event Log. Exchange server

error information can be obtained by viewing the Windows Event Log.
- Tivoli Storage Manager registry hive export.
- Exchange Server registry hive export.
- Tivoli Storage Manager Server activity log. The Data Protection client logs information to the server activity log. A Tivoli Storage Manager administrator can view this log for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- If the Data Protection client is configured for LAN-free data movement, also collect the options file for the Tivoli Storage Manager storage agent. The default name for this file is: dsmsta.opt.
- If the Tivoli Storage Manager client scheduler is being used, also collect the client schedule log.

Return to steps to diagnose Data Protection for Exchange

What should I gather if the silent install failed?

The silent install may not report information about the cause of the failure. To isolate or diagnose a failed silent installation, additional steps must be taken.

If a silent installation fails, gather the following information to assist Customer Support when evaluating your situation:

- Operating system level.
- List of service packs or other fixes applied for the operating system.
- A description of the hardware configuration.
- Installation package (CD-ROM or electronic download) and level.
- Any Windows event log relevant to the failed installation.
- Windows services active during the failed installation (for example, anti-virus software).
- Whether or not you are logged on to the local machine console (not through terminal server).
- Whether or not you are logged on as a local administrator, not a domain administrator. Tivoli does not support cross-domain installations.
- You can create a detailed log file (setup.log) of the failed installation. Run the setup program (setup.exe) in the following manner: setup /v"l*v setup.log".

Return to steps to diagnose Data Protection for Exchange

Main Menu | Data Protection for Informix

Data Protection for Informix

Data Protection for Informix(R) provides support for protecting Informix databases.

Follow these steps for diagnosing or reporting problems with Data Protection for Informix:

Steps to diagnose Data Protection for Informix

1. How do I trace the Data Protection client?
2. Where do I go to locate solutions for the Data Protection client?
3. What information should I gather before calling IBM?
4. What files should I gather before calling IBM?

How do I trace the Data Protection client?

The Data Protection client uses the Tivoli Storage Manager API to communicate with the Tivoli Storage Manager server and to provide data management functions. Refer to the Backup/Archive client section for a list of available trace classes.

To enable tracing for Data Protection for Informix, use the Tivoli Storage Manager API tracing capabilities. Add the following lines to the tsm-api options file, usually dsm.opt file:

trace_file=*trace file name*
trace_flags=*trace flags*

where

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. Trace flags are separated by a space. There are no trace flags specific to Data Protection for Informix. The recommended trace flag is *appl*. You can also specify other Tivoli Storage Manager backup/archive client and Tivoli Storage Manager API trace flags.

For example:
tracefile /home/dpInformix/log/trace.out
traceflags api api_detail appl verbinfo verbdetail timestamp

Return to steps to diagnose Data Protection for Informix

Where do I go to locate solutions for the Data Protection client?

A number of resources are available to learn about or to diagnose the Data Protection client.

The IBM Support Solutions database contains a knowledge base of articles and information on issues related to backup-restore problems. Refer to Support for Data Protection for Informix to review this information. Click the **Hints and Tips, Solutions, and Support Flashes** links in the self- help table to access the search tool. Enter a term in order to search through solutions to previously encountered issues.

For Informix, visit the Informix Web site.

Return to steps to diagnose Data Protection for Informix

What information should I gather before calling IBM?

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

Gather as much of the following information as possible before contacting IBM Support. This information will assist IBM Support with resolving your problem.

- What operating system is the problem being experienced on.
- Exact level of the operating system, including all service packs and hotfixes that have been applied.
- Exact level of the Informix Dynamic server including all service packs and hotfixes that have been applied.
- Exact level of Data Protection for Informix.
- Exact level of the Tivoli Storage Manager API.
- Exact level of the Tivoli Storage Manager server.
- Exact level of the Tivoli Storage Manager backup-archive client.
- Exact level of the Tivoli Storage Manager storage agent (if LAN-free environment).
- Tivoli Storage Manager server platform and operating system level.
- Output from the Tivoli Storage Manager server query system command.
- List of other applications running on the system.
- List of the steps needed to re-create the problem (if the problem is re-creatable).
- If the problem is not re-creatable, list the steps that caused the problem.

- Is the problem occurring on other Informix databases?

Return to steps to diagnose Data Protection for Informix

What files should I gather before calling IBM?

A number of log files and other data may be collected by the Data Protection client.

Gather as many of the following files as possible before contacting IBM Support. The contents of these files will assist IBM Support with resolving your problem.

- Data Protection for Informix Tivoli Storage Manager API options file. For Windows(R), the default options file is dsm.opt or the file referenced by the DSMI_CONFIG environment variable. For UNIX(R), the default options file is dsm.sys in the directory referenced by the DSMI_DIR environment variable.
- Tivoli Storage Manager API error log file. The default API error log is dsierror.log.
- Output from the Informix Dynamic Server online.log.
- onBar Activity Log and Debug Log Files.
- Any Trace files created for the API and Data Protection for Informix.
- Output from failed command or operation. This may be either the console output redirected to a file or an actual screen image of the failure.
- Tivoli Storage Manager Server activity log. The Data Protection client logs information to the server activity log. A Tivoli Storage Manager administrator can view this log for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- If the data protection client is configured for LAN-free data movement, also collect the options file for the Tivoli Storage Manager storage agent. The default name for this file is: dsmsta.opt.
- Any command files and scripts used to run Data Protection for Informix.
- If the Tivoli Storage Manager client scheduler is being used, also collect the client schedule log.

Return to steps to diagnose Data Protection for Informix

Main Menu | Data Protection for mySAP.com

Data Protection for mySAP.com

Data Protection for mySAP.com provides support for protecting Oracle or DB2(R) databases.

Follow these steps for diagnosing or reporting problems with Data Protection for mySAP.com:

Steps to diagnose Data Protection for mySAP.com
1. How do I trace the Data Protection client? 2. Where do I go to locate solutions for the Data Protection client? 3. What information should I gather before calling IBM? 4. What files should I gather before calling IBM?

How do I trace the Data Protection client?

The Data Protection client uses the Tivoli Storage Manager API to communicate with the Tivoli Storage Manager server and to provide data management functions. Refer to the Backup/Archive client section for a list of available trace classes.

To enable tracing, specify both of the following options in the user setup file when invoking the Data Protection for ESS (Oracle and DB2) command-line executables:

TRACE=*trace flags*

TRACEFILE=*backint_%BID.trace*

where

trace flags=
The list of *trace flags* to enable. Trace flags are separated by a space. The Data Protection for mySAP.com trace flags are: ALL FILEIO FILEIO_MAX COMPR COMPR_MAX SYSCALL SYSCALL_MAX MUX MUX_MAX TSM TSM_MAX ASYNC ASYNC_MAX APPLICATION APPLICATION_MAX COMM COMM_MAX DEADLOCK DEADLOCK_MAX PROFILE PROFILE_MAX BLAPI BLAPI_MAX.

If an error occurs early within an operation it is a good idea to set the flag ALL. But setting this flag to trace a complete backup, which may be hundreds of gigabytes, can result in a large trace file. The trace file in this case may be hundreds of megabytes.

backint_%BID.trace=
The value of this parameter specifies a directory and file where all the run logs and trace file are placed. This should be a fully qualified directory. The four letters '%BID' will be replaced during execution by the current backup ID.

Both parameters are also described with more details in the *Data Protection for mySAP.com Technology Installation and User's Guide*.

An example:
tracefile /home/Data Protectioness/log/trace.out
traceflags tdph tdph_detail api api_detail appl verbinfo timestamp

Be sure to remove the corresponding tracefile and traceflags options from the files pointed to by the DSM_CONFIG and DSMI_CONFIG environment variables.

Return to steps to diagnose Data Protection for mySAP.com

Where do I go to locate solutions for the Data Protection client?

A number of resources are available to learn about or to diagnose the Data Protection client.

The first reference to consider is the *Data Protection for mySAP.com Technology Installation and User's Guide*. This is available from the IBM Tivoli Publications Web site. Specifically, refer to Appendix G for details about troubleshooting.

For the latest news, the IBM Support Solutions database contains a knowledge base of articles and information on issues related to backup and restore problems. Refer to IBM Tivoli Storage Manager for Enterprise Resource Planning. Click the **Hints and Tips, Solutions,** and **Support Flashes** links in the self-help table to access the search tool. Enter a term in order to search through solutions to previously encountered issues.

Additional information can be found at these websites:

- For SAP, you can find more information at SAP Service Marketplace.
- For DB2 UDB, you can visit the DB2 Product Family Web site.
- For Oracle, you can search Oracle Technology Network, or visit the Oracle web site.

Return to steps to diagnose Data Protection for mySAP.com

What information should I gather before calling IBM?

The Data Protection client is dependent upon the operating system as well as the Oracle or DB2 application. Collecting all the necessary information about the environment can significantly assist in determining the problem.

Gather as much of the following information as possible before contacting IBM Support. This information will assist IBM Support with resolving your problem.

- What operating system is the problem being experienced on?
- Exact level of the operating system, including all service packs and hotfixes that have been applied.
- Exact level of the Oracle or DB2 Universal Database(TM) (DB2 UDB), including all fix packs that have been applied.
- Exact level of Data Protection for mySAP.
- Exact level of the Tivoli Storage Manager API.
- Exact level of the Tivoli Storage Manager server.
- Exact level of the Tivoli Storage Manager backup-archive client.
- Exact level of the Tivoli Storage Manager storage agent (if LAN-free environment).
- Tivoli Storage Manager server platform and operating system level.
- Output from the Tivoli Storage Manager server query system command.
- List of other applications running on the system.
- List of the steps needed to re-create the problem (if the problem is re-creatable).
- If the problem is not re-creatable, list the steps that caused the problem.
- Is the problem occurring on other Oracle or DB2 UDB databases?

When using DB2 UDB:

- Is Data Protection for mySAP running in a DB2 UDB EEE or ESE environment?

Return to steps to diagnose Data Protection for mySAP.com

What files should I gather before calling IBM?

A number of log files and other data may be collected by the Data Protection client.

Gather as many of the following files as possible before contacting IBM Support. The contents of these files will assist IBM Support with resolving your problem.

- Data Protection for mySAP profile file. The default profile is init<SID>.utl.
- Log files of SAPs BR*Tools
- If using DB2, the Data Protection for mySAP log file is tdpdb2.<SID>.<NODE>.log
- If using Oracle RMAN, the Data Protection for mySAP default log file is sbtio.log.
- If using DB2 Data Protection for mySAP vendor environment file, the default vendor environment file is vendor.env.
- Data Protection for mySAP Tivoli Storage Manager API options file. The default options file is dsm.opt.
- Tivoli Storage Manager API error log file. The default API error log is dsierror.log.
- Output from failed command or operation. This may be either the console output redirected to a file or an actual screen image of the failure.
- With DB2, use the utility 'db2support' for collecting environment data about either a client or server machine. This program generates information about a DB2 server, including information about its configuration and system environment. Any files in the DB2 dump directory, first of all the DB2 diagnostic log db2diag.log.
- Tivoli Storage Manager Server activity log. The Data Protection client logs information to the server activity log. A Tivoli Storage Manager administrator can view this log for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- If the Data Protection client is configured for LAN-free data movement, also collect the options file for the Tivoli Storage Manager storage agent. The default name for this file is: dsmsta.opt.
- Any command files and scripts used to run Data Protection for mySAP.
- If the Tivoli Storage Manager client scheduler is being used, also collect the client schedule log.

Return to steps to diagnose Data Protection for mySAP.com

Main Menu | Data Protection for Oracle

Data Protection for Oracle

Data Protection for Oracle provides support for protecting Oracle databases.

Follow these steps for diagnosing or reporting problems with Data Protection for Oracle:

Steps to diagnose Data Protection for Oracle

1. How do I trace the data protection client?
2. Where do I go to locate solutions for the data protection client?
3. What information should I gather before calling IBM?
4. What files should I gather before calling IBM?

How do I trace the data protection client?

The data protection client uses the Tivoli Storage Manager API for communicating with the Tivoli Storage Manager server and providing data management functions. Refer to the Backup/Archive client section for a list of available trace classes.

To enable tracing for Data Protection for Oracle, add the following lines to the tdpo.opt file:

tdpo_trace_file *trace file name*
tdpo_trace_flags *trace flags*

where

tdpo_trace_file=
The name of the file to write the trace data.
tdpo_trace_flags=
The list of *trace flags* to enable. Trace flags are separated by a space. The trace flags specific to Data Protection for Oracle are orclevel0, orclevel1, orclevel2.

For example:
tdpo_trace_file /home/dpOracle/log/trace.out
tdpo_trace_flags orclevel0 orclevel1 orclevel2

Return to steps to diagnose Data Protection for Oracle

Where do I go to locate solutions for the data protection client?

A number of resources are available to learn about or to diagnose the data protection client.

The IBM Support Solutions database contains a knowledge base of articles and information on issues related to backup-restore problems. Refer to Support for Data Protection for Oracle to review this information. Click the **Hints and Tips, Solutions,** and **Support Flashes** links in the self-help table to access the search tool. Enter a term in order to search through solutions to previously encountered issues.

For Oracle, you can search Oracle's Oracle Technology Network, or visit the Oracle Web site.

Return to steps to diagnose Data Protection for Oracle

What information should I gather before calling IBM?

The data protection client is dependent upon the operating system as well as the Oracle. Collecting all the necessary information about the environment can significantly assist in determining the problem.

Gather as much of the following information as possible before contacting IBM Support. This information will assist IBM Support with resolving your problem.

- What operating system is the problem being experienced on.
- Exact level of the operating system, including all service packs and hotfixes that have been applied.
- Exact level of the Oracle database, including all service packs and hotfixes that have been applied.
- Exact level of Data Protection for Oracle.
- Exact level of the Tivoli Storage Manager API.
- Exact level of the Tivoli Storage Manager server.
- Exact level of the Tivoli Storage Manager backup-archive client.
- Exact level of the Tivoli Storage Manager storage agent (if LAN-free environment).
- Rman output.
- Tivoli Storage Manager server platform and operating system level.
- Output from the Tivoli Storage Manager server query system command.
- Output from the Data Protection for Oracle tdpoconf showenv command.
- List of other applications running on the system.
- List of the steps needed to re-create the problem (if the problem is re-creatable).
- If the problem is not re-creatable, list the steps that caused the problem.
- Is Data Protection for Oracle running in a Oracle OPS or RAC environment?
- Is the problem occurring on other Oracle databases?

Return to steps to diagnose Data Protection for Oracle

What files should I gather before calling IBM?

A number of log files and other data may be collected by the data protection client.

Gather as many of the following files as possible before contacting IBM Support. The contents of these files will assist IBM Support with resolving your problem.

- Data Protection for Oracle configuration file. The default configuration file is tdpo.opt.
- Data Protection for Oracle log file. The default log file is tdpoerror.log.
- Data Protection for Oracle Tivoli Storage Manager API options file. The default options file is dsm.opt.
- Tivoli Storage Manager API error log file. The default API error log is tdpoerror.log.
- Output from failed command or operation. This may be either the console output redirected to a file or an actual screen image of the failure.
- Any .trc files in Oracle's user_dump_destination directory.
- Sbtio.log, should be in the $ORACLE_HOME/rdbms/log or user_dump_destination.
- Tivoli Storage Manager Server activity log. The data protection client logs information to the server activity log. A Tivoli Storage Manager administrator can view this log for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- If the data protection client is configured for LAN-free data movement, also collect the options file for the Tivoli Storage Manager storage agent. The default name for this file is: dsmsta.opt.
- Any command files and scripts used to run Data Protection for Oracle.
- If the Tivoli Storage Manager client scheduler is being used, also collect the client schedule log.

Return to steps to diagnose Data Protection for Oracle

Main Menu | Data Protection for SQL

Data Protection for SQL

Data Protection for SQL provides support for protecting SQL databases.

Follow these steps for diagnosing or reporting problems with Data Protection for SQL:

Steps to diagnose Data Protection for SQL

1. How do I trace the Data Protection client?
2. Where do I go to locate solutions for the Data Protection client?
3. What information should I gather before calling IBM?
4. What files should I gather before calling IBM?
5. What should I gather if the silent install failed?

How do I trace the Data Protection client?

The Data Protection client uses the Tivoli Storage Manager API for communicating with the Tivoli Storage Manager server and providing data management functions. Refer to the Backup/Archive client section for a list of available trace classes.

To enable tracing, specify both of the following options on the command line when launching the Data Protection for SQL command-line or GUI executable:

/TRACEFILE=*trace file name*
/TRACEFLAG=*trace flags*

where

trace file name=
The name of the file to write the trace data.
trace flags=
The list of *trace flags* to enable. Trace flags are separated by a comma. Use the SERVICE and API trace classes to turn on a subset of trace flags that the service group normally requires.

See the following examples:

Command line client:
TDPSQLC BACKUP pubs FULL /TRACEFILE=trace.log /TRACEFLAG=SERVICE,API

GUI client:
TDPSQL /TRACEFILE=trace.log /TRACEFLAG=SERVICE,API

Return to steps to diagnose Data Protection for SQL

Where do I go to locate solutions for the Data Protection client?

A number of resources are available to learn about or to diagnose the Data Protection client.

The IBM Support Solutions database contains a knowledge base of articles and information on issues related to backup-restore problems. Refer to Support for Data Protection for SQL to review this information. Click the **Hints and Tips, Solutions,** and **Support Flashes** links in the self-help table to access the search tool. Enter a term in order to search through solutions to previously encountered issues.

The Microsoft Knowledge Base contains articles related to backup-restore problems. To review the Microsoft Knowledge Base, visit Microsoft Support.

Return to steps to diagnose Data Protection for SQL

What information should I gather before calling IBM?

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

Gather as much of the following information as possible before contacting IBM Support. This information will assist IBM Support with resolving your problem.

- Exact level of the Windows operating system, including all service packs and hotfixes that have been applied.
- Exact level of the SQL Server, including all service packs and hotfixes that have been applied.
- Exact level of Data Protection for SQL.
- Exact level of the Tivoli Storage Manager API.
- Exact level of the Tivoli Storage Manager server.
- Exact level of the Tivoli Storage Manager backup-archive client.
- Exact level of the Tivoli Storage Manager storage agent (if LAN-free environment).
- Tivoli Storage Manager server platform and operating system level.
- Output from the Tivoli Storage Manager server query system command.
- Output from the Data Protection for SQL **TDPSQLC QUERY SQL *** command.
- Device type (and connectivity path) of the SQL databases and logs.
- Permissions and name of the user ID being used to run backup-restore operations.
- List of third-party SQL applications running on the system.
- List of other applications running on the system.
- List of the steps needed to re-create the problem (if the problem is re-creatable).
- If the problem is not re-creatable, list the steps that caused the problem.
- Is Data Protection for SQL running in a Microsoft Cluster Server (MSCS) environment?
- Is the problem occurring on other SQL servers?

Return to steps to diagnose Data Protection for SQL

What files should I gather before calling IBM?

A number of log files and other data may be collected by the Data Protection client.

Gather as many of the following files as possible before contacting IBM Support. The contents of these files will assist IBM Support with resolving your problem.

- Data Protection for SQL configuration file. The default configuration file is tdpsql.cfg.
- Data Protection for SQL log file. The default log file is tdpsql.log. This file indicates the date and time of a backup, data backed up, and any error messages or completion codes. This file is important and should be monitored daily. This log file is commonly renamed in scheduled backup scripts by the /LOGFILE parameter specified with the tdpsqlc backup command.
- Data Protection for SQL Tivoli Storage Manager API options file. The default options file is dsm.opt.
- Tivoli Storage Manager API error log file. The default error log file is dsierror.log.
- SQL Server log file. The SQL Server logs information to the SQL Server log file. SQL Server log information can be viewed using Enterprise Manager by selecting Server->Management->SQL Server Logs->Current or Archive #n.
- Windows Event Log for the Application and System event logs.
- Tivoli Storage Manager registry hive export.
- SQL Server registry hive export.
- Tivoli Storage Manager Server activity log. The Data Protection client logs information to the server activity log. A Tivoli Storage Manager administrator can view this log for you if you do not have a Tivoli Storage Manager administrator user ID and password.
- If the Data Protection client is configured for LAN-free data movement, also collect the options file for the Tivoli Storage Manager

storage agent. The default name for this file is: dsmsta.opt.
- If the Tivoli Storage Manager client scheduler is being used, also collect the client schedule log.

Return to steps to diagnose Data Protection for SQL

What should I gather if the silent install failed?

The silent install may not report information about the cause of the failure. To isolate or diagnose a failed silent installation, additional steps must be taken.

If a silent installation fails, gather the following information to assist Customer Support when they evaluate your situation:

- Operating system level.
- List of service packs or other fixes applied for the operating system.
- A description of the hardware configuration.
- Installation package (CD-ROM or electronic download) and level.
- Any Windows event log relevant to the failed installation.
- Windows services active during the failed installation (for example, anti-virus software).
- Whether or not you are logged on to the local machine console (not through terminal server).
- Whether or not you are logged on as a local administrator, not a domain administrator. Tivoli does not support cross-domain installations.
- You can create a detailed log file (setup.log) of the failed installation. Run the setup program (setup.exe) in the following manner: setup /v"l*v setup.log".

Return to steps to diagnose Data Protection for SQL

Main Menu | Hints and tips for device drivers

Hints and tips for device drivers

For problems with device drivers, there are many possible causes. The problem may be with the operating system, the application using the device, the device firmware, or the device hardware itself.

The first question to ask whenever a device problem is encountered is "Has anything been changed?" Changes include anything from the machine trying to use the device to a suspect device - especially if the device worked prior to the change, then stopped working after the change.

If the information presented below doesn't adequately address your device driver issue and/or this is the initial setup of your system's device drivers, please refer to the *Tivoli Storage Manager Administrator's Guide* and *Tivoli Storage Manager Installation Guide*. Finally, double-check that your hardware devices are supported by Tivoli Storage Manager.

Consider the following steps when evaluating problems with device drivers:

Diagnosing a device driver

- Has the operating system changed?
- Has the HBA or SCSI adapter connecting to the device been changed, updated, or replaced?
- Has the adapter firmware changed?
- Has the cabling between the computer and device changed?
- Are any of the cable connections loose?
- Has the device firmware changed?
- Are there error messages in the system error log for this device?
- Do 64-bit/32-bit Linux kernel modules support 32-bit/64-bit applications?
- What are the requirements for running a Tivoli Storage Manager Linux server on x86_64 architecture?
- What change has happened for the HBA drivers on the Linux 2.6.x kernels?
- Additional notes about ddtrace on Linux, using Tivoli Storage Manager Version 5.3.2.
- How is device information of host systems updated on a dynamic SAN without rebooting the host systems?
- What are requirements for Adaptec SCSI and Qlogic Fibre Channel HBA BIOS settings on Linux?

Has the operating system changed?

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 this is not possible, check for device driver updates that may be needed based on this fix level, release or version of the operating system.

Return to diagnosing a device driver

Has the HBA or SCSI adapter connecting to the device been changed, updated, or replaced?

A device driver communicates to a given device through an adapter. If it is a fibre channel-attached device, the device driver will communicate using a host bus adapter (HBA). If the device is SCSI attached, the device driver will communicate using a SCSI adapter. In either case, if the adapter firmware has been updated or the adapter itself has been replaced, the device driver may have trouble using the device.

Work with the vendor of the adapter to verify that it is installed and configured appropriately. Other possible steps include:

- If the adapter was changed, try reverting back to the previous adapter to see if that resolves the issue.
- If other hardware in the computer was changed or the computer was opened, reopen the computer and check to make sure that the adapter is properly seated in the bus. By opening and changing other hardware in the computer, the adapter cards and other connections in the computer may have been loosened which may cause intermittent problems or total failure of devices or other system resources.

Return to diagnosing a device driver

Has the adapter firmware changed?

If the adapter firmware has changed, this may 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.

Return to diagnosing a device driver

Has the cabling between the computer and device changed?

If cabling between the computer and the device has been changed, this often accounts for intermittent or persistent failures.

Check any cabling changes to verify that they are correct.

For SCSI connections, a bent pin in the SCSI cable where it connects to the computer or to the device can cause errors using that device or any device on the same SCSI bus. A cable with a bent pin should be repaired or replaced. Similarly, SCSI buses must be terminated. If a SCSI bus is improperly terminated, devices on the bus may exhibit intermittent problems or data transferred on the bus may be or appear to be corrupted. Check the SCSI bus terminators to insure that they are correct.

Return to diagnosing a device driver

Are any of the cable connections loose?

If a connection is loose from the computer to the cable or from the cable to the device, this may result in problems with 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 may result in difficult problems with one or more devices on that bus.

Return to diagnosing a device driver

Has the device firmware changed?

If the device firmware has changed, this may 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.

Return to diagnosing a device driver

Are there error messages in the system error log for this device?

A device may try to report an error to a system error log. Examples of various system error logs are: errpt for AIX(R), Event Log for Windows(R), and System Log for z/OS(R). The system error logs can be useful because the messages and information recorded may be useful to report the problem or the messages may 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.

Return to diagnosing a device driver

Do 64-bit/32-bit Linux kernel modules support 32-bit/64-bit applications?

The Linux kernel modules dictate the bit mode of the Linux SCSI generic (sg) device driver, all different Host Bus Adapter (HBA) drivers, etc. All of these kernel modules only support applications which have same bit mode with running kernel modules. In other words, the 64-bit kernel modules only support 64-bit applications on 64-bit Linux systems and vice versa.

If a 32-bit application runs on a 64-bit Linux system and invokes a 64-bit kernel module, the 32-bit application will cause a kernel segmentation fault. The same will happen if a 64-bit application invokes a 32-bit kernel module on a 32-bit Linux system.

In order to avoid these kinds of problems, the bit mode of the Linux kernel module and its applications must be same. That means 32-bit applications only can invoke 32-bit kernel modules on 32-bit Linux systems. 64-bit applications only can invoke 64-bit kernel modules on 64-bit Linux systems.

Return to diagnosing a device driver

What are the requirements for running a Tivoli Storage Manager Linux server on x86_64 architecture?

The AMD64 and EM64T systems are 64-bit systems. 32-bit and 64-bit Linux operating systems can run on these 64-bit systems.

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

A 64-bit Tivoli Storage Manager server QUERY SAN command requires a 64-bit HBA API on an AMD64/EM64T system. If an AMD64 system is equipped with a Qlogic HBA, it could create a problem since, by default, Qlogic only provides 32bit HBA API on AMD64 system. So the 64-bit HBA API must be installed on the system before using the 64-bit QUERY SAN command.

Return to diagnosing a device driver

What change has happened for the HBA drivers on the Linux 2.6.x kernels?

1. All drivers have a new suffix "ko" at 2.6.x kernels.

2. Driver's name and location in 2.6.x kernels:

Adaptec: The driver name is aic7xxx.ko and is located at /lib/modules/kernel-level/drivers/scsi/aic7xxx/ directory.

Emulex: The driver name is lpfcdd.ko and is located at /lib/modules/kernel-level/drivers/scsi/lpfc/ directory.

Qlogic: Its driver names are qla2xxx.ko, qla2100.ko, qla2200.ko, qla2300.ko, qla2322.ko etc. There is a certain order to load the HBA drivers. The qla2xxx.ko is a base driver. It should be loaded first. After loading the qla2xxx.ko, the system should then load the qla2300.ko driver if it is equipped with a Qla2300 card. All of drivers are located at /lib/modules/kernel-level/drivers/scsi/qla2xxx/ directory.

What are requirements for multiple LUN support on Linux kernels?

In order 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 add this option manually to the running kernel.

1. To set up multiple LUN enabled on IA32 architecture:

Add one parameter to a boot loader configuration file.

For LILO boot loader:

Add append="max_scsi_luns=128" to the /ect/lilo.conf file.

Run lilo.

For GRUB boot loader:

Add max_scsi_luns=128 after the kernel image list at /etc/grub.conf file for RedHat distribution.

Add max_scsi_luns=128 after the kernel image list at /boot/grub/menu.1 file for SuSE distribution.

2. Reboot the system.

Return to diagnosing a device driver

How can I do a Tivoli Storage Manager device driver trace from Tivoli Storage Manager 5.2.2?

Trace using one of these methods:

- The ddtrace command.
- The ddtrace tool.

Use the following with the ddtrace command and ddtrace tool: ddtrace flag=[EE, FULL, BASE|SYSLOG]

The **ddtrace** command is new for the Linux passthru driver in Tivoli Storage Manager Version 5.3.2. It allows a device driver trace to be performed from the Tivoli Storage Manager server console. If this is issued after **trace enable** and **trace begin**, the trace output it generates is not only stored in the latter part of the Tivoli Storage Manager server trace output file, it is also stored in the trace.out file in the Tivoli Storage Manager device directory. If the ddtrace command is issued without issuing trace enable and trace begin, the device driver trace messages will be the only output in the trace.out file.

For the ddtrace command on the Linux passthru driver, the driver trace output is mixed with the Tivoli Storage Manager server trace in the server trace output file. Trace.out in the Tivoli Storage Manager device directory is an empty file. If you only want the Linux passthru driver trace, set the lpdd option in the trace enable command.

The **ddtrace** tool is a Tivoli Storage Manager driver tool located in the Tivoli Storage Manager device directory. It allows you to get a Tivoli Storage Manager device driver trace at the command line of a Tivoli Storage Manager server host system. The tool only generates the ddtrace.out in the Tivoli Storage Manager device directory. All driver trace messages are stored in this file. This method is not applied to the Linux passthru driver.

For Tivoli Storage Manager Version 5.2.2 and later, the ddtrace command and ddtrace tool are supported for the Tivoli Storage Manager device driver on AIX, HP, Linux (SCSI only), Sun Solaris, and Windows systems.

With Tivoli Storage Manager Version 5.3.2, the ddtrace command is also supported for the Linux Passthru driver.

No additional parameters are required to operate in basic default mode for either the ddtrace command or the ddtrace tool. However, three trace options (trace classes) are available for ddtrace and ddtrace to retrieve additional trace information. They are:

- EE
- FULL
- SYSLOG|BASE

All of these options are used to set special trace flag combinations.

Return to diagnosing a device driver

Additional notes about ddtrace on Linux, using Tivoli Storage Manager Version 5.3.2.

In Tivoli Storage Manager Version 5.3.2 and later, the passthru device driver can be traced through the ddtrace command. To enable trace, issue the following commands from the server console or admin client.

trace enable lpdd <other server trace classe names>

trace begin <file name>

Select one of the following three options:

i. ddtrace start librarydd tapedd (to trace both library and drive)

ii. ddtrace start librarydd (library trace only)

iii. ddtrace start tapedd (drive trace only)

Notes:

DDTRACE GET and DDTRACE END are not required.

The Tivoli Storage Manager passthru device driver trace cannot be enabled through the ddtrace utility.

Return to diagnosing a device driver

How is device information of host systems updated on a dynamic SAN without rebooting the host systems?

When devices in a SAN (Storage Area Network) environment change, the information about this changed environment is not automatically sent to host systems attached to the SAN. If the device information has not been updated to host systems attached to the SAN, previously defined device paths will not exist anymore. If you use the existing device information to define device paths, backup or restore data, these operations may fail. In order to avoid these kinds of failures, you can use different methods for different platforms to update the device information on the SAN without rebooting host systems.

On AIX, issue cfgmgr command to force the operating system to reconfigure itself. Then, run SMIT to reconfigure your Tivoli Storage Manager devices.

On HP, issue ioscan command with –kfn option to force the operating system to rescan SCSI buses and fibre channels. Then, run tsmddcfg to configure your Tivoli Storage Manager devices.

On Linux, there is no system command to reconfigure the operating system. In order 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 autoconf or tsmscsi to reconfigure Tivoli Storage Manager devices on Linux. The lspci command is used to find out which SCSI adapter and fibre channel adapter is available on the system. The rmmod command is used to unload a driver from the kernel and modprobe command to load a driver to the kernel.

HBA adapters and corresponding drivers for all architectures of Linux:

HBA Adapters	HBA Driver Name	Available Architectures
Adaptec 7892	aix7xxx	IA32, AMD64
Qlogic 22xx	qla2200	IA32, AMD64
Qlogic 23xx	qla2300	IA32, AMD64
Qlogic 2362	qla2362	EM64T
Emulex	lpfcdd	IA32, iSeries, pSeries

Return to diagnosing a device driver

What are requirements for Adaptec SCSI and Qlogic Fibre Channel HBA BIOS settings on Linux?

By default, Adaptec SCSI adapters set the multiple LUN (Logical Unit Number) option as off in their BIOS. It will make the SCSI adapter driver not be able to probe a SCSI unit with multi-LUN properly, so this option must be turned on. To do this:

- Press the Ctrl and A keys at the same time.

- Select **SCSI Device Configuration** in **Configure/View Host Adapter Setting**.
- Change No to Yes at Bios Multiple LUN support.

By default, Qlogic Fibre HBAs set the tape enable option as off in their BIOS. This affects the execution of some SCSI commands on several SCSI tape devices, so this option must be turned on. To do this:

- Press the Alt and Q keys at the same time.
- Select Advanced Settings.
- Change Disable to Enable at Fibre Channel Tape Support.

Return to diagnosing a device driver

Main Menu | Hints and tips for hard disk drives and disk subsystems

Hints and tips for hard disk drives and disk subsystems.

To better understand the use of *hard disk drives* and *disk subsystems* by Tivoli Storage Manager, here are definitions:

Hard disk drive
: A hard disk drive storage device, typically installed inside a given computer and used for storage by a Tivoli Storage Manager server on that machine.

Disk subsystem
: An external disk subsystem that connects to a computer through a SAN or some other mechanism. Generally, disk subsystems are outside of the computer to which they are attached and may be located in close proximity or they may be located much farther away. These subsystems may also have some method of caching the I/O requests to the disks. Finally, disk subsystems often have their own configuration and management software.

The Tivoli Storage Manager server may define hard disk drives and disk subsystems to be used by the computer or operating system on the machine where Tivoli Storage Manager is installed. Typically, a hard disk drive or disk subsystem is defined to the computer where Tivoli Storage Manager is installed as a drive or file system. After the hard disk drive or disk subsystem is defined to the operating system, Tivoli Storage Manager may use this space by allocating a database, recovery log, or storage pool volume on the device. At this point, the Tivoli Storage Manager volume looks like another file on that drive or file system.

The Tivoli Storage Manager server requires specific behaviors of hard disk drives, disk subsystems, third-party file systems, and remote file systems. These required behaviors allow Tivoli Storage Manager to appropriately manage and store data by insuring the integrity of the Tivoli Storage Manager server itself.

Tivoli Storage Manager requires the following conditions for hard disk drives, disk subsystems, third-party file systems, and remote file systems:

- When Tivoli Storage Manager opens database, recovery log, and storage pool volumes, they are opened with the appropriate operating system settings to require data write requests to bypass any cache and be written directly to the device. By bypassing cache during write operations, Tivoli Storage Manager can maintain the integrity of client attributes and data. This is required because if an external event, such as a power failure, causes the Tivoli Storage Manager server or the computer the server is installed on to halt or crash while the server is running, the data in the cache may or may not be written to the disk. If the Tivoli Storage Manager data in the disk cache is not successfully written to the disk, information in the server database or recovery log may not be complete or data that was supposed to be written to storage pool volumes may be missing.

  This is less of an issue for hard disk drives installed in the computer where the Tivoli Storage Manager server is installed and running. In this case, the operating system settings used when Tivoli Storage Manager opens volumes on that hard disk drive generally manage the cache behavior appropriately and honor the request to prevent caching of write operations.

  Typically, the use and configuration of caching for disk subsystems is a greater issue. The reason for this is that disk subsystems often do not receive information from the operating system about bypassing cache for write operations or else they ignore this information

when a volume is opened. Because of this, the caching of data write operations may result in corruption of the Tivoli Storage Manager server database or loss of client data, or both, depending upon which Tivoli Storage Manager volumes are defined on the disk subsystem and the amount of data lost in the cache. It is recommended that disk subsystems be configured to not cache write operations when a Tivoli Storage Manager database, recovery log, or storage pool volume is defined on that disk. Another alternative is to use non-volatile cache for the disk subsystem. Non-volatile 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.

- The size and location of Tivoli Storage Manager database, recovery log, and storage pool volumes (files) can not be changed after they are defined and used by the server. If the size is changed or the file is moved, internal information that Tivoli Storage Manager uses to describe the volume may no longer match the actual attributes of the file. If you need to move or change the size of a Tivoli Storage Manager database, recovery log, or storage pool volume, you should move any existing data to other volumes and then delete it from the server prior to altering or moving it.

FILE directory mapping between storage agents and servers for shared files.

Tivoli Storage Manager servers and storage agents can access the same data in File device classes. This is done by defining a set of directories that should be used within a device class definition.

For each storage agent that will be sharing FILE access, the PATHs defined to each DRIVE seen by the storage agent needs to provide access to the same set of directories. When the PATHs are defined, the directories for each storage agent must match in number and ordering for the directories as listed in the device class definition on the server. If these definitions are out of sync, the storage agent may be unable to access the FILE volumes. This will result in the LAN restores being successful while the LANFREE restores result in mount failures.

Main Menu | Hints and tips for tape drives and libraries

Hints and tips for tape drives and libraries

For problems with tape drives and libraries, there are many possible causes. The problem may be with software on the machine trying to use the device, the connections to the device, or the device itself.

The first question to ask whenever a device problem is encountered is, "Has anything been changed?" Anything from the machine trying to use the device to the device itself may be suspect, especially if the device worked prior to a given change then stopped working after that change.

Consider the following steps when evaluating problems with tape drives and libraries. The steps work from the computer trying to access the device out to the device itself. Ask these questions to diagnose the problem:

Diagnosing tape drives and libraries
• Has the operating system changed? • Has a device driver changed? • Has the adapter in the computer been replaced or other hardware changed or fixed? • Has the adapter firmware changed? • Has the cabling between the computer and device changed? • Are any of the cable connections loose? • Has the device firmware changed? • Are there error messages in the system error log for this device?

Has the operating system changed?

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 this is not possible, check for device driver updates that may be needed based on this fix level, release or version of the operating system.

Return to how to diagnose tape drives and libraries

Has a device driver changed?

Upgrading a device driver can result in a tape drive or library can cause a device not to work.

Try reverting to the previous (or earlier) version of the device driver to see if the problem was introduced by the newer version of the driver.

Return to how to diagnose tape drives and libraries

Has the adapter in the computer been replaced or other hardware changed or fixed?

The connecting point for the device to the computer is usually referred to as an adapter. Another term for adapter is card. In the case of a SCSI connection to the device, this is a SCSI adapter. In the case of a fibre-channel (optical) connection to the device, this is a host bus adapter (HBA).

In either case, if actual adapter was changed or the computer was opened and other hardware changed or fixed, this may be the cause of the problem.

- If the adapter was changed, try reverting back to the previous adapter to see if that resolves the issue.
- If other hardware in the computer was changed or the computer was opened, reopen the computer and check to make sure that the adapter is properly seated in the bus. By opening and changing other hardware in the computer, the adapter cards and other connections in the computer may have been loosened which may cause intermittent problems or total failure of the devices or other system resources.

Return to how to diagnose tape drives and libraries

Has the adapter firmware changed?

If the adapter firmware has changed, this may 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.

Return to how to diagnose tape drives and libraries

Has the cabling between the computer and device changed?

If cabling between the computer and the device has been changed, this often accounts for intermittent or persistent failures.

Check any cabling changes to verify that they are correct.

Return to how to diagnose tape drives and libraries

Are any of the cable connections loose?

If a connection is loose from the computer to the cable or from the cable to the device, this may result in problems with 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 may result in difficult problems with one or more devices on that bus.

Return to how to diagnose tape drives and libraries

Has the adapter firmware changed?

If the device firmware has changed, this may 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.

Return to how to diagnose tape drives and libraries

Are there error messages in the system error log for this device?

A device may try to report an error to a system error log. Examples of various system error logs are: errpt for AIX(R), Event Log for Windows(R), and System Log for z/OS(R). The system error logs can be useful because the messages and information recorded may be useful to report the problem or the messages may 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.

Return to how to diagnose tape drives and libraries

Main Menu | Hints and tips for miscellaneous server information

Hints and tips for miscellaneous server information

This document outlines general problems with the Tivoli Storage Manager server that do not quite fit in any of the other server sections.

Running the AIX Tivoli Storage Manager server from a location other than the Installation Directory.

An attempt to run the Tivoli Storage Manager server from a directory other than the installation directory may result in the server failing to start. The following messages might be displayed:

- ANR0109E Attempt number 1 Unable to load cryptography module from '.'.
- ANR0109E Attempt number 2 Unable to load cryptography module from '.'.
- ANR9999D cryptog.c(253): ThreadId<0> ICC load failed.

In order to run a Tivoli Storage Manager server from a directory other than the Tivoli Storage Manager server installation directory, the **DSMSERV_DIR** environment variable must be set to the path of the server installation directory. The Tivoli Storage Manager server uses the installation directory as a reference to locate executables and other files in the Tivoli Storage Manager server installation.

To run a Tivoli Storage Manager server from a directory other than the installation directory do the following:

1. Set the DSMSERV_DIR environment variable to the base installation directory of the Tivoli Storage Manager server (typically: /usr/tivoli/tsm/server/bin).
2. Add the path of the installation directory (/usr/tivoli/tsm/server/bin) to your PATH environment variable.
3. Follow the steps on "Running Multiple Servers on a Single Machine" in the *Tivoli Storage Manager for AIX Installation Guide.*

Hints and tips for a storage area network (SAN)

For problems with a storage area network, there are many possible causes. SAN problems may be with software on the machine trying to use the device, the connections to the device, or the device itself.

The first question to ask whenever a SAN problem is encountered is "Has anything been changed?" Any kind of changes from the machine trying to use the device to the device itself may be suspect, especially if the device worked prior to a given change, then stopped working after that change.

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

Fibre channel
: Fibre channel denotes a fibre-optical connection to a device or component. This is typically abbreviated as FC.

Host bus adapter
: A host bus adapter is used by a given machine to access a storage area network. A host bus adapter is similar in function to a network adapter and how it provides access for a machine to a local area network or wide area network. This is typically abbreviated as HBA.

Storage area network
: A storage area network is a network of shared devices that can typically be accessed using fibre. Often, a storage area network is used to share devices between many different machines. This is typically abbreviated as SAN.

The following should be considered when evaluating problems with a SAN:

Diagnosing a SAN

- Know your SAN configuration.
- Supported devices.
- Considerations for a HBA.
- HBA configuration issues.
- FC switch configuration issues.
- Verify data gateway port settings.
- Verify the SAN configuration between devices.
- Monitor the fibre channel link error report.

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
: This is the simplest configuration. The devices are connected directly to the 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.

Return to diagnosing a SAN

Devices supported.

Many devices or combinations of devices may not be supported in a given SAN. These limitations arise from the ability of a given vendor to certify their device using Fibre Channel Protocols.

For a given device, verify with the device vendor that it is supported in a SAN. This includes whether or not it is supported by the HBAs used in your SAN environment, which means verifying with the vendors of the hubs, gateways, and switches that make up the SAN that this device is supported.

Return to diagnosing a SAN

Considerations for a host bus adapter (HBA).

The HBA is critical device for the proper functioning of a SAN. Problems that arise relating to HBAs range from improper configuration to outdated bios or device drivers.

For a given HBA, check the following:

BIOS
: Host bus adapters have an embedded BIOS that can be updated. The vendor for the HBA will have utilities for updating the BIOS in an HBA. Periodically, the HBAs in use on your SAN should be checked to see if there are BIOS updates that should be applied.

Device driver
: Host bus adapters use device drivers to work with the operating system to provide connectivity to the SAN. The vendor will typically provide a device driver for use with their HBA. Similarly, the vendor will provide instructions and any necessary tools or utilities for updating the device driver. Periodically the device driver level should be compared to what is available from the vendor and if needed, it should be updated to pick up the latest fixes and support.

Configuration
: Host bus adapters typically have a number of configurable settings. The settings below typically affect how Tivoli Storage Manager functions with a SAN device.

Return to diagnosing a SAN

HBA configuration issues.

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

The vendor for the HBA should provide information about the settings for your HBA and the appropriate values for these settings. Similarly, the HBA vendor should provide a utility and other instructions on how to configure your HBA. The settings that typically affect using Tivoli Storage Manager with a SAN are:

SAN topology
: The HBA should be set appropriately based on the SAN topology being used. For example, if your SAN is an arbitrated loop, the HBA should be set for this configuration. If the HBA connects to a Switch, this HBA port should be set to "point to point" and not "loop."

With Tivoli Storage Manager SAN Device Mapping, you can perform SAN discovery on most of the platforms and the persistent binding of the devices are not required. A Tivoli Storage Manager server can find the device if the device path has been changed due to reboot or other reasons.

Go to http://www.ibm.com/support to verify the platform/HBA vendor/driver level support for Tivoli Storage Manager SAN

discovery.

FC link speed
: In many SAN topologies, the SAN can be configured with a maximum speed. For example, if the FC switch maximum speed is 1GB/sec, the HBA should also be set to this value. Or the HBA should be set for automatic (AUTO) negotiation if the HBA supports this capability.

Is fibre channel tape support enabled?
: Tivoli Storage Manager requires that an HBA is configured with tape support. Tivoli Storage Manager typically uses SANs for access to tape drives and libraries. As such, the HBA setting to support tapes must be enabled.

Return to diagnosing a SAN

FC switch configuration issues.

A FC switch typically supports many different configurations. The ports on the switch need to be configured appropriately for the type of SAN that is setup as well as the attributes of the SAN.

The vendor for the switch should provide 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 settings that typically affect how Tivoli Storage Manager uses a switched SAN are:

FC link speed
: In many SAN topologies, the SAN can be configured with a maximum speed. For example, if the FC switch maximum speed is 1GB/sec, the HBA should also be set to this value. Or the HBA should be set for automatic (AUTO) negotiation if the HBA supports this capability.

Port mode
: The ports on the switch need to 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 to loop.

Return to diagnosing a SAN

Verify 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, so it is important that the port settings for a data gateway are correct.

The vendor for the data gateway should provide information about the appropriate settings and configuration based upon the SAN topology being deployed and SCSI devices being used. Similarly, the vendor should provide a utility and other instructions on how to configure it. The following settings can be used for the FC 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 may also initiate and manage data movement operations. Specifically, there are extended SCSI commands that allow for third-party data movement. By setting a given port as an initiator, it is eligible to be used for third-party 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 may also initiate and manage data movement operations. Specifically, there are

extended SCSI commands that allow for third-party data movement. By setting a given port as an initiator, it is eligible to be used for third-party data movement SCSI requests.

Return to diagnosing a SAN

Verify the 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 should provide a utility for configuration. As part of this configuration utility, there is typically information about how this device is configured and other information that this device sees about the SAN topology that it is apart of. These vendor utilities can be used to verify the SAN configuration between devices:

Data gateway
: A data gateway should report all the FC devices as well as the SCSI devices that are available in the SAN.

Switch
: A switch should report information about the SAN fabric.

Tivoli Storage Manager Management Console
: The Tivoli Storage Manager management console will display device names and the paths to those devices. This can be useful to help verify that the definitions for Tivoli Storage Manager match what is actually available.

Return to diagnosing a SAN

Monitor 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 will typically report errors also. The errors that are often seen are:

CRC error, 8b/10b code error, and other similar symptoms
: These are usually recoverable errors. The error handling for these cases is usually provided by firmware or hardware. In most cases, the method to recover the device is to retransmit the failing frame. The FC link is still active when these errors are encountered. Applications using a SAN device that encounters 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 is not able to successfully retransmit the data after repeated attempts. The recovery for these types of errors is typically very fast and will not cause system performance to be degraded.

Link failure (loss of signal, loss of synchronization, NOS primitive received)
: This indicates that a link is actually "broken" for a period of time. It is likely a faulty gigabit interface connector (GBIC), media interface adapter (MIA) or cable. The recovery for this type error is disruptive. This error will appear in the application using the SAN device that encountered this link failure. The recovery is at the command exchange level and involves the application and device driver having to perform a reset to the firmware and hardware. This will cause the system to run degraded until the link recovery is complete. These errors should be monitored closely, as they typically will affect multiple SAN devices. Note: often these errors are caused by a 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 this is the case, the time and duration of the error should correspond to when the service activity was performed.

Return to diagnosing a SAN

Main Menu | SAN Device Mapping Problems | SAN Device Mapping Hints and Tips| SAN devices

SAN Device Mapping Hints and Tips

Starting with Tivoli Storage Manager 5.3.0, Storage Area Network device discovery and device mapping are supported on Windows 2000, Windows 2003 (32 bits), AIX, Solaris, and Linux (except Linux zSeries).

The advantages of Tivoli Storage Manager SAN device discovery and device mapping are the following:

Tivoli Storage Manager can display all the devices on the SAN.

The QUERY SAN server command shows all the devices seen by the Tivoli Storage Manager server via all the Fibre Channel HBAs installed on the system. The parameters shown are device type, vendor name, product model name, serial number and the device name. If FORMAT=DETAIL is specified for the query, additional information such as World Wide Name (WWN), port, bus, target, and LUN is displayed. This information will help identify all the tape, disk, and Data Mover devices on the SAN. For AIX, the data mover is transparent and it will not be shown.

Tivoli Storage Manager can update the device path automatically when a device's path changes.

Tivoli Storage Manager does not require persistent binding for the devices it sees via HBA. Instead, the server uses the SNIA HBAAPI to discover and obtain the serial number for all the devices on the SAN. It can also determine each device's path. By comparing a device's serial number recorded in the Tivoli Storage Manager database with the serial number obtained from the device in real time, a change in a device's path can be detected. If the path has been changed, SAN discovery will be performed automatically to obtain the new path for the device desired. The Tivoli Storage Manager database will also be updated with the new path information.

The HBAAPI wrapper library is the wrapper used by the Tivoli Storage Manager server to communicate with the SNIA HBAAPI. The HBAAPI wrapper library is installed in the same directory as the Tivoli Storage Manager executable (unless the full path is given shown below). This following is a list of the HBA wrapper files shipped with the Tivoli Storage Manager server package (except on AIX):

- Windows(R): hbaapi.dll
- AIX(R): /usr/lib/libhbaapi.a (This is provided by AIX with HBAAPI installation)
- 32-bit Linux: libhbaapi32.so
- 64-bit Linux: libhbaapi64.so
- 32-bit Solaris: libhbaapi32.so
- 64-bit Solaris: libhbaapi64.so

If these files are missing, message "ANR1791W HBAAPI wrapper library xxxxxxxxx failed to load or is missing." will be displayed.

If additional help not discussed below is needed, try the SAN Device Mapping Problems page.

Consider the following items when evaluating problems with SAN device mapping or device discovery:

Diagnosing SAN device mapping
• How do I disable SAN device mapping? • Important platform specific information.

How do I disable SAN device mapping?

Sometimes it is necessary to disable SAN device mapping to circumvent a problem or to isolate a problem when trouble-shooting device problems.

To disable SAN device mapping and device discovery, issue the following server command:

- setopt SANDISCOVERY OFF

This will disable SAN discovery. The Tivoli Storage Manager server will not be able to update the device path automatically if the path has changed. This only needs to be issued once.

- setopt SANDISCOVERY ON

  This will enable SAN discovery.

The above setopt SANDISCOVERY commands can be issued as many times as needed.

Another way to disable/enable SAN discovery is to put the following option in the dsmserv.opt file

- SANDISCOVERY OFF

  This will disable SAN discovery.

- SANDISCOVERY ON

  This will enable SAN discovery.

SANDISCOVERY ON is the default for the Windows platform.

SANDISCOVERY OFF is the default for all non-Windows platforms.

Return to diagnosing SAN device mapping

Important platform specific information.

- For a list of supported HBAs and required driver levels by operating system go to:

  IBM Support and search on keyword TSMHBA.

- AIX :

  Query SAN command will NOT show any Gateway devices because Gateway devices are transparent to AIX.

- Linux :

  There are separate libraries, utilities, etc for RHEL3U3. To run them you must also install an Emulex ioctl kernel module in addition to the Emulex driver. Make sure to load the Emulex driver before loading the ioctl module.

  Emulex has provided an Application kit for RHEL3. To find the Emulex application kit: go to Emulex website and click on "drivers, software, and manuals." Under "HBAs installed on OEM equipment," select vendor "IBM." Then, select "HBA" model name. The application kit found here provides some utilities such as apitest64 to test HBAAPI functions.

Return to diagnosing SAN device mapping Main Menu | Help | Help for Backup/Archive client

Help for Backup/Archive client

The command line client interface includes a help facility that provides the following:

- Description and syntax for client commands and options
- Full description of client messages, including **Explanation**, **System Action**, and **User Response**

Help for commands

The command line client help facility is accessed with the **help** command:

```
dsmc help
```

You will be presented with a menu of help topics. Each topic has a corresponding menu number. You may enter the menu number for a topic, or use the **u** and **d** keys to scroll through the menu pages. Use the **q** key to quit the help facility. For instance, to locate help for the **restore** command, use the **d** key until the page containing the RESTORE topic is displayed:

```
The following help topics are available.
Enter the number of the desired help topic or 'q' to quit,
'd' to scroll down, 'u' to scroll up.

57 - QUERY SESSION
58 - QUERY SYSFILES
59 - QUERY SYSTEMINFO
60 - QUERY SYSTEMOBJECT
61 - QUERY SYSTEMSERVICES
62 - QUERY SYSTEMSTATE
63 - QUERY SYSVOL
64 - QUERY WAS
65 - QUERY WMI
66 - RESTART RESTORE
67 - RESTORE
68 - RESTORE ACTIVEDIRECTORY
69 - RESTORE ASR
70 - RESTORE BACKUPSET
71 - RESTORE CERTSERVERDB
72 - RESTORE CLUSTERDB
73 - RESTORE COMPLUSDB
74 - RESTORE EVENTLOG
75 - RESTORE FRS
```

In this example, you would then enter **67** to view information on the **restore** command.

To view information for a message, use the **d** and **u** keys to navigate the menu until you begin to see entries that look like this:

```
The following help topics are available.
Enter the number of the desired help topic or 'q' to quit,
'd' to scroll down, 'u' to scroll up.

The following help topics are available.
Enter the number of the desired help topic or 'q' to quit,
'd' to scroll down, 'u' to scroll up.

266 - ANS1820 - ANS1839v 267 - ANS1840 - ANS1859
268 - ANS1860 - ANS1879
269 - ANS1880 - ANS1899
270 - ANS1900 - ANS1919
271 - ANS1920 - ANS1939
272 - ANS1940 - ANS1959
273 - ANS1960 - ANS1979
274 - ANS2000 - ANS2019
275 - ANS2020 - ANS2039
276 - ANS2040 - ANS2059
277 - ANS2600 - ANS2619
278 - ANS2620 - ANS2639
```

```
279 - ANS3000 - ANS3019
280 - ANS4000 - ANS4019
281 - ANS4020 - ANS4039
282 - ANS4040 - ANS4059
283 - ANS4900 - ANS4919
284 - ANS4980 - ANS4999
```

Notice that each topic covers a range of message numbers. For example, to look up help information for message ANS1895I, you would select menu item **269** because ANS1895I is within the range of messages ANS1880 - ANS1899. This will display help information for all messages within that range. Use the navigation keys to move forward and backward through the information until you find the text for message ANS1895I:

```
local filesystem.

-----------------------------------------------------------------------
ANS1895I Highest macro return code was return code value.

Explanation: This message is issued after all commands in a client macro
have completed. The return code represents the highest return code that
was issued during processing of the macro. In order of increasing
severity, the return code meanings are:
0 - The command completed successfully
4 - One or more files were skipped
8 - One or more warning messages were issued
12 - One or more error messages (except for skipped files) were
issued

System Action: None.

User Response: For return codes other than 0, the user may wish to
verify the results of the client operation(s) and take diagnostic and
repair actions, as necessary.


>> Press enter to continue, 'q' to return to help menu
'u' to scroll up, or 'd' to scroll down: <<
```

Help for the Windows NT(R) Service Configuration Utility (dsmcutil)

To obtain help information, run dsmcutil as follows:

```
dsmmcutil help
```

You will be prompted to enter **C** for a basic command summary or **F** for full help information.

The basic command summary is just that: a listing of available dsmcutil subcommands and options. This information is most useful if you are already familiar with using dsmcutil.

The full help information will display more comprehensive information within the Windows(R) help utility.

Help for the GUI and Web GUI clients

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

Main Menu | Administration Center | Help does not display

Help does not display

Topics

Verify the following:

1. The help system is running.
2. Help is installed correctly.

Other topics of interest:

- Starting and stopping the help system.
- Adding publications to the Integrated Solutions Console help server.
- Reporting a problem with a help topic.

The help system is running.

When the help system is not running, you will not be able to view the help. The first step is to verify the help server process is running.

On Windows:
Open the Task Manager and select the process tab.

On UNIX:
Run a command like *ps -ef | grep eclipse*.

Look for either eclipse.exe or EclipseSrv.exe in the list of processes. When you do not see either process name in the list, the help system is not running. You will need to start the help system.

When there are multiple processes with the same names in the list, or continued attempts to start the help server fail, further investigation is necessary. This situation can arise if you are using the publications CD on the same system, or some other application is using the port. For example, if the random port number the publications CD chooses is 8423 (the port for the help system), then the Integrated Solutions Console help server will fail to start.

Try pointing your Web browser to http://<machine name>:8423. For example, http://wscheid.tucson.ibm.com:8423 would show the help system on wscheid.tucson.ibm.com. You should see the title, Integrated Solutions Console, on the top left hand corner of the screen. If you do not see that title, then some other help system or program is running on port 8423. Try stopping that application and starting the Integrated Solutions Console help server again.

Help is installed correctly.

The Integrated Solutions Console uses Eclipse to provide a help system. The help system contents are stored in a number of plug-ins, which are directories with a defined structure. The Administration Center uses two plug-ins for its help topics: com.tivoli.dsm.admincenter.help.doc.5.3.0.0 and com.tivoli.dsm.admincenter.demo. The first of these plug-ins contains the help content for the Administration Center at Version 5.3.0. (Future versions of the Tivoli Storage Manager Administration Center might change the version number, the numerical characters in the plug-in name.)

Verify that the help plug-in exists and contains a doc.zip file. Plug-ins are located in <isc root>/PortalServer/ISCEclipse/plugins where <isc root> is the Integrated Solutions Console installation directory. For example, on Windows look in c:\program files\IBM\ISC\PortalServer\ISCEclipse\plugins for the directory com.tivoli.dsm.admincenter.help.doc.5.3.0.0. On UNIX look in /opt/IBM/ISC/PortalServer/ISCEclipse/plugins for that directory. If you see the directory and it contains the doc.zip file, the plug-in is installed correctly. Otherwise, you will need to run the Administration Center installation again to install the help topics.

Starting and stopping the help system.

On Windows
The help server is a service on Windows. Open the service manager and locate the ISC Help Server service. To start the service, right click on the server and select Start from the pop-up menu. To stop the service, right click on the service name and select Stop from the pop-up menu.

On UNIX
To start or stop the help server on UNIX, change directory to <isc root>/PortalServer/ISCEclipse. You should see two scripts, startEclipse.sh and stopEclipse.sh. Run startEclipse.sh to start the help server. Run stopEclipse.sh to stop the help server.

See The help system is running to determine if the server is running.

Adding the publications to the Integrated Solutions Console help server

The publication CD contains a number of Eclipse plug-ins that you can copy from the publication CD to your Integrated Solutions Console help server.

To install the publications into the help server do the following:

1. Copy any directory that starts with com.ibm from ibm_help_win\eclipse\plugins on the publication CD to <isc root>/PortalServer/ISCEclipse/plugins, where <isc root> is the installation location of the Integrated Solutions Console.
2. Stop the help server. See Starting and stopping the help system for more information
3. Start the help server. See Starting and stopping the help system for more information

The help system should now contain the publications. When you run into any problems, verify that the directory structure on the publication CD matches the one in <isc root>/PortalServer/ISCEclipse/plugins. You might not need all publications; you can choose which plug-ins to copy. However, you must copy the following plug-ins. Otherwise you will have problems when attempting to access the publications in the help system.

- com.ibm.help.doc
- com.ibm.help.doc.nl1
- com.ibm.help.ic.doc
- com.ibm.support.itsm.core.doc.nl1_1.0.0
- com.ibm.support.itsm.core.doc_1.0.0
- com.ibm.support.itsm.doc_1.0.0

Reporting a problem with a help topic.

When reporting a problem with the help system, collect the following information before calling support:

- Record what you clicked on to get the help. For example, if you clicked the question mark for a portal, record the name of the portal.
- View the source of the help pop-up window. On most browsers a right mouse click will provide a menu with a View Source option. A window should then appear with the HTML source code for that window. Write down the title of that window, which will be the URL or the name of the file that the help system is attempting to display.

Main Menu | Help | Help for server or storage agent

Help for server or storage agent.

The server and storage agent both provide a help facility. It provides the following information:

- Description and syntax for server commands
- Full description of server messages which includes **Explanation, System Action, and User Response**

Help for commands.

For help on a server command, issue:

HELP *commandName*

where *commandName* is the server command for which you want information. For example,

HELP REGISTER NODE

This command displays the description of the REGISTER NODE command as well as the syntax and information about the command parameters.

Help for messages.

For help on a server message, issue:

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 2011, it assumes the message prefix ANR and reports the help information for ANR2011. If the message number is specified with the prefix, for example HELP ANR2011, it will report the help information for that message. The following is an example of the information displayed for help on a message:

HELP ANR2011

Help output:

ANR2011W Server is out of DB space in adding entries to the Activity Log. Console messages will not be logged until DB space is available.

Explanation: The process that monitors the default (console) output stream and maintains the activity log cannot update the activity log due to a lack of database space.

System Action: The server does not update the activity log.

User Response: An authorized administrator can issue the DEFINE DBVOLUME command to add volumes for use by the database, and can issue the EXTEND DB command to extend the size of the database so that the new volumes are used.

Main Menu | Administration Center | Health Monitor

Health Monitor

Topics

- How the Health Monitor works.
- Conditions that can cause an Unknown server status.
- Conditions that can cause a Warning or Critical database status.
- Conditions that can cause a Warning or Critical storage status.
- When to resynchronize the ADMIN_CENTER administrator ID password.

How the Health Monitor works.

When the Tivoli Storage Manager Server 5.3.0 is installed or the server is upgraded to 5.3.0 or a later version, the administrator ID ADMIN_CENTER is created, and set to a lock status. When you configure the health monitor or add a server connection and check the unlock ADMIN_CENTER administrator check box, the ID becomes unlocked to allow the health monitor to contact the server for status. If for some reason the ADMIN_CENTER does not exist, the health monitor registers the id on the server. The initial

password is the same as the ID. It is recommended that you change the password via the Configure Health Monitor action from the Health Monitor Server table. The password and the time interval set are saved in the health monitor configuration file "dsmhealthmon.properties". This file is located in the {ISCROOT}/PortalServer/shared/app/config/services directory.

The health monitor uses this ID and password to contact each server that has been added to the Administration Center. When it completes contacting each server, the health monitor "sleeps" for the duration specified in the refresh interval. Only unique instances of servers are contacted. For instance, if two administrators each have WIN64SRV at address 9.11.232.90 on port 1500 configured, that server is only contacted once. The length of time it takes for the health monitor to complete contacting each server in the list may vary widely if servers are down. This is because the health monitor needs to wait for each connection to time-out before proceeding to the next server in the list.

The refresh interval countdown displayed on the health monitor work page is not intended to match the actual cycle the health monitor is using to contact servers. Instead, this is when the work page is refreshed. For example, there are 177 servers connections added to the Administration Center. Of these 177 servers, over half are test servers and are not active. It takes the health monitor approximately 10 minutes to work through the list of 177 servers to complete the check. If the refresh interval for the health monitor is 15 minutes and it takes 10 minutes to cycle through the servers, the servers are contacted about once every 25 minutes (10 minutes for the completion of the contact cycle and 15 minutes of sleep). In the meantime, administrator Eric has left his screen on the health monitor work page. His screen refreshes every 15 minutes whether a new status is available or not.

When server connections are added, the health monitor cycles through its list of servers rather than waiting for the refresh interval, so you should avoid adding or removing server connections unnecessarily.

Updating the configuration settings of the health monitor causes the health monitor to "wake up" and start contacting the list of servers, so you should avoid updating configuration settings unnecessarily.

NOTE: The health monitor needs to be configured only one time, regardless of how many administrators and servers you have defined. That configuration applies to all servers.

Conditions that can cause an Unknown server status

- The server is down: if the server is down, the health monitor is unable to obtain status.
- The ADMIN_CENTER ID not defined: If the ADMIN_CENTER administrative ID has been deleted from the server, the health monitor is unable to obtain status. Redefine the ID.
- The ADMIN_CENTER password for that server does not match the password for the health monitor. All servers must use the id ADMIN_CENTER and the same password. If the password the health monitor uses to contact the server is invalid, the health monitor is unable to obtain status.
- The ADMIN_CENTER ID is locked. The ID must be unlocked or the health monitor is unable to obtain status. Use the Administrator tab of the server properties notebook to update the ID. When adding additional server connections, selecting the Unlock ADMIN_CENTER ID check box unlocks the ID.
- The health monitor worker is no longer running. Tracing is the only way to determine if the worker thread is running. Follow the Administration Center tracing instruction to activate tracing for the HEALTH trace class. Search the PortalServer/logs/trace.log file for strings like, "Time for a nap" or "Back to work." Finding those messages in the trace file indicates the health monitor is running. You may need to wait some time before seeing those messages, depending if the health monitor worker is sleeping for the given interval. When you don't see any activity for a period of time longer than the refresh interval, try restarting the Integrated Solutions Console. This restarts the health monitor worker.

Conditions that can cause a Warning or Critical database status.

The health monitor determines server database status by evaluating a set of rules. These rules are fixed (not changeable by you). No single condition causes the change in status but rather the health monitor looks at a variety of items to determine the overall status. Conditions are evaluated on a point system and if the score is between 0 and 4, status is Normal, points of 5-9 are Warning, and points 10 and over indicate Critical status. The rules may change over time as based on customer feedback.

The following items are evaluated:

1. Don't have a space trigger - 2 points
2. Extension of database is 0 and maximum utilization is at 90% - 5 points

3. No space trigger, no extendable space for the DB, and current utilization of DB over 90% - 10 points
4. Extension space exists and utilization is over 85% but less than 95% - 5 points
5. Extension space exists and utilization is over 94% - 10 points
6. No database backup in the last 24 hours - 3 points
7. Cache hit ratio is less than 98% - 3 points
8. Log utilization is currently over 80 % but less than 90% - 3 points
9. Log utilization is currently over 89 % but less than 95% - 5 points
10. Log utilization is currently over 94% - 5 points
11. Log maximum utilization is over 90% - 3 points

Conditions that can cause a Warning or Critical storage status.

Storage status works on the same principles as the database status.

1. Determine the percentage of drives offline in a given library.
2. Determine the percentage of paths offline for each drive in a given library.
3. Find the average of the numbers in the above step. This represent the score for the drive path status.
4. Determine the percentage of library paths offline for the given library.
5. The score for the library is the maximum of the three scores. The number should be somewhere in the range from zero and ten.
6. Repeat the above steps for every library on the server.

The overall health of the storage hardware is the average of all libraries' scores. This is the overall health of the storage hardware on a scale from zero to ten. A score of four or less represents Normal status. A score greater than four but less than eight represents Warning status. A score greater than seven represents a Critical status.

Taking all drives in the library offline makes the status critical for the library. Taking all paths for all drives in a single library offline causes the status to be Critical. Taking the only path to the library offline makes the library status Critical. Depending on the number of libraries on the server, the Critical status for a single library is diluted by other healthy libraries. For example, if the status one library is Critical and two libraries are Normal, the storage status is Warning. To view the individual status of the library, use the Storage Devices work item.

When to resynchronize the ADMIN_CENTER administrator ID password.

When the health monitor is configured, the password for the ADMIN_CENTER administrator ID is saved in the health monitor configuration. The same password is used to contact all the servers on the list. As long as the password being updated is done via the Configure Health Monitor action from the Health Monitor Server table, the password is synchronized for all the servers. When the ADMIN_CENTER password is changed from the command line or the Server Properties notebook, however, the health monitor does know the new password to update the configuration file and the following occurs:

1. The health monitor contacts the server and the server returns an invalid credential message.
2. After the third try, the ADMIN_CENTER ID is locked.

Unlocking the ADMIN_CENTER ID via the command line or the Server Properties notebook does not work because the health monitor does not have a valid password. This action simply unlocks the ID; then after three tries, the ID is locked again.

Reconfiguring the health monitor to update the password also does not work because the health monitor attempts to log on using the password it knows about - which is no longer valid on that server.

In either situation, resynchronize the ADMIN_CENTER password on all the servers; this should solve the problem. The health monitor will update the ADMIN_CENTER password to the current health monitor password for all the servers.

Main Menu | gt script

gt script

```
#!/bin/ksh

# this script uses gdb to extract thread info from a core file
# input parameters are the path/name of the executable (default ./dsmserv)
#                  and the path/name of the core file (default ./dsmcore)
# output file is dsm_gdb.info
#
# note that files named dsm_gdb.cmd and dsm_gdb.info will be overwritten
#
# if you see the following error:
#   ./dsm_gdb.cmd:9: Error in source command file:
#   No symbol table is loaded.  Use the "file" command.
# then comment out the line that prints buildStringP
#
# if you see other errors, you're on your own ...

exe=${1:-"./dsmserv"}           # get parm 1 (executable file path/name), set default
core=${2:-"./dsmcore"}  # get parm 2 (core file path/name), set default

echo " "

# look for the executable file ... quit if not found

if [[ -f $exe ]]; then
  echo "using executable file:" $exe
else
  echo "didn't find executable file ("$exe") ... exiting"
  exit
fi

# look for the core file, if not found, look for ./core ... quit if not found

if [[ -f $core ]]; then
  echo "using core file:" $core
else
  if [[ -f ./core ]]; then
    echo "didn't find core file ("$core") but found ./core ... renaming to" $core
    mv ./core $core
    echo "using core file:" $core
  else
    echo "didn't find core file ("$core") ... exiting"
    exit
  fi
fi

echo " "

# make gdb command file to get thread info

nl="\0134\0156"    # octal codes for \n (so echo won't think it's \n)

echo "# dsm gdb command file" >|dsm_gdb.cmd

echo "define doit" >>dsm_gdb.cmd
echo "info registers" >>dsm_gdb.cmd # show register values
echo "echo" $nl >>dsm_gdb.cmd
echo "where" >>dsm_gdb.cmd  # show function traceback
```

```
echo "echo" $nl"==========================================================================="$nl
>>dsm_gdb.cmd
echo "end" >>dsm_gdb.cmd

echo "echo" $nl"==========================================================================="$nl
$nl >>dsm_gdb.cmd
echo "x/s buildStringP" >>dsm_gdb.cmd
echo "echo" $nl"==========================================================================="$nl
$nl >>dsm_gdb.cmd
echo "info threads" >>dsm_gdb.cmd # show thread info
echo "echo" $nl"==========================================================================="$nl
>>dsm_gdb.cmd
echo "thread apply all doit" >>dsm_gdb.cmd
echo "quit" >>dsm_gdb.cmd

echo "invoking gdb to get thread info (watch for errors) ..."
echo "if you see:"
echo ". warning: The shared libraries were not privately mapped; setting a"
echo ". breakpoint in a shared library will not work until you rerun the program"
echo "that's ok."

echo "if you see:"
echo ". ./dsm_gdb.cmd:x: Error in source command file:"
echo "then type 'quit', edit this script, and read the comments at the top"

gdb -se $exe -c $core -x ./dsm_gdb.cmd >|dsm_gdb.info
rm dsm_gdb.cmd   # done with this now

exit
```

Main Menu | ODBC driver

ODBC driver

For issues with the ODBC driver, review this information:

Diagnosing ODBC driver issues
1. Supported ODBC applications and functions 2. Defining data sources 3. Known problems and limitations 4. Troubleshooting and diagnostics 5. Configuration details 6. Tracing

Main Menu | Help facilities

Help facilities

Help facilities for...?

1. Backup/Archive client
2. Server or storage agent

Main Menu | Other sources for information

Other sources for information

Reference	Description	Link if available
ADSM-L	Marist University provides a listserv for Tivoli Storage Manager. Many experienced Tivoli Storage Manager users monitor and contribute to the listserv. This is used as a forum to exchange information about the product and to ask questions of other Tivoli Storage Manager users. This is not maintained, managed, or owned by IBM.	Information about subscribing to ADSM-L
Administrator's Guide	General concepts about the product as well as specific discussions on the functions and features provided by Tivoli Storage Manager.	IBM Tivoli Publications
Administrator's Reference	Reference for all supported server commands.	IBM Tivoli Publications
Macintosh - Backup/Archive Clients Installation and User's Guide	Installation and information about using the backup/archive client for Macintosh.	IBM Tivoli Publications
Netware - Backup/Archive Clients Installation and User's Guide	Installation and information about using the backup/archive client for Netware.	IBM Tivoli Publications
Quick Start	General concepts about the server as well as setup and configuration information.	IBM Tivoli Publications
Redbooks for Tivoli Storage Manager	In-depth analysis and examples of how to use and implement Tivoli Storage Manager.	IBM Redbooks
Storage Agent User's Guide	General concepts about using storage agents for lanfree data movement as well as setup and configuration information.	IBM Tivoli Publications
IBM Technical Support Search	From the IBM Web site (www.ibm.com), it is possible to search for known problems or issues. The technical support Web site provides this search capability.	IBM Technical Support Search
UNIX - Backup/Archive Clients Installation and User's Guide	Installation and information about using the backup/archive client for UNIX.	IBM Tivoli Publications
Using the Application Programming Interface	Tivoli Storage Manager provides an application programming interface (API). This book discusses how to write an application that uses the API and the functions it provides.	IBM Tivoli Publications
Windows - Backup/Archive Clients Installation and User's Guide	Installation and information about using the backup/archive client for Microsoft(R) Windows (R) operating systems.	IBM Tivoli Publications
IBM Software Support Handbook	Support guidelines and recommendations for IBM software products.	IBM Software Support Handbook
Personalized Support Page	Using your registered IBM support ID and password, this Web site provides personalized support and information.	Personalized Support Page

IBM Software Support	The main IBM software support Web site. All IBM software products are supported from this site.	IBM Software Support
IBM Passport Advantage Site	For Passport Advantage(R) customers, this is the main site for the services that this provides.	IBM Passport Advantage Site
Web/ETR Problem Submission site	Web submission of problems. Problems submitted from this site go directly to the product Level 2 (L2) support.	Web/ETR Problem Submission site
IBM Support FAQS	IBM support site for frequently asked questions (FAQS).	IBM Software Support FAQS
IBM End of Service (EOS)	A list of end-of-service dates for various IBM software products.	IBM End of Service (EOS)
Tivoli Storage Manager Requirements	The Web site for submitting requirements for future enhancements or changes to Tivoli Storage Manager.	Tivoli Storage Manager Requirements
Tivoli Support Page	The Web site for Tivoli Storage Manager support.	Tivoli Storage Manager Support
Tivoli Storage Manager Customer Document Repository	Anonymous document ftp upload site for customers to upload documentation for problems being diagnosed with service. Documents uploaded to this site will be purged after three days.	Anonymous FTP Site
Tivoli(R) Storage Manager Service Strategy	Service strategy for Tivoli Storage Manager. This provides information about how frequently service will be provided and the ways that fixes are shipped.	Service Strategy
Tivoli Storage Manager Data Protection product Users' Guides	Installation and information about using the IBM Tivoli Storage Manager for applications, databases, enterprise resource planning, hardware, and mail.	IBM Tivoli Publications.

Main Menu | Problem Determination Steps

Problem Determination Steps

The following questions should be considered and if possible answered when trying to diagnose a problem:

1. What is the problem?
2. Where did it occur?
3. When did it begin happening?
4. What action was being performed?
5. Were any messages issued?
   - Check the server activity log and the client logs for error messages.
   - If error messages are in the server activity log, check 30 minutes before and after the time that the error message was issued. Often the problem encountered is actually a symptom of another problem and seeing the other error messages that were issued may help to isolate this.
   - Did the Explanation or User Response section of the Tivoli Storage Manager message offer any suggestions on how to resolve the problem? See, Understanding Tivoli Storage Manager Messages for additional information.
6. How frequently does this error occur?
7. Check any system error logs:

   On Windows(R)

   Check the Windows Event logs. These are comprised of the System and Application event logs, but can also include other event logs depending on the Windows operating system.

   On AIX(R) and other UNIX(R) platforms

Check the error report.

8. Check with others that may have made changes in the environment that could affect Tivoli Storage Manager. Some others in a typical IT environment include:
   - SAN Administrator
   - Network Administrator
   - Database Administrator
   - Client or machine owners
9. Check the Tivoli Storage Manager logs:

   dsmserv.err
   : Server error file. This is located on the same machine as the server. The dsmserv.err file is typically in the server install directory. Note that the storage agent may also create a dsmserv.err file to report errors.

   dsmsvc.err
   : Server error file. This is located on the same machine as the server. This file is equivalent to the dsmserv.err file when Tivoli Storage Manager is running as a Windows Service.

   dsmcrash.log, dsmcrash.dmp
   : Client crash log. This is located on the same machine as the client. These files are for Windows clients only. dsmcrash.log and dsmcrash.dmp contain diagnostic information that is written to the client machine when a client program exception (crash) occurs. When reporting a Windows client crash problem to IBM Technical Support, send these files to the IBM support engineer for analysis.

   dsmerror.log
   : Client error log. This is located on the same machine as the client.

   dsmsched.log
   : Client log for scheduled client operations. This is located on the same machine as the client.

   dsmwebcl.log
   : This log file contains messages related to the client acceptor daemon (dsmcad) and the remote client agent (dsmagent).

   db2diag.log, db2alert.log, userexit.log
   : DB2(R) log files. These are useful when troubleshooting a problem when backing up a DB2 database using Tivoli Data Protection for DB2. These are located on the same machine where DB2 is installed. See the DB2 documentation for additional information about what is in these files and where they are located.

   tdpess.log
   : Default error log file used by the Data Protection for Enterprise Storage Server(R) client.

   tdpexc.log
   : Default error log file used by the Data Protection for Exchange client.

   dsierror.log
   : Default error log for the client API.

   domdsm.log
   : Default error log for the Data Protection for Domino client.

   tdpoerror.log
   : Default error log for the Data Protection for Oracle client.

   tdpsql.log
   : Default error log for the Data Protection for SQL client.
10. Verify that devices are still accessible to the system and to Tivoli Storage Manager.
11. Search the online Knowledge Base for matching error messages or problem descriptions.
12. Test other operations to better determine the scope and impact of the problem. This may also help to determine if it is a specific sequence of events that causes the problem.
13. Other useful tips:

    Check if the problem is specific to a particular Tivoli Storage Manager client component. For example, if the problem is observed in the scheduler or GUI client, test whether the problem can be reproduced with the command line client.

    Eliminate complexity. For example, try reducing the options file to a minimum:

    ```
    commmethod tcpip
    tcpserveraddress your.tsm.server.address
    nodename yournodename
    *nodename optional if you use the machine name
    passwordaccess generate
    *passwordaccess optional if you don't use generate
    ```

then see if the problem reoccurs.
If the problem occurs when running a script, try running each command separately to see if the failure can be detected.

Consider the entire Tivoli Storage Manager environment, including all Tivoli Storage Manager software components, the client and server machines and their hardware, and everything in between. While a problem can surface through a particular Tivoli Storage Manager component, that does not mean that the diagnostic effort should be limited to that component. For example, investigating a client TCP/IP connection failure can include an examination of the Tivoli Storage Manager client logs, the Tivoli Storage Manager server activity log, the Tivoli Storage Manager client and Tivoli Storage Manager server machines' network adapters and cables, other network hardware, firewall settings, router firmware levels, etc.

Main Menu | Reporting a problem

Reporting a problem.

For support for this or any Tivoli(R) product, you can contact IBM Customer Support in one of the following ways:

- Visit the Tivoli Storage Manager technical support Web site.
- Submit a problem management record (PMR) electronically at IBMLink.
- Submit a problem management record (PMR) electronically at IBM Software Support.
- Customers in the United States can also call 1-800-IBM-SERV (1-800-426-7378).

International customers should consult the Web site for customer support telephone numbers.

Hearing-impaired customers should visit the TDD/TTY Voice Relay Services and Accessibility Center Web site.

You can also review the IBM Software Support Guide or the Tivoli Software information center.

When you contact IBM Software Support, be prepared to provide identification information for your company so that support personnel can readily assist you. Company identification information is needed to register for online support available on the Web site.

The support Web site offers extensive information, including a guide to support services (IBM Software Support Guide); frequently asked questions (FAQs); and documentation for all IBM Software products, including Release Notes, Redbooks, and white papers, defects (APARs), and solutions. The documentation for some product releases is available in both PDF and HTML formats. Translated documents are also available for some product releases.

All Tivoli publications are available for electronic download or order from the IBM Publications Center.

We are very interested in hearing about your experience with Tivoli products and documentation. We also welcome your suggestions for improvements. If you have comments or suggestions about our documentation, please complete our customer feedback survey at www.ibm.com/software/sysmgmt/products/ support/IBMTivoliStorageManager.html by selecting the Feedback link in the left navigation bar.

Submitting a problem manager record (PMR).

When service receives a PMR from the field, there is a specific set of data that service will need immediately in order to start working on diagnosing the problem. There is a utility now available for the Tivoli Storage Manager server and storage agents, which eases and speeds up the process of gathering this data. The utility is a small application called **tsmdiag.** It is strongly recommended that users run **tsmdiag** and submit the data that is collected by this tool with the PMR. See Installing and running tsmdiag.

Have the following information ready when you report a problem:

1. The data collected by tsmdiag on the server or storage agent.
2. The Tivoli Storage Manager client version, release, modification, and service level number. You can get this information by entering dsmc at the command line.
3. The activity that you were doing when the problem occurred, listing the steps that you followed before the problem occurred.

4. A description of the symptom or error encountered.
5. The exact text of any error messages relating to the symptom or error encountered.
6. Be prepared to provide any error logs or other related documentation for the problem.

Main Menu | Show commands

Show commands

SHOW commands are undocumented and unsupported diagnostic commands used to display information about in-memory control structures and other run-time attributes. These are used by development and service only as diagnostic tools. Depending upon the information that a show command displays, there may be cases where the information is changing or cases where it may cause the application (client, server, or storage agent) to crash. These should only be used at the recommendation of development or service.

The show commands listed are those that are most typically requested or used for diagnosing problems. This list does not discuss all possible show commands that are available.

SHOW commands for...?

- Backup/archive client
- Server or storage agent

Main Menu | Tracing

Tracing

How do I...?

Enable tracing for...

- Backup/Archive client
- Device driver
- Server or storage agent
- Administration Center

Find trace classes for...

- Backup/Archive client
- Server or storage agent
- Administration Center

Main Menu | Using this guide

Using this guide

Welcome to the *IBM Tivoli Storage Manager Problem Determination Guide*. This guide is intended for use by administrators and service personnel, managing and supporting Tivoli Storage Manager.

The *IBM Tivoli Storage Manager Problem Determination Guide* consists of the following sections:

Information about...
A quick reference section to key topics relating to problem determination. Items discussed in this section are tracing and other diagnostic tools available for Tivoli Storage Manager as well as how to contact IBM to report a problem.

Hints and tips...
Hints and tips for tuning and diagnosing aspects of your Tivoli Storage Manager environment that are external to Tivoli Storage Manager itself.

How do I diagnose...?
Diagnosis and troubleshooting recommendations and tips for Tivoli Storage Manager as well as discussion of common problems that are encountered.

The format of the "How do I diagnose...?" sections is a table with a brief description of the symptom. Click the symptom and it will display complete information about the problem in the form:

Symptom

Steps to diagnose that this is the problem encountered.

Steps to resolve this problem.

Main Menu | Administration Center | Administration Center Installation

Administration Center Installation

If the Integrated Solutions Console Installation failed, verify the following:

1. There is enough disk space on the system. To install the console on a system for the first time requires the following approximate space allocations:
   - Install media location: 505 MB
   - Temporary space: 75 MB
   - Installation location: 650 MB
2. The host system is configured to use a static IP address.
3. The media and installation path names for UNIX platforms do not contain space characters.
4. Full path names are limited to 32 characters for the installation path.

If the Tivoli Storage Manager Administration Center Installation failed, verify the following:

1. The Integrated Solutions Console is installed.
2. The configuration information listed in the *'Review Integrated Solutions Console Configuration Information'* section of the Tivoli Storage Manager Administration Center Installation Wizard is correct.
3. Integrated Solutions Console Server and Console Help Server are running.
   - Starting and stopping the console server and the console help server

Logs

The table describes the files created when installing and uninstalling, and it recommends when and where to check the file for information that might assist in troubleshooting problems.

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

For the Integrated Solutions Console Version 6.0.1, the installation logs are collected in ISCLogs_<date>_<time>.jar; this file is in the directory: <ISC_root>/Tivoli/dsm/logs.

Logs

File Name	Description	Symptoms	Location
ISCRuntimeInstall.log	Contains messages for the Integrated Solutions Console Runtime Installation.	Check this log if the Integrated Solutions Console Runtime installation failed.	This log is located in the OS temporary directory. For Windows, this is the value of the %TEMP% system variable. For all other systems, this is the /tmp directory.
ISCRuntimeUninstall.log	Contains trace output for the runtime uninstall program.	Check this log if the runtime uninstall failed. Look for messages that indicate that a command failed.	This log is located in the OS temporary directory. For Windows, this is the value of the %TEMP% system variable. For all other systems, this is the /tmp directory.
EWASE.rsp	Contains the settings specified for the Integrated Solutions Console Runtime installation.	Check this log and verify that the settings are correct if the installation was not successful.	This log is located in the OS temporary directory. For Windows, this is the value of the %TEMP% system variable. For all other systems, this is the /tmp directory.
ac_install.log	Contains output from the Integrated Solutions Console Runtime installation and Administration Center installation.	Check this log if the Administration Center installation failed or when the ISCRuntimeInstall.log does not exist.	This log is located in the isc_runtime_root/Tivoli/dsm/logs directory.
log.txt	Contains error messages for the ISMP installation wizards for both the Integrated Solutions Console and the Administration Center.	Check this file to determine the value of the exception thrown. The exception error thrown helps determine where the install failed.	This log is located in the installation root directory (isc_runtime_root).

Manually Uninstalling the Integrated Solutions Console

If the install program fails, you can try to manually remove the Integrated Solutions Console and then install again. See the *Tivoli Storage Manager Installation Guide* (http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp) for instructions on manually uninstalling the Integrated Solutions Console.

How to Check for Previous Installation of the Integrated Solutions Console

To check for previous installations of the Integrated Solutions Console, open the system vpd.properties file and search for "PortalServer". It will be preceded by the installation directory for the Integrated Solutions Console. The vpd.properties file is usually located in the root home directory for UNIX platforms and in the WINNT directory for windows.

Configuring the IP Address

The Administration Center assumes that the host system is using a static IP address instead of a dynamically assigned IP address. A static IP address is necessary because the Administration Center server must be listed on the domain name servers, which map the host name to the physical address of the system.

On a system that is not connected to the network, you must configure your system so that the IP loopback port is mapped to the fully-qualified host name. To enable that mapping, perform the following procedure.

1. Go to the system where the Administration Center will be installed.

2. Locate the TCP/IP host file on your system.

- For Windows systems, look in the path WINNT\system32\drivers\etc.
- For AIX, Linux, and Solaris systems, look in the path/etc.3.

Use a text editor to open the host file.

4. At the end of the host file, add lines similar to the following:

- 127.0.0.1 localhost
- 127.0.0.1 your.server.name

where your.server.name is the fully-qualified host name for the Administration Center system.

5. Save the host file.

Main Menu | Message refers to the Problem Determination Guide for more information

Message refers to the Problem Determination Guide for more information

Some messages in the Administration Center state to see the Problem Determination Guide for more information. Select a message from the table below to see a more detailed description about the error message and some steps to try and resolve the error. If the problem still cannot be resolved then see Enabling trace for the Administration Center for instructions on how to gather a trace before contacting customer support. The information about each message is in the form:

MESSAGE

Message number and full message text

Detailed description of some of the things that can cause the message to appear

Follow the steps outlined to try and resolve the messages below:

List of messages with Problem Determination Guide information
• ANRW0089E • ANRW0090E • ANRW0126E • ANRW0127E • ANRW0128E • ANRW0129E • ANRW0130E • ANRW0131E • ANRW0132E • ANRW0209E • ANRW0212E • ANRW0213E • ANRW0216E

- ANRW0221E
- ANRW0257E

ANRW0089E

ANRW0089E The following command generates a confirmation request from the server: "<command>." The program cannot process this request. See the Administration Center section of the Problem Determination Guide for more information.

This error message indicates that a command was issued that was expecting a reply. Some commands require a confirmation. This can be seen on the administrative client as commands that prompt for additional responses such as, "Are you sure you want to delete this?" The API that the Administration Center uses does not have a mechanism to process confirmation requests.

Try issuing the same command using the command line.

Return to list of messages

ANRW0090E

ANRW0090E The following command is unknown to the server: "<command>". See the Administration Center section of the Problem Determination Guide for more information.

This error message indicates that an unknown command was issued to the server. This could be caused by a variety of reasons. The command may not exist on the version of server the command is being issued on. It may be an unsupported command on the platform the server is running on. The command may not have been constructed correctly.

Visually inspect the command that was attempted. If the command is run together such as, `QUERYNODE` instead of `QUERY NODE`, this indicates the command was not constructed correctly. Attempt to reconstruct the command and issue it using the command line.

Return to list of messages

ANRW0126E

ANRW0126E The administrative API encountered an internal error while processing the request. See the Administration Center section of the Problem Determination Guide for more information.

This error message indicates that the administrative API that the Administration Center uses had an unrecoverable error while attempting to issue a command to the server.

Try issuing the same command using the command line.

Return to list of messages

ANRW0127E

ANRW0127E The administrative API is unable to process the command document sent from the server. See the Administration Center section of the Problem Determination Guide for more information.

The Administration Center transfers a command document file from the server. This XML file indicates all the valid commands and their syntax for each version and platform of the server. This file is only transferred once during the first connection with a server of a particular version and platform. The error reports that the file could not be processed and may be corrupt. The file is located in the ISC install directory under `PortalServer\shared\app\dsm\CmdFileCache`. The name of the file depends on the version and platform. For example a 5.3.0 server for AIX would have the name `aix_5_3_0.xml`.

Since the Administration Center is having trouble reading this file, it may be corrupt. Navigate to the `CmdFileCache` directory. Since all the files in this directory are only cached files, it is safe to delete them. Delete the file in question, or delete all of the files in the cache directory. This will force the Administration Center to pull the file from each server. Also check the permissions for this directory and make sure that the Administration Center can read and write files in this directory.

Return to list of messages

ANRW0128E

ANRW0128E The administrative API is unable to process the command because it cannot be found in the command document. See the Administration Center section of the Problem Determination Guide for more information.

The Administration Center transfers a command document file from the server. This XML file indicates all the valid commands and their syntax for each version and platform of the server. This file is only transferred once during the first connection with a server of a particular version and platform. The error reports that the command in question could not be found in the command file. This indicates the command file may be corrupt or out of date (wrong version). The file is located in the ISC install directory under `PortalServer\shared\app\dsm\CmdFileCache`. The name of the file depends on the version and platform. For example a 5.3.0 server for AIX would have the name `aix_5_3_0.xml`.

Since the Administration Center is are having trouble reading this file, it may be corrupt. Navigate to the `CmdFileCache` directory. Since all the files in this directory are only cached files, it is safe to delete them. Delete the file in question, or delete all of the files in the cache directory. This will force the Administration Center to pull the file from each server. Also check the permissions for this directory and make sure that the Administration Center can read and write files in this directory.

Return to list of messages

ANRW0129E

ANRW0129E The following command contains one or more invalid parameters: "<command>." See the Administration Center section of the Problem Determination Guide for more information.

This error message indicates that the Administration Center was attempting to execute a command on a server but the call to the API contains one or more invalid parameters. This type of error should typically be caught during validation before the command is submitted to the API for execution.

Check the parameters in the command. If you are filling out text entry fields on a page, you may be able to find the error in the parameters and correct it. Viewing the activity log in the server's properties notebook may also help determine the problem.

Return to list of messages

ANRW0130E

ANRW0130E The administrative API encountered invalid parameters while processing the request. See the Administration Center section of the Problem Determination Guide for more information.

This error message indicates that a command was being executed through the administrative API but one of the parameters to an API method was invalid.

Although this is typically an internal error, it may be caused by unusual parameters while constructing the command. Check the parameters in the command. If you are filling out text entry fields on a page, you may be able to find the error in the parameters and correct it. If you are entering a value with unusual characters such as: <. >, & this may be the source of the problem. Viewing the activity log in the server's properties notebook may also help determine the problem.

Return to list of messages

ANRW0131E

ANRW0131E The administrator's authority level on this server cannot be determined. See the Administration Center section of the Problem Determination Guide for more information.

This error message indicates that the authority level for the administrator used for connecting to this server cannot be determined.

Try using a different administrator id. Viewing the activity log in the server's properties notebook may also help determine the problem. To use a different administrator id, use the Modify Server Connection table action. Then specify a different administrator id.

Return to list of messages

ANRW0132E

ANRW0132E The administrative API cannot locate the command document. See the Administration Center section of the Problem Determination Guide for more information.

The Administration Center transfers a command document file from the server. This XML file indicates all the valid commands and their syntax for each version and platform of the server. This file is only transferred once during the first connection with a server of a particular version and platform. The error reports that this file could not be located. The file is located in the ISC install directory under `PortalServer\shared\app\dsm\CmdFileCache`. The name of the file depends on the version and platform. For example a 5.3.0 server for AIX would have the name `aix_5_3_0.xml`.

Check the directory permissions for the `CmdFileCache` directory. Administration Center must have permission to read and write to this directory.

Return to list of messages

ANRW0209E

ANRW0209E An unknown failure occurred while the health monitor was evaluating the server. See the Administration Center section of the Problem Determination Guide for more information.

The health monitor encountered an error that it could not recover from.

Not much can be done in this case except to verify the health monitor is set up correctly. You also may want to try to stop and restart ISC. The only way the health monitor thread can come back is by restarting ISC in some cases.

Return to list of messages

ANRW0212E

ANRW0212E The health monitor cannot evaluate the health of the server. Correcting problems that are indicated in the Detailed Server Health information might solve the problem. See the Administration Center section of the Problem Determination Guide for more information.

The health monitor encountered an error while running reports on the servers. This error can occur if file specifying how to run the health monitor reports is modified or corrupted.

Verify the health monitor reports file is not modified. This file is located in the ISC install directory under `PortalServer\shared\app\dsm\reports\dsmhealthmon.xml`. If this file has been modified or corrupted, replace this file with a valid version of the file. This can be done by installing the Administration Center again.

Return to list of messages

ANRW0213E

ANRW0213E The health monitor is unable to locate information about the health of the server. See the Administration Center section of the Problem Determination Guide for more information.

This error can occur when a server has recently been removed from the Administration Center, but the health monitor still thinks the server is there. This can happen if the health monitor and the rest of the Administration Center become out of sync on which servers are defined. It can also occur if the report files generated by the health monitor cannot be found.

If the health monitor and the rest of the Administration Center become out of sync then restarting ISC will help. Verify the health monitor report directory exists and has read and write privileges. This directory is located in the ISC install directory under `PortalServer\shared\app\dsm\reports`. You may also want to try to run the new **resync** health monitor password (action from the Health Monitor table of servers entry screen). This password tries to resolve any credential issues with the system.

Return to list of messages

ANRW0216E

ANRW0216E The health monitor is unable to locate the results of its evaluation of the health of a server. See the Administration Center section of the Problem Determination Guide for more information.

This error indicates there was a problem running the health monitor reports. This error can happen if the health monitor reports file cannot be found or if the health monitor reports directory doesn't exist or does not have enough privileges.

Verify the health monitor reports directory exists and has read and write permissions. This directory is located in the ISC install directory under `PortalServer\shared\app\dsm\reports`. Verify the health monitor reports file exists. This file is located in the ISC install directory under `PortalServer\shared\app\dsm\reports\dsmhealthmon.xml`.

Return to list of messages

ANRW0221E

ANRW0221E The health monitor reports an internal error. See the Administration Center section of the Problem Determination Guide for more information.

The health monitor encountered an internal error that it could not recover from.

Not much can be done in this case except to verify the health monitor is set up correctly. You may also want to try to restart ISC when the health monitor setup is correct.

Return to list of messages

ANRW0257E

ANRW0257E The status information that is displayed is not valid because the health monitor service is unavailable. See the Administration Center section of the Problem Determination Guide for more information.

This error occurs when the health monitor service cannot be found. This can happen when the health monitor service is not set up correctly.

Verify the health monitor service is set up correctly.

Return to list of messages

Main Menu | Performance Tuning Suggestions

Performance Tuning Suggestions

How do I size the number of Tivoli Storage Manager Administration Center installations?

This section includes information for sizing the number of Tivoli Storage Manager Administration Center installations, including the Integrated Solutions Console (ISC), which may be required to support a given number of administrators and Tivoli Storage Manager servers in a large enterprise.

While there are a number of variables that can influence this sizing, most of these variables can be safely ignored if the Administration Center installation meets the recommended installation minimums and the following guidelines:

1. The Administration Center is being installed on a system dedicated to this function and no other applications will be installed on this system that require significant processing or memory resources.
2. The Administration Center is being installed on a system that is based on recent hardware (not more than 18 months has elapsed since general availability of the specific hardware models).

If these conditions are met,then it is very likely that the system will have sufficient processing capacity. Uni-processor systems or dual-processor systems are expected to be sufficient to meet the processing needs for a single Administration Center installation. More processors are not expected to be beneficial, as the major resource demand will be for memory. This is explained below.

The Administration Center disk I/O and network I/O requirements are not particularly demanding and there is no need for sustained high I/O throughput. However, application response time will suffer if network delays or disk I/O delays occur. A low latency network will provide the best administrator response time. Because more of the network traffic during administration activities will occur between the ISC system and the administrator's browser than between the ISC system and a Tivoli Storage Manager server, the Integrated Solutions Console and Administration Center should be installed close to the administrators, rather than close to the Tivoli Storage Manager servers.

How can I best optimize performance of the Administration Center?

The main resource requirement for supporting a significant number of Tivoli Storage Manager administrators in the Administration Center is memory. Performance testing has shown that the Administration Center process working set memory can be determined using the following equation:

$$M_{WorkingSet} = 100 + M_{JavaHeapSize}$$

where:

$M_{WorkingSet}$ is the Administration Center process working set memory in MB.

$M_{JavaHeapSize}$ is the Administration Center Java heap size memory in MB.

This working set memory must remain less than 2 GB, which is the maximum supported process virtual memory for many 32-bit operating systems. This is not likely to be a problem, since the maximum Java heap size allowed is 1536 MB. It is very important that the required working set memory be available in real memory (RAM) or significant response time degradation may occur as the result of system memory paging. Thus, the process working set memory requirement is effectively determined by the amount of Java heap memory required.

The system memory sizing must include memory for the operating system and any other applications executing on the system. Most operating systems require 256 MB, so this much memory should be added to the process working set memory requirement to obtain the system memory requirement.

See the Administration Center Support Utility section in this guide for more tuning suggestions.

Where do I go for more information on Performance Tuning?

For Tivoli Storage Manager Administration Center performance information, search the knowledge base for Tivoli Storage Manager from the ibm.com support website http://www.ibm.com/support. Use the search term 1193443.

Main Menu | Blank Pages in Administration Center

Why are there Blank Pages after logging onto the Administration Center?

After installing the Integrated Solutions Console (ISC) and Administration Center, when logging into the ISC, the Administration Center work items on the left show up, but when clicked on, the pages are blank. This issue has been reported by customers that are using equipment at or close to the minimum hardware requirements for this product or on a system under heavy load.

Error seen in <isc_home>\PortalServer\log\SystemOut.log:

Exception: com.ibm.websphere.management.exception.DocumentNotFoundException for admincenter<some_number>.ear

Errors seen in <isc_home>/Tivoli/dsm/logs/ac_install.log:

[31 August 2005 08:29:20 GMT+00:00] **** Precompile Startup ****
[31 August 2005 08:29:20 GMT+00:00] WPS script file [/opt/IBM/ISC/AppServer/bin/wpsSetupCmd.sh] exists - UNIX
[31 August 2005 08:29:20 GMT+00:00] ERROR: AdminCenter failed to deploy, unable to locate enterprise name.
[31 August 2005 08:29:20 GMT+00:00] **** Precompile Failed ****
[31 August 2005 08:29:20 GMT+00:00] Completed step number: 10
[31 August 2005 08:29:20 GMT+00:00] WARNING: Failed to precompile the JSP pages. There will be a performance hit while running.

Key:

<isc_home> = The root directory in which the ISC was installed (default for Windows: C:\Program Files\IBM\ISC Default for UNIX: /opt/IBM/ISC)
<userID> = The user ID used to install the ISC (default is iscadmin).
<password> = The password used to install the ISC.
localhost:8421 = The host name or IP address of where the ISC was installed and the port that was used (8421 is default).
<fully qualified path> = This is the fully qualified path to the AdminCenter.war file that should be located where the Admin Center package was un-zipped or un-tared.

Workaround for Windows:

1. Make sure the ISC server is running. This can be verified by looking in Windows Services for 'IBM WebSphere Application Server V5 - ISC Runtime Service'
2. Open a command prompt and go to the following directory: <isc_home>\PortalServer\bin
3. Enter the following: iscremove com.tivoli.dsm.admincenter <userID> <password> http://localhost:8421/deploy/deploy
4. This should return with a '0' when successfully removed.
5. Cd to <isc_home>\PortalServer\shared\app\config\services
6. Open the following file: DeploymentService.properties
7. Change *was.notification.timeout = 60* to *was.notification.timeout = 300*
8. Stop and restart the ISC Service by completing the following steps:

    Cd to <isc_home>\PortalServer\bin
    Enter: stopISC.bat ISC_Portal
    Enter: startISC.bat ISC_Portal

9. Enter the following: iscdeploy <fully qualified path>\AdminCenter.war <userID> <password> http://localhost:8421/deploy/deploy
10. This should return with a '0' when successfully deployed.

Workaround for UNIX (AIX, Solaris, and Linux):

1. Make sure the ISC server is running. This can be verified by going to <isc_home>/AppServer/bin and issuing: serverStatus.sh ISC_Portal
2. Open a command prompt and go to the following directory: <isc_home>/PortalServer/bin
3. Enter the following: ./iscremove.sh com.tivoli.dsm.admincenter <userID> <password> http://localhost:8421/deploy/deploy
4. This should return with a '0' when successfully removed.
5. Cd to <isc_home>/PortalServer/shared/app/config/services
6. Open the file: DeploymentService.properties
7. Change *was.notification.timeout = 60* to *was.notification.timeout = 300*
8. Stop and restart the ISC Service by completing the following steps:

    Cd to <isc_home>/PortalServer/bin
    Enter: ./stopISC.sh ISC_Portal
    Enter: ./startISC.sh ISC_Portal

9. Enter the following: /iscdeploy.sh <fully qualified path>/AdminCenter.war <userID> <password> http://localhost:8421/deploy/deploy
10. This should return with a '0' when successfully deployed.

Main Menu | Unable to access Server from a Web browser

Unable to access Server from a Web browser

A problem is encountered when attempting to access Integrated Solutions Console (ISC) and the Tivoli Storage Manager Administration Center from a Web browser. The following are examples of the messages that could be received:

- The page cannot be displayed.
- The page you are looking for is currently unavailable.
- Error: Cannot find server or DNS Error.
- The connection was refused when attempting to contact host:port.

These types of errors can occur for a number of reasons. The following is a list of ways to correct these errors.

- Ensure that system on which the browser is running has connectivity to the network and the target system on which ISC is installed.
- Ensure that the browser settings are correct.
- If firewall, Internet security or proxy software is being used, try temporarily disabling it. If ISC is now accessible, the firewall settings should be investigated.
- Ensure that ISC is installed on the target system.
- Ensure that ISC is running on the target system. ISC may have been terminated either normally or abnormally.
- Ensure that the Tivoli Storage Manager Administration Center has been deployed into ISC.

ISC is actually composed of two servers, the console server and the console help server. If the system on which ISC is installed is shutdown, any ISC servers that were started will be stopped.

If the system was configured to automatically start the console server and console help server when the system restarted, you should not have to take any action.

The following section describes the commands that can be used to manually start and stop the servers.

The following table lists the commands for starting and stopping the servers on AIX, Linux, and Solaris systems.

Servers	Start command (on a single line)	Stop command (on a single line)
Console server and console help server	*your_isc_root* /PortalServer/bin/startISC.sh ISC_Portal	*your_isc_root* /PortalServer/bin/stopISC.sh ISC_Portal *administrator password*

Console server only	*your_isc_root/* AppServer/bin/startServer.sh ISC_Portal -username *administrator* -password *password*	*your_isc_root/* AppServer/bin/stopServer.sh ISC_Portal -username *administrator* -password *password*
Console help server only	*your_isc_root/* PortalServer/ISCEclipse/ StartEclipse.sh	*your_isc_root* /PortalServer/ISCEclipse/ StopEclipse.sh
Note: In the commands, the variables have the following meanings: • *your_isc_root* is the root directory for your Integrated Solutions Console installation. • *administrator* is the administrator user ID for Integrated Solutions Console if the installation uses the WebSphere Application Server embedded in the console runtime. If the installation uses a separate WebSphere Application Server installation, *administrator* is the administrator user ID for the WebSphere Application Server installation. • *password* is the password for the administrator.		

The following table lists the commands for starting and stopping the servers on Windows systems.

Servers	Start command (on a single line)	Stop command (on a single line)
Console server and console help server	*your_isc_root* \PortalServer\bin\startISC.bat ISC_Portal	*your_isc_root* \PortalServer\bin\stopISC.bat ISC_Portal *administrator password*
Console server only	*your_isc_root* \AppServer\bin\startServer.bat ISC_Portal -username *administrator* -password *password*	*your_isc_root* \AppServer\bin\stopServer.bat ISC_Portal -username *administrator* -password *password*
Console help server only	*your_isc_root* \PortalServer\ISCEclipse \StartEclipse.bat	*your_isc_root* \PortalServer\ISCEclipse \StopEclipse.bat
Note: In the commands, the variables have the following meanings: • *your_isc_root* is the root directory for your Integrated Solutions Console installation • *administrator* is the administrator user ID for Integrated Solutions Console if the installation uses the WebSphere Application Server embedded in the console runtime. If the installation uses a separate WebSphere Application Server installation, *administrator* is the administrator user ID for the WebSphere Application Server installation. • *password* is the password for the administrator.		

Main Menu | ODBC driver | Configuration details.

Configuration details

Installed files

The Tivoli Storage Manager ODBC driver files are located in the directory specified during installation. The default installation directory is:

(32-bit Windows) `x:\Program Files\Tivoli\TSM\odbc`

(64-bit Windows) `x:\Program Files (x86)\Tivoli\TSM\odbc`

where 'x:' is the drive letter where the operating system is installed (usually C:).

The installed files include:

- dscenu.txt
- dsmntapi.dll
- odbcadsm.dll
- tsmutil1.dll

One or more of the following files may also be installed, depending on the language version of Microsoft(R) Windows(R) that you are running and any additional language support that you specified during installation.

dscchs.txt
: Chinese (Simplified)

dsccht.txt
: Chinese (Traditional)

dscdeu.txt
: German

dscesp.txt
: Spanish

dscfra.txt
: French

dscita.txt
: Italian

dscjpn.txt
: Japanese

dsckor.txt
: Korean

dscptb.txt
: Portuguese (Brazilian)

Example 1: If your machine is running the German version of Windows, then the dscdeu.txt file will be installed by default.

Example 2: If your machine is running the German version of Windows and you also select French and Italian language support, then dscdeu.txt, dscfra.txt, and dscita.txt will be installed.

Example 3: If your machine is running the English version of Windows and you also select Spanish and Japanese support, then dscesp.txt and dscjpn.txt will be installed.

NOTE: Prior versions of the Tivoli Storage Manager ODBC driver placed its files in the winnt\system32 directory. The following files may be deleted from that directory if they still exist:

- dsgameng.txt
- odbcadsm.dll
- odbcadsm.txt

Environment Variables

Starting with Version 5.3 of the Tivoli Storage Manager ODBC driver, the environment variables DSMG_DIR and DSMODBC_DIR are no longer supported and must not be used. The Tivoli Storage Manager ODBC driver installation program will attempt to remove DSMG_DIR and DSMODBC_DIR if they exist on your system. However, if they still exist after the installation, then they should be removed manually.

Registry settings

IMPORTANT: Do NOT change any of these values unless it is at the direction of IBM Development or Service.

The Tivoli Storage Manager ODBC driver installation settings are located in subkeys of the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\IBM\ADSM\CurrentVersion
```

The following subkeys and associated values are pertinent to the Tivoli Storage Manager ODBC driver installation:

```
ComponentsMsi
```

Name	Explanation
`ODBC Driver`	A REG_SZ value that must be `Odbc`

`Odbc`

Name	Explanation
Path	The directory in which you installed the Tivoli Storage Manager ODBC driver. This is the parent directory for the **odbc** subdirectory that contains the odbcadsm.dll file. For example, if odbcadsm.dll is in directory `C:\Program Files\Tivoli\TSM\odbc`, then this value should be `C:\Program Files\Tivoli\TSM`. This is a REG_SZ value.
`PtfLevel`	A REG_SZ value indicating the 4-node version number of the Tivoli Storage Manager ODBC driver. For example, if the Tivoli Storage Manager 5.1.5.0 ODBC driver is installed, then this value should be "5.1.5.0" (without the quotation marks).

`Odbc\nls\xx_XX`

xx_XX is a language abbreviation as follows (with the language name shown to the right for your reference):

de_DE
: German

en_US
: U.S. English

es_ES
: Spanish

fr_FR
: French

it_IT
: Italian

ja_JP
: Japanese

ko_KR
: Korean

pt_BR
: Portuguese (Brazilian)

zh_CN
: Chinese (Simplified)

zh_TW
: Chinese (Traditional)

You must always have the `Odbc\nls\en_US` subkey.

If you installed the Tivoli Storage Manager ODBC driver on a non-English Windows system, then you may also have a subkey that corresponds to the language version of Windows that you are running. For example, if you are running the German language version of Windows, then you may also have the `Odbc\nls\de_DE` subkey.

You may have other subkeys as well if, during the installation, you selected multiple language support.

Name	Explanation
`FileName`	A REG_SZ value indicating the name of the Tivoli Storage Manager message repository that corresponds to the subkey's language setting. For example, in subkey `Odbc\nls\de_DE`, the `FileName` value should be `dscdeu.txt`.
`Language`	A REG_SZ value indicating the language of the message repository specified in the `FileName` value. For example, if `FileName` is `dscdeu.txt`, then the Language setting must be `German`.

The following summarizes the `Language` value settings:

Subkey	File name	Language
`Odbc\nls\de_DE`	`dscdeu.txt`	`German`
`Odbc\nls\en_US`	`dscenu.txt`	`English (US)`
`Odbc\nls\es_ES`	`dscesp.txt`	`Spanish`
`Odbc\nls\fr_FR`	`dscfra.txt`	`French`
`Odbc\nls\it_IT`	`dscita.txt`	`Italian`
`Odbc\nls\ja_JP`	`dscjpn.txt`	`Japanese`
`Odbc\nls\ko_KR`	`dsckor.txt`	`Korean`
`Odbc\nls\pt_BR`	`dscptb.txt`	`Portuguese (Brazilian)`
`Odbc\nls\zh_CN`	`dscchs.txt`	`Chinese (Simplified)`
`Odbc\nls\zh_TW`	`dsccht.txt`	`Chinese (Traditional)`

The Tivoli Storage Manager ODBC driver settings are located in subkeys of the following registry key:

(32-bit Windows) `HKEY_LOCAL_MACHINE\SOFTWARE\ODBC`

(64-bit Windows) `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ODBC`

The following subkeys and associated values are pertinent to the Tivoli Storage Manager ODBC driver (these are all located under `HKEY_LOCAL_MACHINE\SOFTWARE\ODBC`):

`ODBC.INI\ODBCADSM`

Name	Explanation

ADSMTrace	A REG_SZ value that indicates whether Tivoli Storage Manager ODBC driver tracing is enabled. Values may be one of: No Yes If tracing is enabled (ADSMTrace is set to Yes), then trace output will be written to the file specified by the TraceFileName value (see below).
Language	A REG_SZ value that indicates which language version of the Tivoli Storage Manager message repository is in use. Values may be one of the following: German English (US) Spanish French Italian Japanese Korean Portuguese (Brazilian) Chinese (Simplified) Chinese (Traditional)
TraceFileName	A REG_SZ value that indicates the name of the Tivoli Storage Manager ODBC driver trace file.

ODBCINST.INI\ODBC Drivers

Name	Explanation
Tivoli Storage Manager ODBC Driver	A REG_SZ value indicating that the Tivoli Storage Manager ODBC driver is installed. This value is always set to Installed.

ODBCINST.INI\TSM ODBC Driver

Name	Explanation
APILevel	A REG_SZ value that is always set to 0 (zero). This indicates that the Tivoli Storage Manager ODBC driver implements ODBC Core-level functions only.

`ConnectFunctions`	A REG_SZ value that is always set to `YYN`. This string indicates whether the driver implements the SQLConnect(), SQLDriverConnect(), and SQLBrowseConnect() functions, respectively. 'Y' (Yes) means that the function is implemented in the driver; 'N' (No) means that the function is not implemented in the driver. Thus, the Tivoli Storage Manager ODBC driver currently implements SQLConnect() and SQLDriverConnect(); the driver does not implement SQLBrowseConnect().
`Driver`	A REG_SZ value that specifies the full path to the Tivoli Storage Manager ODBC driver. Thus if the driver is installed to `C:\Program Files\Tivoli\TSM\odbc`, the value will be: `C:\Program Files\Tivoli\TSM\odbc\odbcadsm.dll`
`DriverODBCVer`	A REG_SZ value that indicates which version of the ODBC specification the driver supports. This value is set to `03.51`.
`FileUsage`	A REG_SZ value that indicates how the driver treats files in a data source. Because the Tivoli Storage Manager ODBC driver is not a file-based driver, this value is always `0` (zero).
`Setup`	A REG_SZ value that specifies the full path to the driver's Data Source configuration program. The Tivoli Storage Manager ODBC driver's configuration program is in the Tivoli Storage Manager ODBC driver itself, so this value should always match that of the `Driver` setting (see above).
`SQLLevel`	A REG_SZ value that indicates the SQL-92 grammar supported by the driver. Because Tivoli Storage Manager supports a subset of the SQL-92 entry-level grammar, this value is always `0` (zero).
`UsageCount`	A REG_DWORD value that indicates how many components use this driver. Because the driver is used by only one component (the Tivoli Storage Manager ODBC driver itself), this value should be `1`.

The Tivoli Storage Manager ODBC driver data source settings are located in the following registry key:

`HKEY_CURRENT_USER\Software\ODBC\ODBC.INI\<`*data source name*`>`

where <*data source name*> is a user-defined data source name. Note that the Tivoli Storage Manager ODBC driver installation program configures a sample data source called **Tivoli Storage Manager ODBC Sample User DSN**.

Each instance of this subkey contains the following values:

Name	**Explanation**
`CommMethod`	A REG_SZ value indicating the communications method that the Tivoli Storage Manager ODBC driver uses to communicate with the Tivoli Storage Manager server. Because the only supported communications protocol is TCP/IP, this value is always set to `TCP/IP`.

`Driver`	A REG_SZ value that specifies the full path to the Tivoli Storage Manager ODBC driver. Thus if the driver is installed to `C:\Program Files\Tivoli\TSM\odbc`, the value will be: `C:\Program Files\Tivoli\TSM\odbc\odbcadsm.dll` NOTE: This value always matches that of the Driver value specified in registry subkey `HKEY_LOCAL_MACHINE\SOFTWARE\ODBC\ODBCINST.INI\TSM ODBC Driver`.
`DsnDescription`	A REG_SZ value containing a user-specified free-form comment (or description) for the data source definition.
`FetchSize`	A REG_SZ value indicating the maximum number of rows that the ODBC driver requests (buffers) from the Tivoli Storage Manager server at a time. Note that this is not the same as the SQL_ATTR_ROW_ARRAY_SIZE ODBC statement attribute. The former governs the number of records that the ODBC driver will buffer from the server, while the latter governs the number of records that the application will request from the ODBC driver. The default setting is `50` (50 records). Larger settings may improve performance.
`TcpPort`	A REG_SZ value indicating the TCP/IP communications port that the Tivoli Storage Manager ODBC driver uses to connect to the Tivoli Storage Manager server.
`TcpSrvAddr`	A REG_SZ value indicating the TCP/IP address that the Tivoli Storage Manager ODBC driver uses to connect to the Tivoli Storage Manager server.
`UID`	A REG_SZ value indicating the name of the Tivoli Storage Manager Administrator ID used to connect to the Tivoli Storage Manager server.

ODBC .INI Files

The operating system records information from ODBC registry keys in two files located in the `x:\winnt` directory (x:\windows for Windows XP users):

- `ODBC.INI`
- `ODBCINST.INI`

IMPORTANT: Do NOT rename, delete, or otherwise modify these files or their contents unless it is at the direction of IBM Development or Service.

Main Menu | ODBC driver | Troubleshooting and diagnostics

Troubleshooting and diagnostics.

Before proceeding with the items described below, be sure to check Known problems and limitations to see if your problem is already known. If your problem is not described in that section, then continue with this section.

Installation problems.

The Microsoft(R) Software Installer (MSI) is used to install the Tivoli Storage Manager ODBC driver. If you are having trouble installing the Tivoli Storage Manager ODBC driver, you should contact IBM/Tivoli support and provide the following information:

- Operating system version and service pack number (for example, Windows(R) 2000 SP 3, Windows XP SP 1a)
- Exact version/release/level/patch of the Tivoli Storage Manager ODBC driver that you are trying to install (for example, 5.2.0.0). The version information can be found at the top of this README file.
- Step-by-step, re-create the scenario, including the last dialog box that you see before the failure occurs, and the exact text of any error messages that you see.
- Examine the system Event Log and report any event log errors that coincide with the time that the installation failure occurred.
- An MSI setup.log file. To create this file, run setup.exe again as follows:

  setup /v"/l*v setup.log"

  Follow the prompts as usual to recreate the failure. After the setup program ends, you will have a file called setup.log which contains detailed diagnostic information about the setup.

Data source configuration problems.

If you are having trouble creating or configuring a data source for the Tivoli Storage Manager ODBC driver, see ODBC configuration details and compare your system settings with those described in that section. If you do find any differences, and if you do not know how to repair them, or are uncomfortable doing so, then contact IBM support and provide the information described below the next paragraph (you do not need to provide trace data at this point).

If the configuration looks correct, attempt to obtain a trace of the Tivoli Storage Manager ODBC driver. See Tivoli Storage Manager ODBC Driver Tracing for instructions on tracing. After you have the trace data, contact IBM support and provide the following information:

- Operating system version and service pack number (for example, Windows 2000 Server SP 3, Windows XP SP 1)
- Exact version/release/level/patch of the Tivoli Storage Manager ODBC driver that you are trying to install (for example 5.2.0.0).
- Step-by-step recreation of the scenario, including the last dialog box that you see before the failure occurs, and the exact text of any error messages that you see.
- Your ODBC configuration information. Include the following items:

  > X:\WINNT\ODBC.INI

  > X:\WINNT\ODBCINST.INI

  > The odbchkcu.reg and odbchklm.reg files created by the following commands:

  regedit -e odbchkcu.reg hkey_current_user\software\odbc
  regedit -e odbchklm.reg hkey_local_machine\software\odbc

  > A directory listing of the files in the directory where the Tivoli Storage Manager ODBC driver is installed. For example, if the driver is installed to directory C:\Program Files\Tivoli\TSM\odbc, then open an OS command prompt, change into that directory, and issue the following command:

  dir > odbcdir.out

  Then you can send in the odbcdir.out file.

  > Trace data.

Application problems.

If you have been able to install the Tivoli Storage Manager ODBC driver and configure a data source, but you are having trouble getting your application to work, then do the following:

- Obtain a trace as described in ODBC tracing.
- Contact IBM support and provide the following information:
  - Operating system version and service pack number (for example, Windows 2000 Server SP 3, Windows XP SP 1)
  - Exact version/release/level/patch of the Tivoli Storage Manager ODBC driver that you are trying to install (for example, 5.2.0.0).
  - Name and exact version/service pack of the ODBC application that you are using.
  - Step-by-step, re-create the scenario, including the last dialog box that you see before the failure occurs, and the exact text of any error messages that you see.
  - Your ODBC configuration information. Include the following items:
    - X:\WINNT\ODBC.INI
    - X:\WINNT\ODBCINST.INI
    - The odbchkcu.reg and odbchklm.reg files created by the following commands:

      regedit -e odbchkcu.reg hkey_current_user\software\odbc
      regedit -e odbchklm.reg hkey_local_machine\software\odbc

- A directory listing of the files in the directory where the Tivoli Storage Manager ODBC driver is installed. For example, if the driver is installed to directory C:\Program Files\Tivoli\TSM\odbc, then open an OS command prompt, change into that directory, and issue the following command:

  ```
  dir > odbcdir.out
  ```

  Then you can send in the odbcdir.out file.

- Trace data.

Main Menu | ODBC driver | Defining data sources.

Defining data sources.

The ODBC Data Source Administrator in the Control Panel can be used to add, remove, and configure data source entries. ODBC applications can then be used to connect to the Tivoli Storage Manager server. Note for 64-bit Windows(R) XP and 2003 users: The Tivoli Storage Manager ODBC driver is only available in a 32-bit version. Therefore you must use the 32-bit version of the ODBC Data Source Administrator (odbcad32.exe) to configure Tivoli Storage Manager ODBC driver DSNs. This executable is located in directory *x*:\WINDOWS\SysWOW64, where *x*: is the drive letter where Windows is installed (usually C:).

The Tivoli Storage Manager ODBC setup program will create a default user data source name (DSN) called "TSM ODBC Sample User DSN". You can either update this entry, or remove it and create one or more new DSNs to connect to your Tivoli Storage Manager server or servers. When creating new DSNs, the driver name to use is "Tivoli Storage Manager ODBC Driver".

The ODBC data source setup panel for Tivoli Storage Manager has the following controls:

- Data source name: (REQUIRED) Specify the name that you want to use for the DSN. This is the name that will appear in the "Name" column when you run the ODBC Data Source Administrator. The maximum allowed length is255 characters.
- Description: (OPTIONAL) This is a free-form text field where you may enter a comment or description about this DSN. The maximum allowed length is 255 characters.
- Language: (REQUIRED) Select the language in which ANSxxxxx messages should be displayed. The available languages are those that were selected at the time the ODBC driver was installed. Note: Only the ANSxxxxx messages are available in other languages. The ODBC driver dialogs and other non-ANSxxxxx text strings are currently displayed in English only.
- Administrator name: (OPTIONAL) Specify the Tivoli Storage Manager Administrator ID to use with this DSN.
- Fetch size: (REQUIRED) Specify the number of records that the ODBC driver requests from the Tivoli Storage Manager server per fetch request. Larger values may improve performance, especially when fetching a large number of records. This value must be between 1 and 1000000 (inclusive). The default value is 50.
- TCP/IP address: (REQUIRED) Specify the TCP/IP address of the Tivoli Storage Manager server that you are connecting to. If you do

not know which address to use, see your Tivoli Storage Manager administrator.

- TCP/IP port: (REQUIRED) Specify the TCP/IP port number of the Tivoli Storage Manager server that you are connecting to. If you do not know which port number to use, see your Tivoli Storage Manager administrator.
- Enable trace: (OPTIONAL) Selecting this check box activates Tivoli Storage Manager ODBC driver tracing. Clearing this check box deactivates tracing. Normally this check box is not selected. Tracing should only be performed at the request of Tivoli Storage Manager Service or Development.
- Trace file name: (OPTIONAL) This field is available only when "Enable trace" has been selected. Here you can specify a fully-qualified file name to which trace data is written. Tracing should only be performed at the request of Tivoli Storage Manager Service or Development.

When creating file DSNs, you can use the following keywords:

DRIVER=*driverName*
: *driverName* must be Tivoli Storage Manager ODBC Driver.

UID=*administratorName*
: *administratorName* is a valid Tivoli Storage Manager Administrator ID.

TCPADDR=*tcpipAddress*
: *tcpipAddress* is the TCP/IP address for your Tivoli Storage Manager server.

TCPPORT=*tcpipPortNumber*
: *tcpipPortNumber* is the TCP/IP port number used to connect to your Tivoli Storage Manager server.

FETCHSIZE=*fetchSize*
: *fetchSize* is the number of records that the ODBC driver requests from the Tivoli Storage Manager server per fetch request. This value must be between 1 and 1000000 (inclusive). The default value is 50.

For example: if your administrator ID is storman, your Tivoli Storage Manager server address is storman.tucson.ibm.com, and your Tivoli Storage Manager TCP/IP port number is 1500, then the file DSN should look like this:

```
[ODBC]
DRIVER=Tivoli Storage Manager ODBC Driver
UID=storman
TCPADDR=storman.tucson.ibm.com
TCPPORT=1500
```

Note that UID, TCPADDR, and TCPPORT are not required; if they are not specified, then you will be prompted for their values. The [ODBC] line is added for you automatically when you use the ODBC Data Source Administrator to create the file DSN.

To create a file DSN using the ODBC Data Source Administrator:

1. Start the ODBC Data Source Administrator from the Control Panel. Alternatively you can run odbcad32.exe.
2. Click the File DSN tab, and then click **Add**.
3. Select Tivoli Storage Manager ODBC Driver (located in the Name column). Then click **Next**.
4. Enter the name that you want to use for this DSN. For example, Tivoli Storage Manager File DSN. Click **Next**.
5. The dialog now shows you a summary of the options chosen thus far. Click **Finish**.
6. You should now be presented with the Connect to Tivoli Storage Manager Server dialog. Fill in the TCP/IP address, TCP/IP port, administrator ID, and password, and then click **Connect**. An attempt will be made to connect to the Tivoli Storage Manager server using the information that you supplied above. If the connection is successful, you will be returned to the ODBC Data Source Administrator dialog. If there was a problem connecting to the Tivoli Storage Manager server, then you will be prompted again to enter the correct information.

Note: For security purposes, the password is NOT saved. It is used in this step for the purpose of validating the connection information.

Main Menu | ODBC driver | Known problems and limitations.

Known problems and limitations.

The following are the known problems or limitations with the Tivoli Storage Manager ODBC driver:

- SELECT statements containing the AS keyword in the column specification with a quoted column name will not work. For example, the following will not work:

  select node_name as "Node Name" from nodes

  Instead, use single word column names, which do not require quotes:

  select node_name as NODENAME from nodes

- When using Microsoft Excel or Microsoft Query (msqry32.exe), statements containing the AS keyword in the column specification will not work at all. This is because Microsoft Query (which is what Excel uses to import the data) always puts single quotes around the new column name, even if it does not contain blank spaces; and the Tivoli Storage Manager server does not accept single quotes around the column name. Microsoft Knowledge Base article 298955 describes this issue.
- If your ODBC application provides the option to select one or more columns as unique identifiers, do not choose any columns. For example, when linking tables in Microsoft Access, you might be presented with a "Select Unique Record Identifier" dialog. In this case, just click "OK" without selecting any fields. This feature is usually associated with maintaining integrity for database updates; but since it is not possible to update the Tivoli Storage Manager database with the ODBC driver, there is no need to select any fields. Further, selecting one or more fields might result in erroneous output. One common symptom is seeing "#DELETED" in the result set, in which case the solution is to relink the table and not select any unique identifiers.
- For 64-bit bit Microsoft(R) Windows(R) XP and 2003 users: The ODBC driver is available in a 32-bit version only, and thus is not compatible with 64-bit applications. The ODBC driver can be installed on 64-bit Windows XP and 2003 systems, but must be used with 32-bit applications. Also, the 32-bit version of the ODBC Data Source Administrator must be used to create and configure DSNs (see defining data sources).
- If the ODBC application is running in a different locale from that of the server, data containing characters that do not map into the ODBC application machine's locale may appear unintelligible (garbled).
- When opening Tivoli Storage Manager database tables in Lotus(R) Approach(R), the Open dialog includes a field labeled Files of type. Select ODBC Data Sources (*) for this field. Do not select Tivoli Storage Manager ODBC Driver (*), because Approach will not open the tables.
- If you use a file DSN to link Tivoli Storage Manager database tables into Microsoft Access, you will not be able to open the linked table after Access has been shut down and restarted. This problem is discussed in Microsoft support article Q327268.
- The ANSxxxxx messages can now be displayed in languages other than English. Available languages include Chinese (Simplified and Traditional), Czechoslovakian, English, French, German, Hungarian, Italian, Japanese, Korean, Polish, Portuguese (Brazilian), Russian, and Spanish. However, the Tivoli Storage Manager ODBC configuration dialogs and other non-ANSxxxxx text strings are currently displayed in English only.
- In order to maintain the database integrity, the Tivoli Storage Manager ODBC Driver supports only SELECT statements. You will not be able to manipulate the database with any other statements.
- If there are more than two tables specified in the FROM clause of a SELECT statement, the query results will not be optimized by the Tivoli Storage Manager server. You may see duplicate records in the results.

  Example: The following SQL statement will get duplicate records:

  SELECT * FROM NODES, FILESPACES, CONTENTS

  Note: The SQL statement execution time may be much longer due to preparing large amounts of query results.
- The columns in the ORDER BY clause must be in the select list. The Tivoli Storage Manager ODBC driver provides the information for the application. Some applications do not retrieve the information before generating the SQL query.

  Example: The following SQL statement will fail because PWSET_TIME is not in the select list.

  SELECT ADMIN_NAME FROM ADMINS ORDER BY PWSET_TIME

- Not all Tivoli Storage Manager database tables have a primary key. When selecting Link Tables using Microsoft Access, Microsoft Access will prompt the users for primary key fields if the table does not have a primary key, or if the table has a primary key that is comprised of multiple index fields. You may select primary key fields or click **Cancel** to ignore the prompt.

Main Menu | ODBC driver | Supported ODBC applications and functions

Supported ODBC applications

The Tivoli Storage Manager ODBC driver is supported with any ODBC application, within the constraints specified in the "Supported ODBC functions" and "Known Problems and Limitations" sections of this document, as well as any other constraints that might be specified in the README.ODBC.TXT file that accompanies the Tivoli Storage Manager ODBC driver.

Resolution of problems due to limitations or defects in the ODBC application is the responsibility of that application's vendor.

Supported ODBC functions

The following ODBC 3.x functions are implemented in the Tivoli Storage Manager ODBC driver:

- SQLAllocHandle
- SQLBindCol
- SQLBindParameter
- SQLCancel
- SQLCloseCursor
- SQLColAttribute
- SQLColumns
- SQLConnect
- SQLDescribeCol
- SQLDisconnect
- SQLDriverConnect
- SQLExecDirect
- SQLExecute
- SQLFetch
- SQLFreeHandle
- SQLFreeStmt
- SQLGetConnectAttr
- SQLGetCursorName
- SQLGetData
- SQLGetDiagRec
- SQLGetEnvAttr
- SQLGetFunctions
- SQLGetInfo
- SQLGetStmtAttr
- SQLGetTypeInfo
- SQLNumParams
- SQLNumResultCols
- SQLPrepare
- SQLPrimaryKeys
- SQLRowCount
- SQLSetConnectAttr
- SQLSetCursorName
- SQLSetEnvAttr
- SQLSetStmtAttr
- SQLSpecialColumns
- SQLStatistics
- SQLTables

The ODBC 3.x specification can be found in the Microsoft(R) Developer Network (MSDN) library. At the time of this writing, the URL is: ODBC 3.x specification. From there, navigate to the ODBC specification, using these links (in order):

1. Win32 and COM Development
2. Data Access and Storage
3. Microsoft Open Database Connectivity (ODBC)
4. Microsoft Open Database Connectivity (ODBC)
5. ODBC Programmer's Reference

Main Menu | ODBC driver | Tracing

Tracing.

There are two components from which you must collect Tivoli Storage Manager ODBC driver trace data. When collecting trace data, you need to collect:

- ODBC Driver Manager trace
- Tivoli Storage Manager ODBC driver trace

The following are the steps involved in collecting trace data:

1. Activate Tivoli Storage Manager ODBC driver tracing. To do this, start the ODBC Data Source Administrator (odbcad32.exe), select the data source, and click **Configure**. This will open the Configure a Tivoli Storage Manager Data Source dialog. Select the **Enable trace** check box, and enter a fully-qualified trace file name; for example, C:\TSMODBC.TRC. Then click **OK**. Leave the ODBC Data Source Administrator running because you will need it in the next step.
2. Activate ODBC Driver Manager tracing. To do this, click the Tracing tab in the ODBC Data Source Administrator. In the Log file Path text box, enter a fully-qualified trace file name; for example, C:\ODBCMGR.TRC. Do not use the same name as in step 1 above. Click **Apply**.
3. Next, click **Start Tracing Now**. The button's name will change to Stop Tracing Now. If the button already was labeled Stop Tracing Now, then click the button to change it to Start Tracing Now, then click it again to restart the trace (the button will once again read Stop Tracing Now). Click **OK**.
4. Delete (or rename) any existing Tivoli Storage Manager ODBC driver trace. Using the same name as in step 1 above, delete or rename C:\TSMODBC.TRC (if it exists.)
5. Delete (or rename) any existing ODBC Driver Manager trace. Using the same name as in step 2 above, delete or rename C:\ODBCMGR.TRC (if it exists.)
6. Re-create the problem.
7. After the problem has been recreated, rename the two trace files as different names. Using the sample names above, you could rename them as follows:

   C:\TSMODBC.TRC to C:\TSMODBC.KEEP.TRC
   C:\ODBCMGR.TRC to C:\ODBCMGR.KEEP.TRC

8. Zip up the renamed trace files (they may become large) and be prepared to send them in to IBM/Tivoli support for analysis.
9. Deactivate ODBC Driver Manager tracing. Using the ODBC Data Source Administrator, select the Tracing tab and click **Stop Tracing Now**. Then click **OK**.
10. Deactivate Tivoli Storage Manager ODBC driver tracing. Using the ODBC Data Source Administrator, select the data source and click **Configure**. Clear the **Enable trace** check box, and then click **OK**.

Main Menu | Server crash

Server crash.

The server may crash for one of these 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 execution. 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 validation failures are referred to as an *assert*. If the server terminates due to an assert, the following message will be issued:

  ANR7837S Internal error *XXXNNN* detected.

  where *XXXNNN* is an identifier assigned to the assertion failure. Examples of this identifier are DBREC106, BUF011, or TBUNDO096.

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

Whether the server terminated as a result of an assert or the system trap handler, the following information should be collected and reported to IBM service the situation to be diagnosed:

1. **Server error file (dsmserv.err):**
   When the server crashes, it will append information to the file **dsmserv.err** which is located in the same directory where the server is installed.
   - For the Windows(R) platform, if the server is running as a service, then the file will be named **dsmsvc.err**.
   - For the z/OS(R) platform, there is no separate error file. The information will be in the JES joblog and in the system log. For z/OS they are generally written by WTO with routecode=11.
   - For HP-UX, it is possible to use the following script to get basic thread information from the core on the customer system without needing to send it in to IBM. The procedure is as follows:
     - Make sure gdb is installed on the customer system.
     - Copy the gt shell script to the server bin directory (where the server executable and core file are located).
     - Make sure the script is an executable (chmod a+x gt).
     - Invoke the script with the paths/names of the executable (default is ./dsmserv) and the core (default is ./dsmcore).
     - Output will be in file dsm_gdb.info (which should be sent to IBM).
   - For Linux and AIX, as of 5.3 and later, the trap handler is disabled. This will prevent the function traceback from being printed on the console and in **dsmserv.err**. This change is required in order to ensure that we will get a more complete core file. As part of disabling the trap handler, a new script, getcoreinfo, was added to the Linux packages. The getcoreinfo script will get the function traceback for the failing thread as well as register values and function traceback for all other threads. The amount of information available in the core for other threads is still a little incomplete on some Linux platforms/distributions. See the getcoreinfo script (in the server bin directory) for more details.

2. **System image (core file):**
   Typically, a core file or other system image of the memory is currently in use by Tivoli Storage Manager at the time of the failure. In each case this file should be renamed to prevent it from being overwritten by a later crash. For example a file should be renamed to "core.Aug29" instead of just "core". The type and name of the core file varies depending upon the platform:
   - For UNIX(R) systems, such as AIX(R), HP-UX, Sun Solaris, PASE and Linux, there is typically a file called **core** created. Be sure there is enough space in the server directory to accommodate a dump operation. It is common to have a dump file as large as 2GB for the 32-bit Tivoli Storage Manager server. Additionally, make sure the ulimit for core files is set to unlimited to prevent the dump file from being truncated.
   - For Windows systems before 5.2, a **drwtsn.dmp** file is created if the system is configured to do so. From a DOS prompt issue **drwtsn32 -i** to enable dumps. Beginning with 5.2 the dump will be performed automatically through a system API call. If the server is running as a service the dump file will be called **dsmsvc.dmp**. Otherwise, the dump will be called **dsmserv.dmp**.
   - For a z/OS server, a dump should have been created based upon what was specified in the SYSMDUMP DD statement for the JCL used to start the server. A SYSMDUMP should be used instead of a SYSABEND or SYSUDUMP dump, because a SYSMDUMP will contain more useful debugging information. Ensure that the dump options for SYSMDUMP include RGN, LSQA, TRT, CSA, SUM, SWA, and NUC on systems where a Tivoli Storage Manager address space will execute. To display the current options, issue the z/OS command DISPLAY DUMP,OPTIONS. You can use the CHNGDUMP command to alter the SYSMDUMP options. Note that this will only change the parameters until the next IPL is performed.

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

3. **Libraries and other files:**
   For the UNIX platforms, core files are specific to the application, libraries, and other system resources in use by the application on the system where it was running. To accurately read the core file on our system, we may need all of the below files (located in the in the same directory where the server is installed):
   - dsmserv
   - dsmlicense
   - ndmpspi
   - dsmcored
   - dsmaio
   - centera

The needed library files vary between all the platforms:

- For AIX/PASE systems collect the following files:
  - /usr/ccs/lib/libpthreads.a
  - /usr/ccs/lib/libc.a
  - Also collect any other loaded libraries such as message exits. To see what libraries are loaded, invoke dbx using **dbx dsmserv core_file**. Then, from the dbx prompt use the command **map**. This will show all of the libraries that are loaded and needed for core analysis.

- For HP-UX, use the "chatr dsmserv" command. Send in all the dynamic shared libraries. For example:
  - /usr/lib/libpthread.1
  - /usr/lib/libm.2
  - /usr/lib/libstd.2
  - /usr/lib/libstream.2
  - /usr/lib/libCsup.2
  - /usr/lib/libcl.2
  - /usr/lib/libc.2
  - /usr/lib/libdld.2

- For Linux, use the "ldd dsmserv" command. Send in all the dynamic shared libraries. For example:
  - libm.so.6 => /lib64/libm.so.6
  - libnsl.so.1 => /lib64/libnsl.so.1
  - libpthread.so.0 => /lib64/libpthread.so.0
  - libdl.so.2 => /lib64/libdl.so.2
  - libc.so.6 => /lib64/libc.so.6
  - /lib64/ld64.so.1 => /lib64/ld64.so.1

- For Solaris, run the following commands to collect the needed libraries:
  - sh
  - cd /usr
  - (find . -name "ld.so" -print ; \
  - find . -name "ld.so.?" -print ; \
  - find . -name "libm.so.?" -print ; \
  - find . -name "libsocket.so.?" -print ; \
  - find . -name "libnsl.so.?" -print ; \
  - find . -name "libthread.so.?" -print ; \
  - find . -name "libthread_db.so.?" -print ; \
  - find . -name "libdl.so.?" -print ; \
  - find . -name "libw.so.?" -print ; \
  - find . -name "libgen.so.?" -print ; \
  - find . -name "libCrun.so.?" -print ; \
  - find . -name "libc.so.?" -print ; \
  - find . -name "libmp.so.?" -print ; \
  - find . -name "libc_psr.so.?" -print ; \
  - find . -name "librtld_db.so.?" -print) > runliblist
  - tar cfh runliblist.tar -I runliblist

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

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

6. Package all the data (files) collected and contact IBM service to report this problem. See Reporting a Problem for additional information. For all files from z/OS make sure to use xmit to create a fixed-blocked (FB) file because it will be moved by ftp to an open systems machine, which will cause the loss of DCB attributes for the file. For the open systems platforms, package the files in a zip or tar file.

Main Menu | Database

Database.

Click the symptom in the table or skip the table and scroll through the information below:

List of database symptoms
• ANR0102E, ANR0103E, or ANR0104E message issued. • Maximum reduction value does not increase as the percentage utilized decreases. • Out of database space. • Unable to read the database during server restart.

ANR0102E, ANR0103E, or ANR0104E message issued.

These messages are used to report errors during database operations.

These messages are used to report errors during an insert, update, and delete operations respectively. Database operations typically fail for two primary reasons. These reasons are:

1. There is a timing error between server threads. In order for this error to be encountered, two or more operations have to be working on the same or similar resources within the server. Many times this can be avoided by reducing the number of sessions or processing operating on the same thing at the same time.

2. The database has been damaged. A typical symptom for the ANR0102 message is "Error 1" which indicates a duplicate row. A typical error for the ANR0103E message may be either an "Error 1" which indicates a duplicate row or an "Error 2" which indicates a missing row. A typical error for the ANR0104E message is "Error 2" which indicates a missing row. For the table reported, issue the command SHOW TBLSCAN *tableName* where *tableName* is the name of the table reported in the message. Note that the table name used by this command is case sensitive and it should be entered exactly as it was displayed in the message. This command will scan the entire table and evaluate the structure of that table. If it detects an error in the structure of the table, it will report that an error was found. In cases where the database table is severely damage, this command may cause the server to crash. Also, if this is a large table within the database, this evaluation may require a significant amount of I/O and may degrade the server performance. If SHOW TBLSCAN reports a problem, contact support. Support will first work with you to try to determine the cause of the database damage. In most cases, the cause is external to Tivoli Storage Manager and is usually an error in the storage device where the one or more of the server database volumes are allocated. After helping to determine the cause of the database damage, support will instruct you on the necessary steps to repair this damage in the database.

Return to list of database symptoms

Maximum reduction value does not increase as the percentage utilized decreases.

Issue QUERY DB over a period of time when the "Pct. Util" has decreased.

The server allocates database space from the beginning of the database, which is the lowest allocated page, to the end of the database, which is the highest allocated page. Because allocation is ascending, it is possible for database space to be freed and become available for use between the lowest and highest allocated pages. The "Maximum Reduction" reported by QUERY DATABASE cannot reduce the database below the highest allocated page even though free space may exist. The server will use the free space between the lowest

and highest allocated page as necessary.

It is possible to reorganize the server database. Reorganizing the database will typically reduce the pages used which results in reduced percentage utilization. Similarly, reorganizing the server database may increase the maximum reduction value after it has been complete. Database reorganization is done by issuing DSMSERV UNLOADDB, DSMSERV LOADFORMAT, and then a DSMSERV LOADDB. The database reorganization requires the server to be offline or not available for client use. Review these commands and this procedure in the *Administrator's Reference* and *Administrator's Guide* in the Tivoli Storage Manager publications for your platform.

Return to list of database symptoms

Out of database space.

Issue QUERY DB. If the "Pct. Util" is very high, such as close to 100%, the server will report out of database space during operations that need to add or change database information.

You can take a number of steps to provide more database space.

- If the "Maximum Extension" value reported by the QUERY DB command is greater than zero, issue EXTEND DB *nn* where *nn* is the available space that the db can be extended as reported by QUERY DB.
- Define additional database volumes using the DEFINE DBVOLUME command. After defining one or more additional database volumes, issue QUERY DB and extend the database by all or some of the available extension amount using EXTEND DB *nn* where *nn* is the extension amount.
- Issue QUERY STATUS and evaluate the "Activity Log Retention Period" and the "Activity Summary Retention Period." By reducing either or both of these values, the database space consumed by this information will be reduced. These may be set to lower values using the SET ACTLOGRETENTION and SET SUMMARYRETENTION commands.

Return to list of database symptoms

Unable to read the database during server restart.

If the server fails to read a required page during startup processing, the server will fail to start.

You can take a number of steps to resolve this:

- Were the database volumes mirrored as DBCOPY volumes? If so, change or add the MIRRORREAD DB option to MIRRORREAD DB VERIFY and try to restart the server. This forces the server to attempt to read the database page from all possible database volumes that would contain that page and to compare them to each other. If it can read a valid copy of the database page from one of the database volumes, it will use that copy and may allow the server to restart in this case.
- Check the disk drives, adapters, and connections between the server computer and the disk storage where the Tivoli Storage Manager database volumes are defined. If any of the hardware between the server computer and Tivoli Storage Manager database volumes is reporting a failure or other error, investigate and resolve the error and then try to restart the Tivoli Storage Manager server.

Return to list of database symptoms

Main Menu | Server diagnostic tips

Server diagnostic tips.

For a server problem, review these steps to try to isolate or resolve the problem:

Diagnosing a server problem

1. Check the server activity log.
2. Check HELP for Tivoli Storage Manager messages issued for the problem.
3. Is the problem re-creatable?
4. Is the problem an error reading or writing to a device?
5. Have any server options or settings been changed?
6. Did a scheduled client operation fail?
7. Did the server run out of space?
8. Are connections by clients or administrators failing?
9. Are there server error messages indicating a code page conversion failure?

Check the server activity log.

Check the server activity log for other messages 30 minutes before and after the time of the error. To review the messages in the server activity log, issue QUERY ACTLOG (see the *Administrator's Reference* for more details).

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

Return to diagnosing a server problem

Check HELP for Tivoli Storage Manager messages issued for the problem.

Check HELP for any messages issued by Tivoli Storage Manager.

Tivoli Storage Manager messages provide additional information beyond just the message itself. The **Explanation, System Action,** or **User Response** sections of the message may provide additional information about the problem. Often, this supplemental information about the message may provide the necessary steps to resolve the problem.

Return to diagnosing a server problem

Is the problem re-creatable?

If a problem can be easily or consistently re-created, it may be possible to isolate the cause of the problem to a specific sequence of events as causing it.

Many problems occur as a result of a combination of events. For example, expiration running along with nightly scheduled backups for twenty clients. In some cases, by changing the timing or order of execution for events, this may prevent the problem from reoccurring. In this example, one way to change the timing would be to run expiration at a time when the nightly scheduled backups for the twenty clients was not running.

Return to diagnosing a server problem

Is the problem an error reading or writing to a device?

If the problem is an error reading or writing data from a device, many systems and devices record information in a system error log. Examples of the system error log are errpt for AIX(R), Event Log for Windows(R), and the System Log for z/OS(R).

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

Return to diagnosing a server problem

Have any server options or settings been changed?

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

If there have been one or more configuration changes to the server, try reverting the settings back to their original values and retry the failing operation. If the operation succeeds, try to make one change at a time and retry the operation until the attribute change that caused the failure has been identified.

Return to diagnosing a server problem

Did a scheduled client operation fail?

Scheduled client operations are influenced by the schedule definitions on the server as well as the scheduling service (Tivoli Storage Manager scheduler) being run on the client machine itself.

In many cases, if a schedule changes on the server, it may require the scheduling service on the client to be stopped and restarted.

Return to diagnosing a server problem

Did the server run out of space?

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

To determine if the database is out of space, issue QUERY DATABASE. If the "%Util" is at or near 100% then more space should be defined. Typically, if the database is running out of space, other server messages will be issued indicating this. To add more space to the database, allocate one or more new database volumes, define them to the server using DEFINE DBVOL, and then extend the available database space using EXTEND DB.

To determine if a storage pool is out of space, issue QUERY STGPOOL. If the "%Util" is at or near 100%, then more storage space should be made available. Typically, if operations fail because a storage pool is out of space, other server messages will be issued indicating this. To add more space to a DISK storage pool, allocate one or more new storage pool and define them to the server using DEFINE VOLUME. 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 using UPDATE STGPOOL command.

Return to diagnosing a server problem

Are connections by clients or administrators failing?

There may be many different causes for connection problems to the server. There are two main cases for connection failure. First, a general failure, where no connections at all are allowed. If no connections at all are possible, it may be necessary to run the server in the foreground so that a server console is available, and additional diagnostic steps can be taken. Second, an isolated failure where some connections are allowed but others fail.

A number of settings should be checked to verify proper configuration for communication to the server:

- Check that the server was able to bind to a port when is started. If it was 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 will be issued:

  ANR8200I TCP/IP driver ready for connection with clients on port 1500.

If the server is configured for HTTP communications for administrative session and successfully binds to a port on startup, the following message will be issued:

ANR8280I HTTP driver ready for connection with clients on port 1501.

If a given communication method is configured in the server options file but a successful bind message is not issued during server startup, then there is a problem initializing for that communication method.

- Verify that the code TCPPORT setting in the server options file is correct. If this was changed advertently, this would cause clients to fail to connect because they 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 different. 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 QUERY STATUS and verify that "Availability: Enabled" is set. If this reports "Availability: Disabled", issue the command ENABLE SESSIONS.
- If only specific clients are not allowed 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 command QUERY NODE *nodeName* where *nodeName* is the name of the node to check. If this reports "Locked?: Yes", then evaluate why this node has been locked. Nodes can only be locked using the administrative command LOCK NODE. If it is appropriate to unlock this node, issue the command UNLOCK NODE *nodeName* where *nodeName* is the name of the node to unlock.
- Work with the local network administrator to determine if there have been network changes that account for why the failing nodes are unable to connect to the server.
- If the computer that the server is running on is having memory or resource allocation problems, it may not be possible to start new connections to the server. This may be temporarily cleared up by either halting and restarting the server or by halting and rebooting the computer itself. This is a temporary solution and diagnosis should be continued for either the operating system or the Tivoli Storage Manager server because this may indicate an error in either.

Return to diagnosing a server problem

Are there server error messages indicating a code page conversion failure?

The Tivoli Storage Manager server uses operating system functions to convert between Unicode and the server code page. If the system is not set up correctly, then this will fail. This problem is most likely to occur when using the Administrative Center to access the server.

To get more information on the failure, begin tracing the UNICODE trace class (see here for information on how to trace) and repeat the action that caused the error message to occur. Check the server README for any platform-specific requirements for language installation. Ensure the locales indicated by the problem code pages are installed and any requirements listed in the README are installed.

Main Menu | Hang or loop

Hang or loop

A hang is when the server does not start or complete a function and is not using any CPU. A hang could be just one session or process that is not processing, or it could be the entire Tivoli Storage Manager server not responding. A loop is when no progress is being made, but the server is using a high amount of CPU. A loop can affect just one session or process, or it could affect the entire Tivoli Storage Manager server.

Collecting documentation to resolve this type of problem depends on whether the server is able to respond to commands.

- For a hang or a loop where the server can respond to commands, issue the following commands to help determine the cause of the hang:
  - QUERY SESSION

  - QUERY PROCESS
  - SHOW THREADS
  - SHOW DEADLOCK
  - SHOW TXNT
  - SHOW DBTXNT
  - SHOW LOCKS
  - SHOW DBV
  - SHOW SSS
  - SHOW CSV (do this only if the problem appears to be related to scheduling)

- In addition to the output from the listed commands, or in the case of a server that cannot respond to commands, collect a dump. The way to collect a dump will depend on the operating system.
  - For UNIX(R) type operating systems, use **kill -11** command on the dsmserv process to create a core file. The process ID to perform the kill is obtained by using the **ps** command. See the section on Server Crash for information on additional libraries or executables that need to be packaged along with the core file.
  - For the Windows operating system refer to Microsoft Knowledge Base item http://support.microsoft.com/default.aspx?scid=kb;en-us;241215 for instructions on installing and using the userdump.exe program to obtain a dump.
  - For the z/OS(R) platform, use the operating system DUMP command. The SDUMP options should specify RGN, LSQA, TRT, CSA, SWA, and NUC.

Main Menu | Recovery log

Recovery log.

Click the symptom in the table or skip the table and scroll through the information below:

List of recovery log symptoms
• "ANR0314W Recovery log usage..." message received with LOGMODE=ROLLFORWARD. • Out of recovery log space. • Unable to read the recovery log during server restart.

"ANR0314W Recovery log usage..." message received with LOGMODE=ROLLFORWARD

The server will begin reporting that the recovery log usage is getting high using the ANR0314W message when the "Pct. Util" exceeds 80%. If the server is running with LOGMODE=ROLLFORWARD, this may put the server in jeopardy of running out of recovery log space.

The server retains all the recovery log information between database backups when the LOGMODE=ROLLFORWARD. If a database backup is not run manually, scheduled, or triggered, the server will run out of recovery log space. Consider these steps to reduce the size of your recovery log:

- Manually run a database backup using the BACKUP DB command. The database backup may be run as TYPE=FULL or TYPE=INCREMENTAL. Note that the recovery log utilization will continue to increase until the database backup completes, so adequate time must be given for the operation to complete before the server runs out of recovery log space.
- Issue QUERY DBBACKUPTRIGGER.
  - If a DBBACKUPTRIGGER is not defined, define one using DEFINE DBBACKUPTRIGGER. This allows the server to automatically start a database backup based upon the trigger setting for the LOGFULLPCT.
  - If a DBBACKUPTRIGGER is defined:
    - Evaluate and consider lowering the LOGFULLPCT.
    - If the LOGFULLPCT is higher than the current "Pct. Util" for the recovery log, wait until the LOGFULLPCT is reached and a database backup is triggered. If the LOGFULLPCT is greater than 80%, the ANR0314W message will

always be issued until the database backup is triggered and completed.
  - Check previous database backups to make sure they were successful. If there are no volumes available for the triggered database backup to use or there is an error processing the database backup, the server will not be able to relieve the filling recovery log situation.
- Evaluate whether or not LOGMODE=ROLLFORWARD is needed in your environment. If it is not needed, issue SET LOGMODE=NORMAL. This allows space in the recovery log to be recovered without requiring a database backup to be run.

Return to list of recovery log symptoms

Out of recovery log space.

Issue QUERY LOG. If the "Pct. Util" is very high (close to 100%) the server will report "out of recovery log space" during operations that need to add or change database information.

You can take a number of steps to provide more recovery log space.

- If the "Maximum Extension" value reported by the QUERY LOG command is greater than zero, issue EXTEND LOG *nn* where *nn* is the available space that the recovery log can be extended as reported by QUERY LOG.
- Define additional recovery log volumes using the DEFINE LOGVOLUME command. After defining one or more additional recovery log volumes, issue QUERY LOG and extend the recovery log by all or some of the available extension amount using EXTEND LOG *nn* where *nn* is the extension amount.
- If the server crashed because it was out of recovery log space and fails to start, it may be necessary to do an emergency extend of the recovery log. This is done by issuing DSMSERV EXTEND LOG *volumeName amount* where *volumeName* is a formatted recovery log volume and *amount* is the amount (in megabytes) to extend the recovery log.

The recovery log has a maximum size of 13 GB. It is recommended that the recovery log only be defined or extended to 12 GB so that there is 1GB in reserve if an emergency extend needs to be done. In the event that an emergency extend is done, reduce the recovery log after restarting the server to preserve the 1 GB reserve amount.

It is possible that the recovery log will appear to be out of space when in fact it is being *pinned* by an operation or combination of operations on the server. A pinned recovery log is where space in the recovery log can not be reclaimed and used by current transactions because an existing transaction is processing slowly or is hung. To determine if the recovery log is pinned, issue SHOW LOGPINNED repeatedly over many minutes. If this reports the same client session or server processes as pinning the recovery log, it may be necessary to take action to cancel or terminate that operation in order to keep the recovery log from running out of space. To cancel or terminate a session or process that is pinning the recovery log, issue SHOW LOGPINNED CANCEL. Server version 5.1.7.0 and later, and 5.2.0.0 and later have additional support for the recovery log to automatically recognize that the recovery log is running out of space and where possible to detect and resolve a pinned recovery log using the SHOW LOGPINNED processing.

Return to list of recovery log symptoms

Unable to read the recovery log during server restart.

If the server fails to read from the recovery log during startup processing, the server will not start.

Consider these steps:

- Check if the recovery volumes are mirrored. If so, change or add the MIRRORREAD LOG option to MIRRORREAD LOG VERIFY and try to restart the server. This forces the server to attempt to read the recovery log information from all possible recovery log volumes that contain the startup information, and compare them to one another. If it can read the requested information from one of the recovery log volumes, it will be used and may allow the server to restart.
- Check the disk drives, adapters, and connections between the server computer and the disk storage where the Tivoli Storage Manager recovery log volumes are defined. If any of the hardware between the server computer and Tivoli Storage Manager recovery log volumes is reporting a failure or other error, investigate and resolve the error and then try to restart the Tivoli Storage Manager server.

Return to list of recovery log symptoms

Main Menu | Processes

Processes.

For more information about server processes, see:

- What is a process?
- Process messages

Click the symptom in the table or skip the table and scroll through the information below:

List of process symptoms

- "ANR1221E COMMAND: Process *processID* terminated - insufficient space in target copy storage pool."
- "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."
- Files are not expired after reducing the number of versions that need to be kept.
- Migration does not run for a sequential media storage pool.
- Migration only uses one process.
- Process running slow.

"ANR1221E COMMAND: Process *processID* terminated - insufficient space in target copy storage pool."

Issue QUERY STGPOOL *stgpoolName* F=D. The following SQL select statement should also be issued from an administrative client to this server: "**select** stgpool_name,devclass_name,count(*) as 'VOLUMES' from volumes group by stgpool_name,devclass_name".

Compare the number of volumes reported by the select statement to the maximum scratch volumes allowed (as reported by QUERY STGPOOL). If the number of volumes reported by the **select** is equal to or exceeds the "Maximum Scratch Volumes Allowed," update the storage pool and allow more scratch volumes by issuing UPDATE STGPOOL *stgpoolName* MAXSCR=*nn* 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. **Note:** 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.

Return to list of process symptoms

"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."

Issue QUERY VOLUME *volumeName* F=D. The following SQL select statement should also be issued 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 may report files as damaged because, at the time that the data was written, the hardware was malfunctioning and did not write the data correctly, even though it reported to the Tivoli Storage Manager server that the operation had been successful. As a result of this device malfunction, many files on many different volumes may be affected. You can take the following steps to resolve this:

- Evaluate the system error logs or other information about this drive to determine if it is still reporting an error. If errors are still being reported, they must be resolved first. To resolve a hardware issue, work with the hardware vendor to correct the problem.
- If this is a copy of a storage pool volume, simply delete this volume using DELETE VOLUME *volumeName* DISCARDDATA=YES.

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 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, backup the files again from the client.
- If this is a primary storage pool volume but the data was put on this volume by MIGRATION, MOVE DATA, or MOVE NODEDATA, then there may be an undamaged copy of the file on the server. If the primary storage pool that contained this file was backed up to a copy storage pool prior to the MIGRATION, MOVE DATA, or MOVE NODEDATA then an undamaged file may exist. If this is the case, issue UPDATE VOLUME *volumeName* ACCESS=DESTROYED and then issue RESTORE VOLUME *volumeName* to recover the damaged files from the copy storage pool for this volume.

Return to list of process symptoms

Files are not expired after reducing the number of versions that need to be kept.

The server policies have been updated to reduce the number of versions of a file to retain. Issue QUERY COPYGROUP *domainName policySetName copyGroupName* F=D. If either the "Versions Data Exists" or "Versions Data Deleted" parameters were changed for a TYPE=BACKUP copy group, this may affect expiration.

If the "Versions Data Exists" or "Versions Data Deleted" values for a TYPE=BACKUP copy group have been reduced, the server expiration process may not immediately recognize this and expire these files. The server only applies the "Versions Data Exists" and "Versions Data Deleted" values to files at the time they are backed up to the server. When a file is backed up, the server will count the number of versions of that file and if that exceeds the number of versions that should be kept, the server will mark the oldest versions that exceed this value to be expired.

Return to list of process symptoms

Migration does not run for sequential media storage pool.

Issue QUERY STGPOOL *stgpoolName* F=D.

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 migratable data in use for the storage pool relative to the total number of volumes that can be used for that storage pool. Because it may be considering unused scratch volumes in this calculation, there may not appear to be sufficient migratable data in the storage pool to require migration processing.

Return to list of process symptoms

Migration only uses one process.

Issue QUERY STGPOOL *stgpoolName* F=D and QUERY OCCUPANCY * * STGPOOL=*stgpoolName*.

Some reasons that only one migration process is being run include:

- The Migration Processes setting for the storage pool is set to one or is not defined (blank). If this is the case, issue UPDATE STGPOOL *stgpoolName* MIGPROCESS=*n* where *n* is the number of processes to use for migrating from this pool. Note that this value should be less than or equal to the number of drives (mount limit) for the NEXT storage pool where migration is storing data.
- If QUERY OCCUPANCY 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 partitions data based on client node and file space. In order for migration to run with multiple processes, data for more than one client node and file space needs to be available in that storage pool.

Return to list of process symptoms

Process running slow.

Issue QUERY DB F=D. If the Cache Hit Ratio reported is less then 98%, the process may be running slow because it is waiting for database I/O processing due to the small size of the buffer pool.

This can be addressed by changing or adding either of the following server options to the options file. Increase the BUFPOOLSIZE server option or set SELFTUNEBUFPOOLSIZE YES. For setting the BUFPOOLSIZE option, it is generally recommended that it be set to half the available physical memory on the machine. After setting either of these options as recommended, halt and restart the server.

Return to list of process symptoms

Main Menu | Processes | Process Messages

Process Messages.

Server processes, whether run in the FOREGROUND or BACKGROUND, will always issue a *process started* message and a *process ended* message. The process ended message issued will vary depending upon whether or not the process needs to report information about no items or bytes processed, items and bytes processed, items processed, or just bytes processed. The following are the server process messages:

List of process messages

- Process started.
- Process ended.
- Process ended with items and bytes.
- Process ended with items.
- Process ended with bytes.

In addition to these general process messages, each process will typically report its own messages relating to when it started and ended.

Process started.

The server may run tasks as processes. Processes are assigned an identification message and report that they have started with the message:

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

Where

process_id	Numeric process identifier.
process_name	The name of the process.
process_state	The process may be run in the FOREGROUND or BACKGROUND. If it is run 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 being run in the BACKGROUND returns immediately to the administrative session that issued the command indicating that a process was started while the process actually still runs. Processes run in the BACKGROUND may be monitored with the "QUERY PROCESS" command.
time	The time the process was started.

Return to list of process messages.

Process ended.

When a process completes and it does not have bytes or number of files to report, it will issue the following messages:

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

Where

process_id	Numeric process identifier.
process_name	The name of the process.
process_state	The process may be run in the FOREGROUND or BACKGROUND. If it is run 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 being run in the BACKGROUND returns immediately to the administrative session that issued the command indicating that a process was started while the process actually still runs. Processes run in the BACKGROUND may be monitored with the "QUERY PROCESS" command.
completion_state	The process reports SUCCESS or FAILURE for the completion state.
time	The time the process was started.

Return to list of process messages.

Process ended with items and bytes.

When a process completes and has bytes and items processed to report, it will issue the following messages:

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*.

Where

Where

process_id	Numeric process identifier.
process_name	The name of the process.
process_state	The process may be run in the FOREGROUND or BACKGROUND. If it is run 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 being run in the BACKGROUND returns immediately to the administrative session that issued the command indicating that a process was started while the process actually still runs. Processes run in the BACKGROUND may be monitored with the "QUERY PROCESS" command.
number_of_items	The number of items processed.
bytes_processed	The number of bytes processed.
completion_state	The process reports SUCCESS or FAILURE for the completion state.
time	The time the process was started.

Return to list of process messages.

Process ended with items.

When a process completes and has items processed to report, it will issue the following messages:

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*.

Where

process_id	Numeric process identifier.
process_name	The name of the process.

process_state	The process may be run in the FOREGROUND or BACKGROUND. If it is run 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 being run in the BACKGROUND returns immediately to the administrative session that issued the command indicating that a process was started while the process actually still runs. Processes run in the BACKGROUND may be monitored with the "QUERY PROCESS" command.
number_of_items	The number of items processed.
completion_state	The process reports SUCCESS or FAILURE for the completion state.
time	The time the process was started.

Return to list of process messages.

Process ended with bytes.

When a process completes and has bytes processed to report, it will issue the following messages:

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

Where

process_id	Numeric process identifier.
process_name	The name of the process.
process_state	The process may be run in the FOREGROUND or BACKGROUND. If it is run 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 being run in the BACKGROUND returns immediately to the administrative session that issued the command indicating that a process was started while the process actually still runs. Processes run in the BACKGROUND may be monitored with the "QUERY PROCESS" command.
bytes_processed	The number of bytes processed.
completion_state	The process reports SUCCESS or FAILURE for the completion state.
time	The time the process was started.

Return to list of process messages.

Main Menu | Processes | What is a process?

What is a process?

A server process is a task performed on the server. This task is assigned to perform a specific operation such as migrating data from a storage pool to the NEXT storage pool in the hierarchy. Or it may be the process for backing up the server database to a sequential media device for the purpose of safeguarding it.

A process has the following general characteristics:

Process name
Name of the process as it appears in the process related messages.

Means of initiation
Server processes may be initiated in a number of ways. Typically, they are either an automated process on the server that may or may not be influenced by a server options or other setting or they are started by a command.

Can run in foreground?
Many server processes can be run in the FOREGROUND or synchronously. Processes that run in the FOREGROUND can be 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 those specified with WAIT=NO are run in the BACKGROUND or asynchronously.

Can run as multiple processes

If a server process is able to initiate multiple processes simultaneously to accomplish the task.

Description

A description of what this task does.

To find information about a server process, click the process in the table below or scroll through the list of processes below the table. The information about processes is in the form:

PROCESS_NAME

Description of the process...

How is the process started?

Other Information about this process...

List of server processes

- AUDIT VOLUME
- BACKUP DB
- BACKUP STORAGEPOOL
- CHECKIN LIBVOLUME
- CHECKOUT LIBVOLUME
- EXPIRATION
- EXPORT
- IMPORT
- LABEL LIBVOLUME
- MIGRATION
- MOVE DATA
- MOVE DRMEDIA
- MOVE MEDIA
- MOVE NODEDATA
- PREPARE
- RECLAMATION
- RESTORE STGPOOL
- RESTORE VOLUME

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.

Initiated by command AUDIT VOLUME.

None.

Return to list of processes.

BACKUP DB

Backup the server database. This may be either a FULL or INCREMENTAL backup.

Initiated by command BACKUP DB. May also be automatically triggered if the server is in ROLLFORWARD mode for recovery log processing and a DBBACKUPTRIGGER is defined.

This may be run as a synchronous process by specifying WAIT=YES.

Return to list of processes.

BACKUP STGPOOL

Backup a primary server storage pool to a copy storage pool for the purpose of making duplicate copies of the data and potentially taking duplicate copies to an off-site location.

Initiated by the command BACKUP STGPOOL.

This may be run as a synchronous process by specifying WAIT=YES. BACKUP STGPOOL may be run using multiple concurrent processes. This is controlled by the MAXPROCESS parameter specified on the BACKUP STGPOOL command.

Return to list of processes.

CHECKIN LIBVOLUME

Check a tape volume into a tape library.

Initiated by the command CHECKIN LIBVOLUME.

None.

Return to list of processes.

CHECKOUT LIBVOLUME

Check a tape volume out from a tape library.

Initiated by the command CHECKOUT LIBVOLUME.

None.

Return to list of processes.

EXPIRATION

Expire (delete) client backup and archive files from the server based on the policies defined to manage those files.

Expiration may be run automatically by specifying EXPINTERVAL other than zero in the server options file. It may also be initiated using the command EXPIRE INVENTORY. Please note that it is not possible to have more than one expiration process running at a time.

This may be run as a synchronous process by specifying WAIT=YES.

Return to list of processes.

EXPORT

Export data to sequential media volumes or directly to another server using TCP/IP communication connections between the servers.

Export processing can be started by any of the following commands: EXPORT ADMIN, EXPORT NODE, EXPORT POLICY, and EXPORT SERVER.

None.

Return to list of processes.

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, and IMPORT SERVER.

Other.

Return to list of processes.

LABEL LIBVOLUME

Label one or more library volumes in a library.

Initiated by command LABEL LIBVOLUME.

None.

Return to list of processes.

MIGRATION

Migrate data from one storage pool to the next in the storage hierarchy.

Migration is started and stopped 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, and as needed, it starts migration processing for that storage pool when the HighMig threshold is exceeded.

Migration may be configured to run multiple concurrent processes. This is controlled by the MIGPROCESS attribute of the storage pool and may be updated by issue UPDATE STGPOOL.

Return to list of processes.

MOVE DATA

Move data from one volume to other volumes in the same storage pool or to a different storage pool.

Initiated by command MOVE DATA.

This may be run as a synchronous process by specifying WAIT=YES.

Return to list of processes.

MOVE DRMEDIA

Manage disaster recovery media by moving onsite volumes offsite or by bringing offsite volumes back onsite. Disaster recovery media is the database backup and storage pool backup volumes necessary to protect and recover the server.

Initiated by command MOVE DRMEDIA.

This may be run as a synchronous process by specifying WAIT=YES.

Return to list of processes.

MOVE MEDIA

Move volumes from a tape library to the overflow location to prevent a library from becoming full.

Initiated by command MOVE MEDIA.

None.

Return to list of processes.

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.

Initiated by command MOVE NODEDATA.

This may be run as a synchronous process by specifying WAIT=YES.

Return to list of processes.

PREPARE

Create a recovery plan file.

Initiated by command PREPARE.

This may be run as a synchronous process by specifying WAIT=YES.

Return to list of processes.

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 will start a reclamation process for that storage pool to reclaim any eligible volumes if it determines that one or more eligible volumes exist.

None.

Return to list of processes.

RESTORE STGPOOL

Restore all files for a given storage pool from a copy storage pool.

Initiated by command RESTORE STGPOOL.

This may be run as a synchronous process by specifying WAIT=YES. RESTORE STGPOOL may be run using multiple concurrent processes. This is controlled by the MAXPROCESS parameter specified on the RESTORE STGPOOL command.

Return to list of processes.

RESTORE VOLUME

Restore all files for a given volume from a copy storage pool.

Initiated by command RESTORE VOLUME.

This may be run as a synchronous process by specifying WAIT=YES. RESTORE VOLUME may be run using multiple concurrent processes. This is controlled by the MAXPROCESS parameter specified on the RESTORE VOLUME command.

Return to list of processes.

Main Menu | Show commands | Show commands for the server or storage agent

Show commands for server or storage agent.

The show commands listed are those that are most typically requested or used for diagnosing problems. This list does not discuss all possible show commands that are available.

To find information about a show command, click the command in the table below or scroll through the list of commands below the table. The information about show commands is in the form:

COMMAND

Description of command

Recommendation for when to use

Other information

List of show commands

- ASQUEUED
- ASVOL
- BFOBJECT
- BUFSTATS
- BUFVARS
- CONFIGURATION
- DBTXNTABLE
- DBVARS
- DEVCLASS
- GROUPLEADERS
- GROUPMEMBERS
- INVOBJECT
- LANFREE
- LIBINVENTORY
- LIBRARY
- LOCK
- LOGPINNED
- LOGVARS
- MEMTREND
- MP
- RESQUEUE
- SESSIONS
- SLOTS
- SSPOOL
- THREADS
- TOCSETS
- TOCVARS
- TXNTABLE
- VERIFYEXPTABLE
- VOLINUSE

ASQUEUED

Displays the mount point queue.

In order to use a drive, a client session or server process must first obtain a mount point. The mount point management on the server allows for queuing waiters for mount points if more mount points are needed than are available. Useful for determining the state of a mount point request especially if a session or process appears to be hung waiting for a mount point.

The syntax is SHOW ASQueued.

Return to list of show commands

ASVOL

Displays assigned volumes.

As sequential media volumes are assigned for use by a session or a process, they are tracked in an in-memory list. Viewing this list is useful to determine the state of in-use volumes as well as hang or deadlock situations where a session or process appears to be stuck waiting for a volume or holding a volume and waiting for something else.

The syntax is SHOW ASVol.

Return to list of show commands

BFOBJECT

Displays information about a bitfile object in the server storage hierarchy.

This command is useful to determine the existence and attributes of a bitfile object in the server's storage hierarchy. You might use this command if you are having trouble restoring, retrieving, expiring, or auditing the object.

The syntax is SHOW BFObject *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 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, etc.

Return to list of show commands

BUFSTATS

Displays usage statistics for the database buffer pool.

Useful for determining if the configured database buffer pool size is large enough.

The syntax is SHOW BUFStats.

Return to list of show commands

BUFVARS

Displays database buffer pool global attributes.

Useful for determining if the configured database buffer pool size is large enough.

This may also be useful for diagnosing cases where the server is hung or when the server runs out of recovery log space. The database buffer pool performance and characteristics can influence the server's running out of recovery log space because the ability to write (flush) the changed pages to the database volumes can impact the ability of the recovery log to manage its space.

The syntax is SHOW BUFVars.

Return to list of show commands

CONFIGURATION

A summary show command that actually issues many different show commands and queries.

Useful for providing a general configuration and other information about the server to IBM service.

The syntax is SHOW CONFIGuration.

Return to list of show commands

DBTXNTABLE

Displays information about transactions that are performing database operations.

For transactions performing database operations, this command displays information about:

- Database tables that are open (in-use)
- Recovery log usage information such as the number of log records written and the recovery log space used
- The first, last, and next recovery log records written
- Whether or not a transaction is valid or being aborted (rolled back)

The syntax is SHOW DBTXNTable.

Return to list of show commands

DBVARS

Displays database global attributes.

Useful for viewing the current state and attributes of the server database.

The syntax is SHOW DBVars.

Return to list of show commands

DEVCLASS

Displays information about device classes.

Useful for displaying the states of allocated drives, device class attributes, and other information. Often used to diagnose problems with devices or hangs waiting for a drive, library, or volume. The command SHOW LIBRARY will also give good complementary information about drives and libraries.

The syntax for this command is SHOW DEVCLass.

Return to list of show commands

GROUPLEADERS

Displays all backup group leaders for an object in the server's inventory.

This command is useful to determine the backup group relationships of an object in the server's inventory. You might use this command if you are having trouble restoring, retrieving, expiring, or auditing the object.

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 being queried. The high-order word is optional; if not specified, a value of zero is assumed. The object must be a backup object.

Only available in 5.2.3.0 or above.

Return to list of show commands

GROUPMEMBERS

Displays all backup group members for an object in the server's inventory.

This command is useful to determine the backup group relationships of an object in the server's inventory. You might use this command if you are having trouble restoring, retrieving, expiring, or auditing the object.

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 being queried. The high-order word is optional; if not specified, a value of zero is assumed. The object must be a backup object.

Only available in 5.2.0.0 or above.

Return to list of show commands

INVOBJECT

Displays information about an inventory object in the server.

This command is useful to determine the existence and attributes of an object in the server's inventory. You might use this command if you are having trouble restoring, retrieving, expiring, or auditing the object.

As of 5.2.2.0, this command will report new information for archive retention protected objects. This command will report if the archive object is in deletion hold. If the object uses event-based retention, this command will report appropriately: if the Expiring. Objects row does not exist, then the object's event has not been signaled; if the basedate in the Expiring.Objects row is later than the insert date, then this is the date the object's event was signaled.

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 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, etc.

Return to list of show commands

LANFREE

A number of definitions are needed on the server in order for a given client to perform LAN-free data movement operations. In cases where these definitions are not present or are incorrect, it may be difficult to determine if the LAN-free environment is configured correctly.

In version 5.2.2 of the Tivoli Storage Manager server, the SHOW LANFREE command was added. This will evaluate all possible destination storage pools for this client node and report whether or not the storage pool is capable of LAN-free data movement operations.

The syntax is SHOW LANFREE *nodeName storageAgent*.

Return to list of show commands

LIBINVENTORY

This command displays the current state of the library inventory for the library specified.

This command is useful if there seems to be a problem with the library inventory information. The command will display current in-memory properties of the library inventory.

The syntax is SHOW LIBINVentory *libraryName* where *libraryName* is optional, and if left out, the command will return the

inventory information for all libraries.

Return to list of show commands

LIBRARY

This command displays the current state of the specified library and all of its drives.

This command is useful to gather a quick view of all in-memory information about a library and its drives. This output should be gathered for any problem related to libraries or drives (e.g. mounting problems).

The syntax is SHOW LIBRary *libraryName* where *libraryName* is optional. If left out, the command will return information for all the libraries.

Return to list of show commands

LOCK

Displays lock holders and waiters.

The server and storage agent use locks as a mechanism to serialize access and updates to information and other constructs. Displaying this information is useful for diagnosing hangs or other resource contention issues.

The syntax is SHOW LOCK.

Return to list of show commands

LOGPINNED

Evaluates and determines whether or not the server recovery log is pinned. A pinned recovery log may cause the recovery log to run out of space and possibly cause the server to crash.

This show command interrogates a number of server control structures and correlates the data to determine if a session, transaction, or process is pinning the recovery log. If it determines that something is pinning the recovery log, it will report this.

The syntax is SHOW LOGPInned. To recover from a pinned recovery log, issue 'SHOW LOGPInned Cancel' to cause the server to cancel or terminate the session, transaction, or process.

Only available in 5.1.7.0 or above.

Return to list of show commands

LOGVARS

Displays recovery log global attributes.

Useful for determining the state of the recovery log.

The syntax is SHOW LOGVars.

Return to list of show commands

MEMTREND

This command will report the memory used by the server, in megabytes, recorded at hourly intervals for the last 50 hours (this is a constant in the server code and is not user-configurable). The command will also display a histogram to help visualize the usage trend.

This command is useful to determine if the server has a memory leak. If the memory usage is constantly increasing, this might indicate a leak. Note that for the measurements to be valid, the measurement period (the last 50 hours) should represent normal, steady-state server activity. The reported usage represents the amount of memory internal server routines request from the pseudo-kernel memory routines. It does NOT represent the total amount of memory the server is using. Nevertheless, it is still useful in determining the server's memory usage trend.

The syntax is SHOW MEMTREnd.

Only available in 5.3.0.0 or above

Return to list of show commands

MP

Displays mount points.

Useful for determining which volume is in-use by a given mount point and other attributes for the assigned mount points. SHOW LIBRARY and SHOW DEVCLASS have useful complimentary information with this command to display the current state of drives and current devclass mp counts.

The syntax is SHOW MP.

Return to list of show commands

RESQUEUE

Displays the resource queue.

The resource queue is used to monitor common resources on the server. If a resource appears to be hung or holding a resource for an unreasonable amount of time, the resource monitoring algorithms for the server will take action and cancel or terminate the resource user. Typically, this is useful for displaying information about transactions, locks, and other resources used by a storage agent on the database server that it is configured to use.

The syntax is SHOW RESQueue.

Return to list of show commands

SESSIONS

Displays information about sessions connected to the server or storage agent.

Useful for diagnosing hangs or other general session problems while a session is still connected to the server. This is also useful in cases where a session is cancelled or terminated and still appears in QUERY SESSION.

SHOW SESSions.

Return to list of show commands

SLOTS

This command displays the current state of the specified library's slot information (e.g. which volumes are in the library and in which slots).

The information displayed is what is saved directly from the library hardware to in-memory values, and can be used to determine if this information is out of sync or incorrect, or to determine if the values returned from the library hardware itself are invalid.

Alternatively, this command can also be used to determine the drive element numbers for a SCSI library if QUERY SAN is unavailable for a particular library (e.g. 3570 library).

The syntax is SHOW SLOTS *libraryName*.

Return to list of show commands

SSPOOL

Displays information storage pools.

Useful for displaying the states and attributes of defined storage pools.

The syntax is SHOW SSPool.

Return to list of show commands

THREADS

Displays information about all threads known to the server.

The server will display information about each thread, typically including the Tivoli Storage Manager thread ID, the system thread id, the thread name, mutexes it holds (if any), and mutex or condition it is awaiting (if any). This command is platform-specific, so each platform may have slightly different information. You might want to use this command if the server or a particular server process appears to be hung in order to see if there are threads waiting for resources held by another thread.

The syntax is SHOW THReads.

Note that on some platforms (e.g. HP), the information reported is obtained without serialization. On a busy system (one that is not hung), the information may be inconsistent: multiple threads may report holding the same mutex, or a thread may report that it is waiting on a mutex held by another thread that doesn't claim to hold it.

On a Tivoli Storage Manager server V5.3 for AIX, SHOW THReads displays a stack trace of each thread. A stack trace reveals information about the server, which can help developers and service diagnose problems. See examples below.

```
Stack trace:
0x09000000000b9a5c select
0x000000010000d06c pkDelayThread
0x00000001007d8290 CsCmdSchedulerThread
0x000000010000a494 StartThread
0x090000000032f2ac _pthread_body
Thread 33: SmAdminCommandThread
tid=8513, ktid=3686529, ptid=21, det=0, zomb=0, join=0, result=0, sess=0
Holding mutex THRV->mutex (0x11173b9f8), acquired at 10000b5c8 (PkShowThreads)

Stack trace:
0x000000010000ba0c PkShowThreads
0x00000001009df84c ShowCommand
0x00000001009df2f0 psCommand
0x0000000100a064a0 AdmShow
0x00000001006f091c AdmCommandLocal
0x00000001006f2758 admCommand
0x000000010087d8d4 SmAdminCommandThread
0x000000010000a494 StartThread
0x090000000032f2ac _pthread_body
```

Return to list of show commands

TOCSETS

Displays all Table Of Contents (TOC) sets known to the server.

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 db. During a restore, one or more TOCs may be loaded into the server database in order to provide file and directory names to the client GUI. This command will display the status of the TOC set (e.g. building or loading), and how much temporary DB space is in use for each TOC set. You might use this command if you are having trouble doing an NDMP backup with the TOC=YES parameter, or having trouble restoring files from an NDMP backup, or if TOC sets are being retained in the server database too long or not long enough.

The syntax is SHOW TOCSets DELETE=*setNum* TOUCH=*setNum*. The DELETE parameter will cause the specified TOC set number to be deleted. The TOUCH parameter will update the last used date of the specified TOC set number. A TOC set will be retained for the TOC retention period following the last used date (see SET TOCRETENTION command).

Only available in 5.2.0.0 or later version.

Return to list of show commands

TOCVARS

Displays information about the Table Of Contents (TOC) component of the server.

This command is useful to determine the status of the toc component. You might use this command if you are having trouble doing an NDMP backup with the TOC=YES parameter, or having trouble restoring files from an NDMP backup.

The syntax is SHOW TOCVars.

Only available in 5.2.0.0 or later version.

Return to list of show commands

TXNTABLE

Displays information about transactions in-use on the server.

Transactions 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 hangs or other transaction related failures while the transaction is still open on the server.

The syntax is SHOW TXNTable.

Return to list of show commands

VERIFYEXPTABLE

Performs verification of the Expiring.Objects table of the server.

This command is useful to find and clean up orphaned entries in the Expiring.Objects table (e.g. entries that don't have corresponding entries in the Backup.Objects or Archive.Objects tables). You might use this command if you are having trouble with inventory expiration.

The syntax is SHOW VERIFYEXPtable.

Make sure the IMEXP trace class is active.

Return to list of show commands

VOLINUSE

Displays whether or not the volume specified is currently in the server's in-use list. This command will display additional information that may be helpful - including whether the volume is currently pending removal from the in-use list.

This command is useful to determine whether or not a volume is currently in the in-use list, and if necessary, to remove it from that list. Note that operations associated with this volume might fail if it is removed from the in-use list.

The syntax is SHOW VOLINUSE *volumeName*. If the volume needs to be removed from the in-use list an additional parameter can be specified to remove the volume from the list: SHOW VOLINUSE *volumeName* REMOVE=YES.

Return to list of show commands Main Menu | Storage Pool

Storage Pool.

Click the symptom in the table or skip the table and scroll through the information below:

List of storage pool symptoms
• "ANR0522W Transaction failed..." message received. • Storage pool experiences high volume usage after increasing MAXSCRATCH value. • Storage Pool has "Collocate?=Yes" but volumes still contain data for many nodes.

"ANR0522W Transaction failed..." message received.

The server was unable to allocate space in the storage pool indicated to store data for the specified client.

There are a number of possible causes for running out of space in a storage pool. Likely causes and solutions include:

- Issue QUERY VOLUME *volname* F=D for the volumes in the referenced storage pool. For any volumes reported with access different from Read/Write, check that volume. A volume may have been marked Read/Only or Unavailable because of a device error. If the device error has been resolved, use the UPDATE VOLUME *volname* ACCESS=READWRITE command to allow the server to select and try to write data to that volume.
- 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. This 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 using UPDATE STGPOOL *stgpoolname* REUSEDELAY=*nn* 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. This is significantly influenced by the number of scratch volumes in a storage pool.

Return to list of storage pool symptoms

Storage pool experiences high volume usage after increasing MAXSCRATCH value.

For collocated sequential storage pools, increasing the MAXSCRATCH value may cause the server to use more volumes.

The server will use more storage pool volumes in this case because of the collocation processing. Collocation will group user data for a client node onto the same tape. During a client backup or archive operation, if no tapes currently have data for this client node, the server will select a scratch volume to store the data. Then for other client nodes storing data, it makes the same decision when selecting a volume and again selects a scratch volume. The reason this does not occur 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 will ignore the collocation request and simply store the data on an available volume.

Return to list of storage pool symptoms

Storage pool has "Collocate=Yes" but volumes still contain data for many nodes.

There are two possible reasons for this:

1. Data was stored on volumes in this storage pool prior to setting "Collocate=Yes".
2. The storage pool ran out of scratch tapes and stored data on the best possible volume even though it ignored the request to collocate.

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

- Issue MOVE DATA for the volume or volumes affected. If scratch volumes are available or volumes with sufficient space are assigned to this client node for collocating their data, the process will read the data from the specified volume and move it to a different volume or volumes in the same storage pool.
- Allow migration to move all the data from that storage pool by setting the HIGHMIG and LOWMIG thresholds to accomplish this. By allowing migration to migrate all data to the NEXT storage pool, the collocation requirements will be honored, provided the NEXT storage pool is set to "Collocate=Yes," it has sufficient scratch volumes, and is assigned volumes to satisfy the collocation requirements.
- Issue MOVE NODEDATA for the client nodes whose data resides in that storage pool. If scratch volumes are available or volumes with sufficient space are assigned to this client node for collocating their data, the MOVE NODEDATA process will read the data from the volumes that this node has data on and move it to a different volume or volumes in the same storage pool.

The key to getting the data collocated is to have sufficient space in the target storage pool for the collocation processing to select an appropriate volume. This is significantly influenced by the number of scratch volumes in a storage pool.

Another alternative is to explicitly define more volumes to the storage pool using DEFINE VOLUME. Again, the key is to have candidate empty volumes for collocation to use rather than using a volume that already has data on it for a different node.

Return to list of storage pool symptoms

Main Menu | Tracing | Enabling stack trace for specific messages for the server or storage agent

Enabling stack trace for specific messages for the server or storage agent

1. A stack trace reveals information about an application that can help IBM Support diagnose problems faster.

   IBM Support may find it helpful to enable stack trace on specific messages issued by the server or storage agent. The following types of Messages on which a stack trace can be enabled include:

   - server console
   - storage agent console
   - administrative client connected to either the server or storage agent
2. To get a stack trace when a specific message is issued by the server or storage agent, enable the message for stack trace. To enable one or more messages for stack trace, issue:

MSgstacktrace enable <*messageNumber*>

Note that <*messageNumber*> may be a space-delimited list of message numbers. For example, this could be entered as MS ENABLE 2017. The MSGSTACKTRACE ENABLE command is cumulative such that extra messages can be enabled by issuing MSGSTACKTRACE ENABLE additional times. For example, if you wanted to add message 985 in addition to those that were already enabled, issue: MS ENABLE 985. Notice that only the number part of the message is allowed in the MSGSTACKTRACE command.

3. To stop getting stack trace for messages issued by the server or storage agent, the stack trace for those messages needs to be disabled. To disable one or more messages, issue:

   MSgstacktrace DISABLE <*messageNumber*>

   Note that <*messageNumber*> may be a space delimited list of message numbers. For example, this could be entered as MSGSTACKTRACE DISABLE 2017 985. Additional messages can also be disabled by issuing MS DISABLE. For example, if you wanted to remove message number 7837 in addition to those that were already disabled, issue: MSgstacktrace DISABLE 7837.

The following messages are enabled for stack trace by default:

7837, 9600, 9601, 9602, 9604, 9605, 9606, 9607, 9608, 6753, 7823, 5010, 5015, 5019, 5021, 5093, 5099, 5100, 5267, 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, and 9999.

Main Menu | Storage agent diagnostic tips

Storage agent diagnostic tips.

For a storage agent problem, review these steps to try to isolate or resolve the problem:

Diagnosing a storage agent problem.

1. Check the server activity log.
2. Check HELP for Tivoli Storage Manager messages issued for the problem.
3. Is the problem an error reading or writing to a device?
4. Have any storage agent options been changed?
5. Have any server options or settings been changed?

Check the server activity log.

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

Storage agents start and manage many sessions to the server. Review the server activity log for messages from the storage agent. To review the activity log messages, issue QUERY ACTLOG. See the Administrator's Reference for more details.

1. If no messages are seen for this storage agent in the server activity log, verify the communication settings:
   - Issue QUERY SERVER F=D on the server and verify that the high-level address (HLA) and low-level address (LLA) set for the server entry representing this storage agent are correct.
   - In the device configuration file specified in the dsmsta.opt file, verify that the SERVERNAME as well as the high-level address (HLA) and low-level address (LLA) are set correctly in the DEFINE SERVER line.
2. If messages are seen on the server for this storage agent, do any error messages give an indication about why this storage agent is not working with the server?

Return to diagnosing a storage agent problem.

Check HELP for Tivoli Storage Manager messages issued for the problem.

Check HELP for any messages issued by Tivoli Storage Manager.

Tivoli Storage Manager messages provide additional information beyond just the message itself. The **Explanation, System Action,** or **User Response** sections of the message may provide additional information about the problem. Often, this supplemental information about the message may provide the necessary steps to resolve the problem.

Return to diagnosing a storage agent problem.

Is the problem an error reading or writing to a device?

If the problem is an error reading or writing data from a device, many systems and devices record information in a system error log. Examples of the system error log are errpt for AIX(R), and the Event Log for Windows(R).

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

Storage agents are particularly vulnerable if path information is changed or not correct. Issue QUERY PATH F=D on the server. For each of the storage agent's paths, verify that the settings are correct. In particular, verify that the device listed matches the system device name. If the path information is not correct, update the path information with the UPDATE PATH command.

Return to diagnosing a storage agent problem.

Have any storage agent options been changed?

Changes to options in the storage agent option file may cause operations that previously succeeded to fail.

Review any changes to the storage agent option file. Try reverting the settings back 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 causes the failure has been identified.

Return to diagnosing a storage agent problem.

Have any server options or settings been changed?

Changes to options in the server option file or changes to server settings using SET commands may affect the storage agent.

Review any changes to server option settings. Try reverting the settings back 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 causes the failure has been identified.

Review server settings by issuing: QUERY STATUS. If any settings reported by this query have changed, review the reason for the change, and if possible, revert it back to the original value and retry the storage agent operation.

Return to diagnosing a storage agent problem.

Main Menu | LAN-free setup

LAN-free setup.

Click the symptom in the table or skip the table and scroll through the information below:

List of LAN-free setup symptoms

- Data is sent directly to the server.
- The z/OS LAN-free session failed.
- My storage pool is configured for simultaneous write but will not work LAN-free.
- SHOW LAN-free - Validating LAN-free definitions on the server.
- Testing LAN-free configuration.

Data is sent directly to the server.

The client summary statistics do not report any bytes transferred LAN-free. The client will report the bytes sent LAN-free using "ANE4971I LAN-free Data Bytes: xx KB". Similarly, the server does not report any instance of "ANR0415I Session SESS_NUM proxied by STORAGE_AGENT started for node NODE_NAME" for this node and storage agent indicating that the LAN-free proxy operation was done for this client node.

The client will only attempt to send data LAN-free using the storage agent if the primary storage pool destination in the server storage hierarchy is LAN-free. A server storage pool is LAN-free enabled for a given storage agent if one or more paths are defined from that storage agent to a SAN device. To determine if the storage pool destination is configured correctly, do the following:

- Issue QUERY NODE *nodeName*. This will report the policy domain to which this node is assigned.
- Issue QUERY COPYGROUP *domainName policySetName mgmtclassName* F=D for the management classes that this node would use from their assigned policy domain. Note that this reports information for backup files. To query copy group information for archive files issue: QUERY COPYGROUP domainName policySetName mgmtclassName TYPE=ARCHIVE F=D.
- Issue QUERY STGPOOL *stgpoolName* where *stgpoolName* is the destination reported from the previous QUERY COPYGROUP queries.
- Issue QUERY DEVCLASS *deviceClassName* for the device class used by the destination storage pool.
- Issue QUERY LIBRARY *libraryName* for the library reported for the device class used by the destination storage pool.
- Issue QUERY DRIVE *libraryName* F=D for the library specified for the device class used by the destination storage pool. If no drives are defined to this library, review the library and drive configuration for this server and use DEFINE DRIVE to define the needed drives. If one or more of the drives report "On line=No", evaluate why the drive is offline, and if possible, update it to online, using UPDATE DRIVE *libraryName driveName* ONLINE=YES.
- Issue QUERY SERVER to determine the name of the storage agent as defined to this server.
- Issue QUERY PATH *stgAgentName* where *stgAgentName* was the name of the storage agent defined to this server and reported in the QUERY SERVER command. Review this output and verify that one or more paths are defined for drives defined for the device class used by the destination storage pool. If no paths are defined for this storage pool, use DEFINE PATH to define the needed paths. Also, review this output and verify that the path is online. If paths are defined but no paths are online, update the path to online using UPDATE PATH *srcName destName* SRCTYPE=SERVER DESTTYPE=DRIVE ONLINE=YES.

Return to list of LAN-free setup symptoms

The z/OS LAN-free session failed.

If a LAN-free session fails, ensure that the storage pool contains volumes that have been formatted for LAN-free operations. To format volumes automatically, set the value of the LFVOLUMEFORMATCOUNT server option from 1 to 255. Volumes are formatted at the time the storage agent makes a request if the storage pool does not contain any formatted volumes. Using the command FORMAT LFVOLUME, you can format volumes before LAN-free sessions begin and have them available in a storage pool.

Return to list of LAN-free setup symptoms

My storage pool is configured for simultaneous write but will not work LAN-free.

The server disqualifies a storage pool as being a LAN-free enabled storage pool if it has been configured for simultaneous write. In this case, data from the client will be sent directly to the server and will not be using LAN-free.

Issue QUERY STGPOOL *stgpoolName* F=D for the destination storage pool for this client. If the storage pool is set for simultaneous write operations, the "Copy Storage Pool(s):" value will reference one or more other storage pool names. If this is the case, Tivoli Storage Manager interprets the simultaneous write operation to be higher priority than LAN-free data transfer. Because simultaneous write is considered a higher priority operation, this storage pool will not be reported as LAN-free enabled and as such the client will send the data directly to the server. The storage agent does not support simultaneous write operations.

Return to list of LAN-free setup symptoms

SHOW LAN-free and VALIDATE LAN-free - Validating LAN-free definitions on the server.

It can be difficult to determine if all the appropriate definitions are defined on the server to allow a given client node to perform LAN-free operations with a storage agent. Similarly, if the configuration has changed, it may be difficult to determine if the definitions on the server for a previously configured LAN-free client are still appropriate.

Version 5.2.2 of the Tivoli Storage Manager server provides a SHOW LANFREE command. Version 5.3.0 of the Tivoli Storage Manager server provides the equivalent VALIDATE LANFREE command. These commands will evaluate the destination storage pools for the domain to which this client node is assigned. The policy destinations are evaluated for BACKUP, ARCHIVE, and SPACEMANAGED operations for this node. For a given destination and operation, SHOW LANFREE and VALIDATE LANFREE will report whether or not it is LAN-free capable, if it is not LAN-free capable, they will give an explanation about why this destination can not be used.

The commands SHOW LANFREE and VALIDATE LANFREE are available from the server. The syntax for these command are:

SHOW LANFREE *nodeName stgAgentName*
VALIDATE LANFREE *nodeName stgAgentName*

where

nodeName
A client node registered to the server.
stgAgentName
A storage agent defined to the server.

An example of the output from this command is:

tsm: SRV1>q lanfree fred sta1
ANR0391I Evaluating node FRED using storage agent STA1 for LAN-free data movement.
ANR0387W Node FRED has data path restrictions.

Node Name	Storage Agent	Operation	Mgmt Class Name	Destination Name	LAN-Free capable?	Explanation
NODE1	STA1	BACKUP	NOLF	OUTPOOL	No	No available online paths.
NODE1	STA1	BACKUP	LF	LFPOOL	Yes	
NODE1	STA1	BACKUP	NOLF_SW	PRIMARY	No	Destination storage pool is configured for simultaneous write.
NODE1	STA1	ARCHIVE	NOLF	OUTPOOL	No	No available online paths.
NODE1	STA1	ARCHIVE	LF	LFPOOL	Yes	
NODE1	STA1	ARCHIVE	NOLF_SW	PRIMARY	No	Destination storage pool is configured for simultaneous write.

ANR1706I Ping for server 'STA1' was able to establish a connection.
ANR0392I Node FRED using storage agent STA1 has 2 storage pools capable of LAN-free data movement and 4 storage pools not capable of LAN-free data movement.

Return to list of LAN-free setup symptoms

Testing LAN-free configuration.

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 may 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 UPDATE NODE *nodeName* DATAWRITEPATH=LAN-free for the client node whose LAN-free configuration you want to test. Next, try a data store operation such as backup or restore. If the client and storage agent attempt to send the data directly to the server using the LAN, the following error message will be received:

ANR0416W Session *sessionNumber* for node *nodeName* not allowed to *operation* using *path* data transfer path.

The *operation* reported will indicate either READ or WRITE depending upon the operation attempted and the path will be reported as LAN-free.

If this message is received when trying a LAN-free operation, evaluate and verify the LAN-free settings. Generally, if data is not sent LAN-free when the client is configured to use LAN-free, the storage pool destination for the policy assigned to this node is not a LAN-free enabled storage pool, or the paths are not defined correctly.

Return to list of LAN-free setup symptoms

Main Menu | Data storage diagnostic tips

Data storage diagnostic tips

For a problem storing or retrieving data, review these steps to try to isolate or resolve the problem:

Diagnosing a problem storing or retrieving data

1. Check the server activity log.
2. Check HELP for messages issued for the problem.
3. Is the problem recreatable?
4. Is the problem an error reading or writing to a device?
5. Review the hints and tips for drives and libraries.
6. Has the storage hierarchy been changed?
7. Have the server policies been changed?
8. Does the problem occur only for a specific volume?
9. Does the problem occur only when using a specific drive to read or write to a volume?

Check the server activity log

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

Return to diagnosing a problem storing or retrieving data

Check HELP for messages issued for the problem

Check HELP for any messages issued by Tivoli Storage Manager.

TSM messages provide additional information beyond just the message itself. The **Explanation, System Action,** or **User Response** sections of the message may provide additional information about the problem. Often, this supplemental information about the message may provide the necessary steps to take to resolve the problem.

Return to diagnosing a problem storing or retrieving data

Is the problem recreatable?

If a problem can be easily or consistently recreated, it may be possible to isolate the cause of the problem to a specific sequence of events. Data read or write problems may be sequence-related in terms of the operations being performed or may be an underlying device error or failure.

Typical problems relating to the sequence of event occur for sequential volumes. One example would be that a volume is in use for a client backup and that volume is preempted by a restore of another client node's data. In this case, this may surface as an error to the client backup session that was preempted. However, that client backup session may succeed if it was retried or if it had not been preempted in the first place.

Return to diagnosing a problem storing or retrieving data

Is the problem an error reading or writing to a device?

If the problem is an error reading or writing data from a device, many systems and devices record information in a system error log. Examples of the system error log are errpt for AIX(R), Event Log for Windows(R), and the System Log for z/OS(R).

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

Return to diagnosing a problem storing or retrieving data

Review the hints and tips for drives and libraries.

The hints and tips for drives and libraries offer many items to check to insure that data read and write operations are successful.

Review the hints and tips for: Tape drives and libraries.

Return to diagnosing a problem storing or retrieving data

Has the storage hierarchy been changed?

The storage hierarchy includes the defined storage pools and the relationships between the storage pools on the server. These storage pool definitions are also used by the storage agent. If attributes of a storage pool have been changed this may affect data store and retrieve operations.

Review any changes to the storage hierarchy and storage pool definitions. Use QUERY ACTLOG to see the history of commands or changes that may affect storage pools. Also, use the following QUERY commands to determine if any changes have been made:

QUERY STGPOOL F=D

Review the storage pool settings. If a storage pool is UNAVAILABLE, then data in that storage pool cannot be accessed. If a storage pool is READONLY, then data cannot be written to that pool. If either of these is the case, review why these values were set and consider issuing UPDATE STGPOOL to set the pool to READWRITE. Another consideration is the number of scratch volumes available for a sequential media storage pool.

QUERY DEVCLASS F=D

The storage pools can be influenced by changes to device classes. Review the device class settings for the storage pools. This may also include checking the library, drive, and path definitions using QUERY LIBRARY, QUERY DRIVE, and QUERY PATH for sequential media storage pools.

Return to diagnosing a problem storing or retrieving data

Have the server policies been changed?

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. Use QUERY ACTLOG to see the history of commands or changes that may affect storage policies. Also, use the following QUERY commands to determine if any changes have been 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 have been 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 may impact data read and write operations for that client. Specifically, the node may now be going to storage pool destinations that are not appropriate based on the requirements of this node. For example, it may be assigned to a domain that does not have any TYPE=ARCHIVE copy group destinations. If this node tries to archive data, it will fail.

Return to diagnosing a problem storing or retrieving data

Does the problem occur only for a specific volume?

If problems occur only for a specific storage volume, there may be an error with the volume itself. This is true whether the volume is sequential media or DISK.

If this is a data write operation, use 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 UPDATE VOLUME *volumeName* ACCESS=READWRITE and again retry the operation. If the operation fails only when using this volume, consider using AUDIT VOLUME to evaluate this volume and MOVE DATA to move the data from this volume to other volumes in the storage pool. After the data is moved off this volume, delete the volume using DELETE VOLUME

Return to diagnosing a problem storing or retrieving data

Does the problem occur only when using a specific drive to read or write to a volume?

For sequential media volumes, does the error occur only when using a specific drive?

If the error occurs only when using a specific drive, update that drive to be offline using UPDATE DRIVE *libraryName driveName* ONLINE=NO. Then take the appropriate steps to correct the drive error or else have the vendor for the drive service the drive. After the drive has been corrected, use the UPDATE DRIVE *libraryName driveName* ONLINE=YES to set the drive online again.

Return to diagnosing a problem storing or retrieving data

Main Menu | SAN devices| SAN Device Mapping Problems | SAN Device Mapping Hints and Tips

SAN devices.

Click the symptom in the table or skip the table and scroll through the information below:

List of SAN device symptoms

- 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 recommended action.
- ANR8957E: *command*: Autodetect is OFF and the serial number reported by the library did not match the serial number in the library definition.
- ANR8958E: *command*: Autodetect is OFF and the serial number reported by the drive did not match the serial number in the drive definition.
- ANR8963E: Unable to find path to match the serial number defined for drive *driveName* in library *libraryName*.
- ANR8965W: The server is unable to automatically determine the serial number for the device.
- ANR8972E: Unable to find element number for drive *driveName* in library *libraryName*.
- For additional error messages related to SAN, try the SAN Device Mapping Problems page.

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 recommended action.

For SAN device errors, this message will be issued in many cases. The CC=205 reports that the device driver has detected a SCSI adapter error. In the case of a SAN-attached device that encountered a link reset caused by link loss, it will be 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 should be updated to ONLINE=NO using the UPDATE PATH command. Do not set the path to ONLINE=YES until the cause for the link reset has been isolated and corrected.

Return to list of SAN device symptoms

ANR8957E: *command*: Autodetect is OFF and the serial number reported by the library did not match the serial number in the library definition.

Tivoli Storage Manager SAN Device Mapping encountered a path for the library that reports a different serial number than the current Tivoli Storage Manager definition for the library. The AUTODETECT parameter was set to NO for the command which prevented the server from updating the serial number for the library.

For Tivoli Storage Manager 5.2 on Windows with systems using QLogic host bus adapters, reissue the command with AUTODETECT=ON. The newly discovered serial number will automatically be updated for this library. The following informational message will be issued:

ANR8953I: Library *libraryName* with serial number *serialNumber* is updated with the newly discovered serial number *newSerialNumber*

For other Tivoli Storage Manager servers, determine the new path and issue UPDATE PATH to correct this.

Return to list of SAN device symptoms

ANR8958E: *command*: Autodetect is OFF and the serial number reported by the drive did not match the serial number in the drive definition.

Tivoli Storage Manager SAN Device Mapping encountered a path for a drive that reports a different serial number than the current Tivoli

Storage Manager definition for that drive. The AUTODETECT parameter was set to NO for the command which prevented the server from updating the serial number for this drive.

For Tivoli Storage Manager 5.2 on Windows with systems using QLogic host bus adapters, reissue the command with AUTODETECT=ON. The newly discovered serial number will automatically be updated for this drive. The following informational message will be issued:

ANR8955I: Drive *driveName* in library *libraryName* with serial number *serialNumber* is updated with the newly discovered serial number *newSerialNumber*.

For other Tivoli Storage Manager servers, determine the new path and issue UPDATE PATH to correct this.

Return to list of SAN device symptoms

ANR8963E: Unable to find path to match the serial number defined for drive *driveName* in library *libraryName*.

For Tivoli Storage Manager 5.2 Windows servers and storage agents, the server will automatically map SAN devices. This SAN Device Mapping feature allows the server to compare its definitions for SAN devices to the actual definition for these devices. If a discrepancy is detected, the server definition for this device is updated to correct this.

The SAN Device Mapping was not able to find a SAN device that was previously defined to the server. The most likely cause for this is that the device itself has been removed or replaced in the SAN. The possible steps to resolve this are:

Device Removed

If the device has been removed from the SAN, simply delete the server definitions that refer to this device. Use QUERY PATH F=D to determine any paths that reference the device. Then, issue DELETE PATH to remove these paths.

Device Replace

As a result of maintenance or an upgrade to a SAN device, it has been replaced with a new device. Please do the following:

- Try not to delete the drive or drive path definition after you replace the drive.
- Issue one of the following server command:
  - UPDate DRive <*libraryName*> <*driveName*> SERIAL=AUTODetect

    This will force the new serial number to be recorded into the server database. Since the drive was replaced, the element number should stay the same.

  - UPDate PATH <*sourceName*> <*driveName*> SRCT=SERVER DESTT=DRIVE LIBRary=<*libraryName*> DEVIce=xxxxx AUTODetect=Yes

    This will force the new serial number to be recorded into the database. Since the drive was replaced, the element number should stay the same.

- If the drive or drive path has been deleted, you should completely redefine this new, replaced drive. You will have to restart the Tivoli Storage Manager server so that the element number / serial number map for the library can be refreshed, because this mapping only occurs at initialization.

Use QUERY PATH F=D to find any paths defined on the server that reference this device and issue the following command to update the path information:

UPDATE PATH AUTODetect=Yes

Return to list of SAN device symptoms

ANR8965E: The server is unable to automatically determine the serial number for the device.

For a Tivoli Storage Manager 5.2 non-Windows server or a Tivoli Storage Manager 5.3 server with SANDISCOVERY option OFF, the Tivoli Storage Manager server will not automatically detect SAN device serial numbers.

If the Tivoli Storage Manager server detects that a device serial number has changed, it will mark the device's path as offline. To correct this:

1. Determine the correct serial number for the device.
2. Use the UPDATE DRIVE or UPDATE LIBRARY to update the device serial number.

Return to list of SAN device symptoms

ANR8972E: Unable to find element number for drive *driveName* in library *libraryName*.

If the ELEMent parameter was set to AUTODetect when defining the drive, Tivoli Storage Manager will try to get the drive's element number automatically. However, if the library does not provide element number/serial number map, this message will be issued.

To correct this:

1. Determine the element number for this tape drive.
2. Use the UPDATE DRIVE command to update the device element number.

Return to list of SAN device symptoms

Main Menu | SCSI devices

SCSI devices

Click the symptom in the table or skip the table and scroll through the information below:

List of SCSI device symptoms
• Did Tivoli Storage Manager issue message ANR8300, ANR8301, ANR8302, ANR8303, ANR8943, or ANR98944?

Did Tivoli Storage Manager issue message ANR8300, ANR8301, ANR8302, ANR8303, ANR8943, or ANR98944?

Tape drives and libraries may report information back to Tivoli Storage Manager about the error encountered. This information is reported in one or more of the messages listed. The data Tivoli Storage Manager reports from these devices may provide sufficient detail to determine the steps necessary to resolve the problem. Generally, when the Tivoli Storage Manager server reports device sense data using these messages, the problem is typically with the device, the connection to the device, or some other related issue outside of Tivoli Storage Manager.

Using the information reported in Tivoli Storage Manager message ANR8300, ANR8301, ANR8302, ANR8303, ANR8943, or ANR8944, refer to the Appendix C of the *IBM Tivoli Storage Manager Messages* manual. This appendix documents information about standard errors that may 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.

Return to list of SCSI device symptoms

Main Menu | SAN Device Mapping Hints and Tips | SAN Device Mapping Problems| SAN devices

SAN Device Mapping Problems

If additional help not discussed below is needed, try the SAN Device Mapping Hints and Tips page.

Click the symptom in the table or skip the table and scroll through the information below:

Diagnosing SAN device mapping
• Some of my SAN devices are missing from QUERY SAN. • ANR1745I Unable to discover SAN devices. Function is busy. • ANR1786W, ANR1787W or ANR1788W messages were issued. • ANR1789W Get HBA target mapping failed. • ANR1790W SAN discovery failed. • ANR1791W HBAAPI wrapper library xxxxx failed to load or is missing. • ANR1792W HBAAPI vendor library failed to load or is missing. • ANR1793W Tivoli Storage Manager SAN discovery is not supported on this platform or this version of OS. • ANR1794W Tivoli Storage Manager SAN discovery is disabled by options. • ANR2034E QUERY SAN: No match found using this criteria. • ANR8226E Error detecting version of HBA-API library. • ANR8227E Fileset devices.common.IBM.fc.hba-api is not at the required level.

Some of my SAN devices are missing from the display of QUERY SAN server command.

The following are possible reasons why the QUERY SAN server command does not display all the devices.

SAN configuration needs to be refreshed.

The configuration has been changed (add/remove device), and the system configuration needs to be updated.

- Update configuration on AIX:
  - For IBM devices:

    Issue cfgmgr command to get the new configuration and the special file name for IBM tape devices (not Tivoli Storage Manager devices) will be added or removed (drv_add and drv_rem).

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

  - For Tivoli Storage Manager devices:

    One way to get the special files updated is by using smitty->devices->Tivoli Storage Manager Devices-> remove all defined devices, then "discover devices supported by Tivoli Storage Manager."

    Special file name: /dev/mt0, /dev/lb0

  - Alternatively, reinstalling the IBM device driver/Tivoli Storage Manager device driver will update all the current special file name.
- Update configuration on Windows:

  With the plug and play, the Windows registry will be updated and the device name may be changed without the need for reboot the machine or device driver's involvement. The Tivoli Storage Manager server can detect the change in a special file name and update the new special file name when it access the tape devices (during server initialization or normal operation). The correct device name will be updated in the Tivoli Storage Manager database.

  Special file name: mtx.x.x.x, lbx.x.x.x (can be used for both Tivoli Storage Manager devices and IBM devices) Tape0, Changer0 (only for IBM devices)

- Update configuration on Linux:

  The HBA will get the most up to date configuration information as result of the RSCN. Sometimes, the machine has to

be rebooted to pick up the configuration changes.

- For IBM devices:

  Issue the IBMtapeconfig command to get the new configuration. The special file name for IBM tape devices (not Tivoli Storage Manager devices) will be added or removed.

  Special file name: /dev/IBMTape0, /dev/IBMChanger0

- For Tivoli Storage Manager devices:

  Users can issue Tivoli Storage Manager device driver auto configure script autoconf, which resides in /opt/tivoli/tsm/devices/bin to get all the current special file names.

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

- Alternatively, reinstalling the IBM device driver/Tivoli Storage Manager device driver will update all the current special file names.

With Linux pass-thru device driver for Tivoli Storage Manager devices, the HBA driver and the generic driver have to be reloaded to get all the current special file names. You have to run the autoconf script so the Tivoli Storage Manager device driver can create configuration files /dev/tsmscsi/lbinfo and /dev/tsmscsi/mtinfo. These files will be used by the Tivoli Storage Manager server to create the special file names after each SAN discovery.

- 32 bits (Linux xSeries)

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

- 64 bits (Linux pSeries)

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

- Update the configuration on Solaris:

The HBA will get the most up-to-date configuration information as result of the RSCN. Most of the time, the machine has to be rebooted to pick up the configuration changes. Reinstall the IBM device driver/Tivoli Storage Manager device driver or issue the following commands to have the special file name updated.

- For IBM devices:

  IBM device driver has already done the rem_drv and add_drv and nothing needs to be done.

  Special file name: /dev/rmt/0st , /dev/rmt/0smc

- For Tivoli Storage Manager devices:
  - modify /usr/kernel/drv/mt.conf and /usr/kernel/drv/lb.conf to add the following:

    - mt.conf:

      name="mt" parent="pseudo" instance=16383;

    - lb.conf

name="lb" parent="pseudo" instance=16383

- Issue the following commands:
  - rem_drv lb
  - add_drv lb
  - rem_drv mt
  - add_drv mt add_drv -m '* 0666 bin bin' mt

  Alternatively, the **autoconf** command can do the above by itself.

  Make sure that the virtual device special file is in /devices/pseudo/mt@16383:tsmmtctl and /devices/pseudo/lb@16383:tsmlbctl

- Issue the command twice
  - sync
  - sync
- Reboot the machine

The following files are used by the Tivoli Storage Manager server to create the special filenames after each SAN discovery.

- 32 bits (Linux xSeries)

  make sure libhbaapi32.so is in the same directory as dsmserv.exe or in /opt/tivoli/tsm/server/bin

- 64 bits (Linux pSeries)

  make sure libhbaapi64.so is in the same directory as dsmserv.exe or in /opt/tivoli/tsm/server/bin

Configuration problem (HBA hardware, HBA driver level, OS level)

Please go to http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html to verify the platform/HBA vendor/driver level support level to make sure the HBA driver level and OS level are compatible and supported by Tivoli for SAN discovery.

Use the HBA vendor utility to check if the device can be seen by the HBA. If the device is not seen by HBA, the device may 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 may have problems with the HBA API.

Vendor support for any particular device in the SAN

Many devices or combinations of devices may not be supported in a given 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. This includes all hardware associated with the SAN, which means verifying this device is supported with the vendors of the HBAs, hubs, gateways, and switches that make up the SAN environment.

Return to diagnosing SAN device mapping

ANR1745I Unable to discover SAN devices. Function is busy.

This occurs if there is another active SAN discovery.

RESULTS: The Tivoli Storage Manager server will not be able to perform SAN discovery.

Retry again after the other SAN discovery is completed.

Return to diagnosing SAN device mapping

ANR1786W, ANR1787W or ANR1788W messages were issued.

These 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*.

RESULTS: The Tivoli Storage Manager server will not be able to perform SAN discovery.

Please go to http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html to verify that the HBA driver is up-to-date and at a supported level.

Return to diagnosing SAN device mapping

ANR1789W Get HBA target mapping failed .

This is the most common error on the SAN with HBAAPI. This means HBA encountered an error while gathering device mapping information by sending various SCSI commands.

Verify that all SAN device are working properly (e.g. a SAN Data Gateway could be hung and need to be rebooted). If all devices appear functional, verify the firmware of device on the SAN as well as HBA driver are at appropriate levels.

RESULTS: The Tivoli Storage Manager server will not be able to perform SAN discovery.

Please go to http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html to verify that the HBA driver is up-to-date and at a supported level.

For IBM tape devices, make sure the latest firmware is installed. Firmware prior to 4772 for IBM 3580 tape devices caused problems with Qlogic HBAAPI.

Return to diagnosing SAN device mapping

ANR1790W SAN discovery failed.

This message is a general message indicating HBAAPI function failed and we will not be able to perform SAN discovery.

Verify that all SAN device are working properly (e.g. a SAN Data Gateway could be hung and need to be rebooted). If all devices appear functional, verify the firmware of device on the SAN as well as HBA driver are at appropriate levels.

For IBM tape devices, make sure the latest firmware is installed. Firmware prior to 4772 for IBM 3580 tape devices causes problems with Qlogic HBAAPI.

RESULTS: The Tivoli Storage Manager server will not be able to perform SAN discovery.

Return to diagnosing SAN device mapping

ANR1791W HBAAPI wrapper library xxxxx failed to load or is missing.

The HBAAPI wrapper libraries is the wrapper used by the Tivoli Storage Manager server to communicate with the SNIA HBAAPI. The HBAAPI wrapper libraries in the same directory as the Tivoli Storage Manager executable (unless full path is given as shown below). The following is a list of the HBA wrapper files shipped with the Tivoli Storage Manager server package (except on AIX). This warning message indicates that the HBAAPI wrapper file is either missing or could not be loaded by Tivoli Storage Manager. Please verify that the

wrapper file is in the same directory as the Tivoli Storage Manager executable. The HBAAPI wrapper library files are:

- Windows(R): hbaapi.dll
- AIX(R): /usr/lib/libhbaapi.a (This is provided by AIX with HBAAPI installation)
- 32-bit Linux: libhbaapi32.so
- 64-bit Linux: libhbaapi64.so
- 32-bit Solaris: libhbaapi32.so
- 64-bit Solaris: libhbaapi64.so

RESULTS: The Tivoli Storage Manager server will not be able to perform SAN discovery.

Return to diagnosing SAN device mapping

ANR1792W HBAAPI vendor library failed to load or is missing.

This message indicates that vendor's library file failed to be loaded. Please verify validity of these files.

UNIX(R) systems store their HBAAPI libraries in the location specified by the **/etc/hba.conf file**, and Window's file are stored in the C:\winnt \system32 directory.

The following are examples of vendor library files:

- C:\winnt\system32\qlsdm.dll (QlLogic's(R) Windows file)
- /usr/lib/libHBAAPI.a (Emulex's(R) AIX file)
- /usr/lib/libqlsdm.so (Qlogic's Linux file)
- /usr/lib/libemulexhbaapi.so (Emulex's Linux 32-bit file)
- /usr/lib64/libemulexhbaapi.so (Emulex's Linux 64-bit file)
- /usr/lib/libqlsdm.so (Qlogic's Solaris file)
- /opt/JNIsnia/Solaris/Jni/64bit/JniHbaLib.so (JNI's(R) Solaris file)

RESULTS: The Tivoli Storage Manager server will not be able to perform SAN discovery.

Return to diagnosing SAN device mapping

ANR1793W Tivoli Storage Manager SAN discovery is not supported on this platform or this version of OS.

This message will only be displayed if Tivoli Storage Manager attempted a SAN device mapping or device discovery operation on an unsupported operating system.

The following platforms are not currently supported by SAN device mapping or device discovery:

- HP-UX
- 64-bit Windows(R) 2003
- Linux zSeries
- 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.

RESULTS: The Tivoli Storage Manager server will not be able to perform SAN discovery.

Return to diagnosing SAN device mapping

ANR1794W Tivoli Storage Manager SAN discovery is disabled by options.

The message indicate that the SAN discovery on Tivoli Storage Manager server is disabled.

RESULTS: The Tivoli Storage Manager server will not be able to perform SAN discovery.

The SAN discovery can be disabled by the following server command:

- setopt SANDISCOVERY OFF

  This will disable the SAN discovery. The Tivoli Storage Manager server will not be able to correct the device path automatically is the path has been changed. This only needed to be issued once.

- setopt SANDISCOVERY ON

  This will enable the SAN discovery.

The above setopt SANDISCOVERY command can be issue as many times as needed.

Another way to disable/enable SAN discovery is to put the following option in the dsmserv.opt file

- SANDISCOVERY OFF

  This will disable the SAN discovery.

- SANDISCOVERY ON

  This will enable the SAN discovery.

SANDISCOVERY is default to ON for Windows platform.

SANDISCOVERY is default to OFF for all non-Windows platforms.

It is recommended you go to http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html to verify the platform/HBA vendor/driver level support level prior to setting SANDISCOVERY ON to enable the SAN discovery.

Return to diagnosing SAN device mapping

ANR2034E QUERY SAN: No match found using this criteria.

The Tivoli Storage Manager server is trying to collect configuration information for the SAN and found nothing.

RESULTS: The Tivoli Storage Manager server will not be able to perform SAN discovery.

There are a number of possible reasons for not being able to find information about the SAN. These possibilities are:

- This is a not support platform or OS level.
- This is not a SAN environment.
- There maybe a problem with the SAN.
- Check FC HBA driver and make sure it is installed and enabled.
- Check the HBA driver level to make sure that it is up to date.
- Use the HBA vendor's utility to check for any reported FC link problem.
- Uninstall and then install the HBA driver. If there is an issue with the HBA configuration, device driver, or compatibility, sometimes uninstalling and re-installing it will correct the problem.
- Check the FC cable connection to the HBA.
- Check the FC cable connection from the HBA to the SAN device (switch, data gateway, or other device).
- Check the GBIC.
- On the SAN device (switch, data gateway, or other device) try a different target port. Sometimes the SAN devices may have a specific port failure.

- Halt the Tivoli Storage Manager server, reboot the machine, and restart the server. If there have been configuration changes in the SAN, sometimes the operating system, device driver, or HBA will require a reboot of the machine before they can communicate with the SAN.
- Recycle the destination port on the SAN device.
- Reseat the HBA card.
- Replace the HBA.

Return to diagnosing SAN device mapping

ANR8226E Error detecting version of HBA-API library.

This is for AIX only. The server attempted to determine the level of the devices.common.IBM.fc.hba-api fileset and encountered an error.

This message indicates that encountered error to detect the HBA-API libraryFileset version on AIX.

RESULTS: The Tivoli Storage Manager server will not be able to perform SAN discovery.

Return to diagnosing SAN device mapping

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

Due to problems in AIX HBAAPI code, the minimum fileset devices.common.IBM.fc.hba-api level needed for successful SAN discovery is the following:

- AIX52 - Need 5.2.0.50
- AIX53 - Need 5.3.0.10

This is for AIX only. It specified file set devices.common.IBM.fc.hba-api is at a level that is incompatible with Tivoli Storage Manager operations. Install the latest maintenance for this file set if you use SAN devices.

RESULTS: The Tivoli Storage Manager server will not be able to perform SAN discovery.

Return to diagnosing SAN device mapping

Main Menu | Sequential media volume (tape)

Sequential media volume (tape).

Click the symptom in the table or skip the table and scroll through the information below:

List of sequential media volume (tape) symptoms
• ANR0542W Retrieve or restore failed for session *sessionNumber* for node *nodeName* - storage media inaccessible. • ANR8778W Scratch volume changed to private status to prevent re-access.

ANR0542W Retrieve or restore failed for session *sessionNumber* for node *nodeName* - storage media inaccessible.

This is often an issue with the drive or connection to the drive that was selected to read this tape volume. To verify that Tivoli Storage Manager can access this volume, do the following:

- Issue QUERY LIBVOL *libraryName volumeName*.
- For a 349X library, issue **mtlib -l /dev/lmcp0 -qV *volumeName***. The device is typically /dev/lmcp0 but if it is different then substitute the correct library manager control point device.

The possible steps to resolve this are:

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 and then run AUDIT LIBRARY *libraryName* on the server.
2. If the volume is not reported by QUERY LIBVOL then the server does not know about this volume in the library. Issue AUDIT LIBRARY *libraryName* to synchronize the library inventory in the server with what is actually in the tape library.
3. If both commands successfully reported this volume, then the cause is likely a permanent or intermittent hardware error. This may 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.

Return to list of sequential media volume (tape) symptoms

ANR8778W Scratch volume changed to private status to prevent re-access.

Review the activity log messages to determine the cause of the problem using 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 was caused by a drive requiring cleaning, or other hardware-specific issue that has been resolved, any volumes that were set to private status as a result of this may be reset to scratch by issuing AUDIT LIBRARY *libraryName*.

Return to list of sequential media volume (tape) symptoms

Main Menu | Administration Center Support Utility

Administration Center Support Utility

The Administration Center Support Utility is a command-line driven utility designed to assist in performing basic support tasks with the Administration Center. This utility is located in the Integrated Solutions Console installation directory under `Tivoli\dsm\bin`. On Windows, start the tool by issuing `supportUtil.bat`. On Unix, start the tool by issuing `supportUtil.sh`.

```
C:\Program Files\IBM\ISC\Tivoli\dsm\bin>supportUtil.bat

Administration Center Support Utility - Main Menu
=================================================

1.  Turn all tracing on
2.  Turn all tracing off
3.  Turn a single trace class on
4.  Update the maximum memory size Administration Center can use
5.  Update the Administration Center session timeout setting
6.  Collect trace files, logs and system information to send to support
7.  View the log file for this utility

9.  Exit

Enter Selection:
```

Main Menu | Administration Center Task or Work Item Fails in an Unexpected Manner

Administration Center Task or Work Item Fails in an Unexpected Manner.

For a Administration Center problem, review these steps and possible causes to try to isolate or resolve the problem:

Diagnosing an Administration Center problem

1. Check the server activity log.
2. Restarting Wizards.
3. Working with Multiple Portlets.
4. Internal Errors.

Check the server activity log.

The Administration Center processes its work by issuing Tivoli Storage Manager query, select, and other commands to the server. Check the activity log of the server that was selected for the failing task. To view the activity log, select a server and select the Server Properties table action. In the properties notebook, click the Activity Log tab. Ensure that commands related to the task appear in the activity log and are being issued to the correct server. If there is no evidence of activity related to the Administration Center task, there may be a connection problem or the incorrect server might have been selected.

Look for failures related to commands that were issued in the time frame of the failing Administration Center task. These failures could be due to incorrect formatting of the command or invalid parameter values. In some cases, a parameter value comes directly from the Administration Center, but in many others it is specified by the user through a text box, radio button, or other input method. Correct the invalid value if possible, but if an invalid parameter or incorrectly constructed command cannot be changed by the user, this might be the result of an internal error.

Return to diagnosing an Administration Center problem

Restarting wizards.

Some issues have to do with the manner in which wizards or portlets are started or restarted. Some of the wizards commit objects to the server database as you proceed through the wizard, instead of only when you finish the wizard. When the wizard defines objects before you click Finish, a summary panel listing the committed objects is displayed. The Back button is disabled, because you cannot return to the previous wizard panel to undo the committed changes. If the wizard is then cancelled, the objects that have been created are not deleted. You cannot restart the wizard from the point at which it was cancelled.

If you cancel and then restart a wizard, you will not be able to duplicate the work done earlier. For example, if you use the "Add a Storage Device" wizard to add a tape device, you will encounter several such summary panels that show the objects that are created before the wizard is finished. The first is for the creation of a device class, a library and paths from the server to the library. If you cancel this wizard following the first summary panel and restart it, you will not be able to create another library with the same name, because it now already exists.

After an object is created with a wizard, any changes to that object must be made through a properties notebook. In the "Add a Storage Device" example, changes to the device class, library, or storage pools must be made through the device class, library, or storage pool properties notebooks. Typically, you can click an object's name to open its properties notebook. The properties notebook contains tabbed pages for modifying all of the attributes created by the wizard. Attributes that are set to the default by the wizard can be changed only in these properties notebooks.

Return to diagnosing an Administration Center problem

Working with multiple portlets.

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

refreshed.

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

Return to diagnosing an Administration Center problem

Internal Errors.

An internal error is an error in the Administration Center itself. This can be more difficult to diagnose. In most cases an error message is displayed. Another symptom is when an OK or Cancel button does not close the current panel. Finally, the server activity log might contain commands issued by the Administration Center that are improperly formatted or contain invalid parameter values.

These errors can only be accurately diagnosed by support personnel. In addition to the server activity log, the Integrated Solutions Console trace log can often provide information relating to the error. If the problem is repeatable, the tracing level of the trace log can be increased. The information in the trace log will likely be useful to support personnel. See tracing for more information about the Integrated Solutions Console trace log and setting the tracing level.

Return to diagnosing an Administration Center problem Main Menu | Administration Center Task Fails with a Message

Administration Center Task Fails with a Message.

For an Administration Center task that fails with a message, review these steps to resolve the problem:

Administration Center Task Fails With a Message
1. Handling tasks that failed with a message. 2. Administration Center messages versus Tivoli Storage Manager server messages. 3. Problems with the Tivoli Storage Manager server command definition file.

Handling tasks that failed with a message.

The Administration Center provides both informational and error messages when a failure occurs. Most failure messages will typically provide information about what went wrong. The Administration Center failure messages often explicitly diagnose the problem and include suggestions about how to handle it. The messages can contain information about Administration Center internal errors or server errors.

In some cases, errors can be resolved by retrying a failed task. Before retrying the task, close any error messages. This ensures that only relevant error messages will be displayed if the task fails again.

Some Administration Center failure messages include one or more error messages returned from the server. The server messages are displayed in the language enabled for the server, which might not match the language enabled for the Web browser. In such cases, consult the Message Manual for the message number identified for the specific server platform. Further distinction between Administration Center messages and Tivoli Storage Manager server messages is outlined in the sub topic Administration Center messages versus Tivoli Storage Manager server messages.

To better understand the problem, check the Activity Log for the server chosen when the failure message occurred. To view the activity log, select a server and select the Server Properties table action. In the properties notebook, click the Activity Log tab. This is very useful for cases where the failure message does not seem related to the behavior of the task, it does not include server messages at all, or only the

last server message is shown when multiple tasks have been performed.

The Administration Center often performs multiple tasks to implement a single interface action. For example, these tasks can include:

- Communication setup
- Validity checks
- Table construction

In these cases, multiple commands are typically issued to the server. Look carefully in the Activity Log for the activity at the time frame of the failure. Keep in mind that the server machine and the machine used to access the Administration Center might be running in different time zones. Activity log information is displayed using the server's time.

If errors are encountered during command routing, check the activity log of the source server first. If it does not provide sufficient information to diagnose the problem, check the activity log of the target server. Make sure that the two servers are running, they are set up correctly to communicate with each other, and the source server has appropriate authority level for the commands performed on the target server at the time of failure.

If you need to contact customer support to resolve an Administration Center issue, you can use the tracing features provided by the Administration Center Support Utility to obtain additional information. You can find further information about the activity log and Administration Center internal errors in Task fails in an Unexpected Manner.

Further information about server and Administration Center tracing is available in Tracing.

Return to Task or Work Item Fails with a Message

Administration Center messages versus Tivoli Storage Manager server messages.

Messages shown in the Administration Center are typically either Tivoli Storage Manager server messages, or messages specific to the Administration Center. Tivoli Storage Manager server messages are displayed in some cases when the Administration Center performed a command on a Tivoli Storage Manager server. In some cases, both types of messages are displayed. To distinguish between Administration Center messages and Tivoli Storage Manager server messages, observe the prefix of the message number. Messages that do not show any message number are Administration Center informational messages specific to the Integrated Solutions Console.

Administration Center messages have the prefix **ANRW**. Informational and error messages are distinguished by the suffix of the message, **I** versus **E**. For example:

ANRW0022I The operation completed successfully.

ANRW0023E An internal error occurred during validation.

Informational messages such as messages numbered ANRW0022I provide some general information or identify limitations or requirements for the task being performed.

Error messages such as messages numbered ANRW0023E identify that there is a problem completing the requested task. The message can prompt the user to perform additional setup (such as server-to-server-communication), identify that the requested action is invalid in the current work item or task, or report an internal error. Resolving an internal error, such as ANRW0023E, generally requires assistance from customer support. You will typically be asked to provide tracing information, as well as information about the sequence of actions performed before the failure occurred.

Messages issued by the Administration Center are displayed in the language selected for the Web browser, or its closest default. Administration Center message numbers are not server dependent, unlike Tivoli Storage Manager server messages. Administration Center messages are generally based on the return code specified when a Tivoli Storage Manager server command has been performed. You can find information about basic Administration Center messages here. In some cases the server return code that is the basis for the Administration Center message will be shown, along with the server message. In rare cases, the return code and message displayed will be inconsistent. If you use tracing to capture information for a server command execution problem, the return codes identified with the server messages can help customer support diagnose the issue.

Some Tivoli Storage Manager server messages are dependent on the server operating system. When the Administration Center displays a

Tivoli Storage Manager server message, the message number can be associated with different messages for servers running on different operating systems. A single message number might have different variations for different platforms.

Tivoli Storage Manager messages use different message prefixes, typically ANR or ANE. Understanding Tivoli Storage Manager Messages provides basic information about the different types of Tivoli Storage Manager server messages.

The following is a sample message that occurs if a DEFINE LIBRARY command fails on a z/OS server:

ANR2022E **DEFINE LIBRARY: One or more parameters are missing.**

Use the command line in the Administration Center or on the server machine where the failure occurred to obtain further information about the server message. To request help for the specific Tivoli Storage Manager server message number, issue the following command:

help ANR2022E

As an alternative, use the Messages Manual and identify the message by its message number. If the message is platform specific, locate the message for the server platform of interest. For z/OS, for example, you should consider the version of the message showing (S/390) right after the message number. For the sample server message above, locate message number ANR2022E in the Message Manual.

For further information about diagnosing server problems, see Server Diagnostic Tips

Return to Task or Work Item Fails with a Message

Problems with the Tivoli Storage Manager server command definition file.

The Administration Center uses command definition files, which specify the commands supported for each server operating system. There is a different command definition file for each server operating system. If this file has been corrupted on the Integrated Solutions Console system or on the server system, parts of the interface might not be able to validate creation of new objects and might be unable to correctly construct server commands. Administration Center messages that identify server command syntax errors or validation errors (for text entry or the selection of interface controls) can sometimes be caused by problems with the command definition files.

The Administration Center acquires a command definition file the first time it establishes a connection with a server. The next time the Administration Center contacts a server running on the same operating system, it uses the command file that was previously obtained for that operating system. If the command file is not available in the appropriate location (due to renaming or deletion) then any request following the disappearance of the file will trigger another transfer of the command file for that specific operating system.

Do not attempt to update the command definition file. If problems occur that could be related to the command definition file, verify that the appropriate command definition file is available to the Administration Center for the server's operating system. If this file does not exist, verify that the Tivoli Storage Manager server can communicate with the Administration Center API and that it has the platform specific command file and dsmcmd.xml available in the server directory hierarchy.

Though it is not recommended, there is nothing to prevent the command definition file from being modified or deleted on the Integrated Solutions Console system, or the Tivoli Storage Manager server system. Verify that the appropriate platform specific command file is available to the Administration Center for the platform of the server of interest. If not, verify that the Tivoli Storage Manager server can communicate with the Administration Center API and it has the platform specific command file and dsmcmd.xml available and valid in the server directory hierarchy. To verify that a connection with the server can be established, see Unable to establish a connection with a Tivoli Storage Manager server

For example, a failure occurs when a task is being performed on an AIX 5.3.0 server. When a connection is established to that server or a command needs to be generated or verified, the Administration Center looks for the command definition file. The platform specific command file is copied from the server to the following location in the Integrated Solutions Console Portal Server hierarchy:

PortalServer\shared\app\dsm\CmdFileCache\aix_5_3_0_0.xml

Other platforms will have similar file name construction, and are located in the same Portal Server directory. Any changes made to the command definition file can corrupt the XML defined objects in the file, which can cause Administration Center tasks to fail.

If you identify a corrupted command definition file and cannot recover the original command definition file, you can use a connection to another server running the same operating system to obtain a valid command definition file, after renaming the original corrupted file.

If other Administration Center .xml or .xsd files in the WEB-INF directory are corrupted, similar issues can occur. For example, the Administration Center might not be able to construct tables or launch portlets.

Return to Task or Work Item Fails with a Message

Return to Main Menu
Main Menu | topic title

Topic Title

Topic Discussion Main Menu | Tracing | Client

Tracing for the client

Enable tracing on the client or client API

1. Determine the trace classes to enable from the table class table below.
2. Enable the trace by adding the following to the client options file: traceflag <trace class name>.
   Please note that <trace class name> may be a comma delimited list of trace classes. For example, this could be entered as traceflag service,verbinfo,verbdetail.
3. Next, configure trace to begin and issue the trace messages to a file by adding the following to the client options file: tracefile <file name>.
4. Perform the operation that is causing the problem.

Please note that tracing may 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.

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 hung and not responding.	None.
VERBINFO	Collect information regarding the client-server protocol used by Tivoli Storage Manager.	To debug protocol violations, transaction processing errors, or in cases where the client is hung and not responding.	None.
VERBDETAIL	Detailed information regarding the client-server protocol used by Tivoli Storage Manager. This will display internal memory buffers containing the verbs sent and received by the client.	To debug session data corruption issues that may e caused by the network.	This generates a large amount of output.

Main Menu | Tracing | Enabling tracing for the Administration Center

Enabling tracing for the Administration Center

1. Start the Administration Center Support Utility. This utility is located in the Integrated Solutions Console directory under `Tivoli\dsm\bin`. Start the utility by typing `supportUtil.bat` on Windows or `supportUtil.sh` on Unix.
2. Choose which trace classes to enable or simply turn all tracing on. See Trace Classes for Administration Center for more information about the different trace classes. Follow the on screen instructions to enable tracing.

```
C:\Program Files\IBM\ISC\Tivoli\dsm\bin>supportUtil.bat
Administration Center Support Utility - Main Menu
==================================================
1.  Turn all tracing on
```

```
2.  Turn all tracing off
3.  Turn a single trace class on
4.  Update the maximum memory size Administration Center can use
5.  Update the Administration Center session timeout setting
6.  Collect trace files, logs and system information to send to support
7.  View the log file for this utility
9.  Exit
Enter Selection: 3
Administration Center Support Utility - Turn Single Class On
==============================================================
1.  ALL - all trace classes
2.  SNAPIN - this includes CONTROL, VIEWS, ACTIONS, UTIL, and PS
3.  CONTROL - traces control objects
4.  VIEWS - traces view objects
5.  ACTIONS - traces action objects
6.  UTIL - traces utility objects
7.  PS - traces Presentation Services extended objects
8.  API - traces communications with the TSM server
9.  HEALTH - traces the health monitor
10. TAGLIB - traces JSP tag library extensions
11. XMLVAL - traces the reading and writing of XML files
99. Return to main menu
Enter Selection: 2
```

3. Perform the operation that is causing the problem. This will cause a trace of the problem to be written to the trace file. The trace file is located in the Integrated Solutions Console installation directory under `PortalServer\log\trace.log`.
4. Disable tracing. Leaving tracing enabled may cause performance problems.
5. Run `Collect trace files, logs and system information to send to support`. This will package up the trace information collected and log files into a single file called `AdminCenterSupport.zip`. This file can then be sent to support.

Main Menu | Tracing | Enabling client trace

Enabling client trace

There are two methods of tracing available for the B/A client. The first method is to configure trace parameters prior to starting the B/A Client and the second more recent method is to enable tracing while the client is running. Choose which method of tracing to enable:

See Enable tracing on command-line or Enable tracing while client is running

Main Menu | Tracing | Enabling client trace on command-line

Enabling client trace on command-line

1. Determine the trace classes to enable. See Trace Classes for Backup/Archive Client for possible trace classes
.
2. Choose which trace classes to enable by adding the following to the client options file: **traceflags <trace class name>**

3. Use a minus sign (-) in front of a trace class to subtract the class from tracing. Make sure that the subtracted trace classes are placed at the end of the trace class list. For example, if you want to collect a SERVICE trace without the SESSION or SESSVERB classes, then specify:

   Correct: traceflags service,-session,-sessverb
   Incorrect: traceflags -session,-sessverb,service

   Note that <trace class name> may be a comma-delimited list of trace classes. For example, this could be entered as:

```
traceflags service,verbdetail
```

4. Choose the location of the trace messages output by adding the following to the client options file: **tracefile <file name>**.
   The *tracefile* name should be fully qualified, for example:

```
tracefile c:\service\trace.out   (Windows(R))
tracefile /home/spike/trace.out  (UNIX(R))
tracefile SYS:\PUBLIC\TRACE.OUT  (NetWare(R))
tracefile trace.txt (Macintosh(R))
```

5. You can also set a maximum size for the trace file between 1 and 4,294,967,295 MB by specifying the following in the client options file:

   tracemax <size in mb>

   If a maximum value is specified, the client will start writing information from the beginning of the trace file (that is, wrapping) when the trace reaches its maximum size. This 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
```

   Once a tracefile reaches the limit specified with tracemax, "Continued at beginning of file" is written to the end of the trace file, and tracing continues from the top of the file. The end of the tracefile is indicated with "END OF DATA". You can locate the end of the trace by searching for this string. If you specify a TRACEMAX size of 1001 or higher and TRACESEGSIZE is not specified, then the trace file will automatically be split into multiple segments of 1000 MB per segment (see TRACESEGSIZE discussion immediately below).

6. You can choose to let the client split the trace into smaller segments between 1 and 1,000 MB per segment by specifying the following in the client options file:

   tracesegsize <trace segment size in MB>

   This allows you to more easily manage very large amounts of trace data, avoiding the problems associated with zipping very large files and eliminating the need to use a separate "file splitter" utility. For example, to specify a trace segment size of 200 MB:

   tracesegize 200

   Trace file segments are named using the name specified with the tracefile option plus an extension indicating the segment number. For example, if you specify tracefile tsmtrace.out and tracesegsize 200 then the trace will be segmented into multiple separate files of no more than 200 MB each, with file names tsmtrace.out.1, tsmtrace.out.2, etc. Note that when specifying the segment size, do not use any "thousands" separators:

   Correct: tracemax 1000

   Incorrect: tracemax 1,000

   If you use the **tracesegsize** option, the trace file segments will be named using the **tracefile** option with an additional extension using the segment number. For example, "trace.out.1"

7. Perform the operation that exhibits the problem

   Note that tracing may also be configured and started by invoking the client from a command prompt and specifying the flags above. For example:

```
dsmc -traceflags=service,verbdetail -tracefile=tsmtrace.out -tracemax=2500 -
tracesegsize=200
```

Main Menu | Tracing | Enable client trace while the client is running

Enable client trace while the client is running

Background:

The Tivoli Storage Manager client's built-in trace facility is a very useful diagnostic aid when working with IBM technical support engineers. Historically, the steps for gathering a Tivoli Storage Manager client trace are as follows:

1. Stop the Tivoli Storage Manager client.
2. Configure the client options file with the desired trace options.
3. Restart the client and reproduce the problem.
4. Stop the client.
5. Remove the trace options from the client options file.
6. Send the resulting trace file in to IBM technical support for analysis.

Beginning with IBM Tivoli Storage Manager Version 5.3, you can also use the new **dsmtrace** utility to start, stop, and configure client tracing dynamically, without having to stop the client or modify the options file. Dynamic tracing is especially useful when you need to trace only the beginning of a long-running client operation, or when you need to start tracing after the client has been running for some time.

The **dsmtrace** utility includes the following features:

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

The table below summarizes the availability of this feature:

Client Component	Unix Program Name	Windows Program Name
Backup-Archive Client (command line)	dsmc	dsmc.exe
Backup-Archive Client (GUI)	N/A	dsm.exe
Client Acceptor Daemon	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

Notes:

- Dynamic tracing is currently unavailable for NetWare clients.
- The Unix column above includes Mac OS X.
- Tracing for the Data Protection components is for the Tivoli Storage Manager API only.
- Tivoli Storage Manager API tracing is available with any multithreaded application that uses the Tivoli Storage Manager API. The executable name is the name of the application program that loads the API.

Requirements:

- **dsmtrace** is included with the Tivoli Storage Manager 5.3 backup-archive client. The backup-archive client must be installed to use dynamic tracing.
- As described above, the *DSMTRACELISTEN YES* option must be in effect when the client is started. On Windows, this option is specified in the client options file (usually dsm.opt). On Unix, this option is specified in the system options file (dsm.sys) in the stanza that the client uses.
- Unix users must be logged in as root in order to use **dsmtrace**.
- Windows users must be logged in as a member of the Administrators group.

How it works:

When the client starts, it launches a separate "trace listener" thread. This thread "listens" on a named pipe, waiting to be contacted by the **dsmtrace** utility. In order to make the named pipe name unique, the client process ID (PID) is a part of the pipe name. When **dsmtrace** is used to configure tracing, it contacts the client through the named pipe on which the client is listening and passes to it the desired trace configuration operation. The client then passes the results of the operation back to **dsmtrace** through another similarly-named output pipe. **dsmtrace** displays the results to the console. The client starts the trace listener thread only when client option *DSMTRACELISTEN YES* is in effect. If *DSMTRACELISTEN NO* is in effect, then the listener thread is not started and dynamic tracing is not available to that client. *DSMTRACELISTEN NO* is currently the default value.

Dynamic Tracing Example:

1) Identify the process ID (PID) of the client you want to trace (make sure that *DSMTRACELISTEN YES* is in effect). Use the following command to show all running instances of the client:

dsmtrace query pids

Example output:

```
D:\tsm>dsmtrace query pids

IBM Tivoli Storage Manager
dsmtrace utility
  dsmtrace Version 5, Release 3, Level 0.0
  dsmtrace date/time: 10/24/2004 21:07:36
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.

PROCESS ID   PROCESS OWNER   DESCRIPTION                   EXECUTABLE NAME
      4020   andy            Backup-Archive Client (CLI)   dsmc.exe

D:\tsm>
```

Important note for Linux users: the threading model for some versions of Linux is to run each thread as a separate process. This means that when you query process information, you may see several processes for each instance of the client. The process you need to identify is the dsmc parent process. For example:

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

IBM Tivoli Storage Manager
dsmtrace utility
  dsmtrace Version 5, Release 3, Level 0.0
  dsmtrace date/time: 10/24/04   08:07:37
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.

PROCESS ID   PROCESS OWNER   DESCRIPTION                   EXECUTABLE NAME
     28970   root            Backup-Archive Client (CLI)   dsmc
     28969   root            Backup-Archive Client (CLI)   dsmc
```

```
     28968   root            Backup-Archive Client (CLI)   dsmc
     28967   root            Backup-Archive Client (CLI)   dsmc

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

In such a situation, use the ps command to identify the parent dsmc process:

```
linuxppc:~ # ps -ef | grep dsmc

root     28967  1151  0 Oct22 pts/16   00:00:00 dsmc
root     28968 28967  0 Oct22 pts/16   00:00:00 dsmc
root     28969 28968  0 Oct22 pts/16   00:00:00 dsmc
root     28970 28968  0 Oct22 pts/16   00:00:00 dsmc
root     24092 24076  0 08:15 pts/93   00:00:00 grep dsmc

linuxppc:~ #
```

Notice that the parent for processes 28969 and 28970 is 28968. The parent for 28968 is 28967. The parent for 28967 is 1151, but the 1151 process does not appear in this display output. Process 1151 is the process that launched dsmc. So, the correct parent process ID is 28967.

2) Enable tracing on the client:

dsmtrace enable 4020 -traceflags=service -tracefile=d:\trace.txt

Example output:

```
C:\program files\tivoli\tsm\baclient>dsmtrace enable 4020 -traceflags=service -
tracefile=d:\trace.txt

IBM Tivoli Storage Manager
dsmtrace utility
  dsmtrace Version 5, Release 3, Level 0.0
  dsmtrace date/time: 10/24/2004 21:45:54
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.

ANS2805I Tracing has been enabled.

C:\program files\tivoli\tsm\baclient>

C:\program files\tivoli\tsm\baclient>
```

Important note when tracing an API application: The -pipenameprefix option must be included.

- On Unix, use prefix /tmp/TsmTraceTargetAPI
- On Windows systems use prefix \\.\pipe\TsmTraceTargetAPI

3) After sufficient trace data has been collected, disable the tracing:

dsmtrace disable 4020

Example output:

```
C:\program files\tivoli\tsm\baclient>dsmtrace disable 4020

IBM Tivoli Storage Manager
```

```
dsmtrace utility
  dsmtrace Version 5, Release 3, Level 0.0
  dsmtrace date/time: 10/24/2004 21:47:43
(c) Copyright by IBM Corporation and other(s) 1990, 2004. All Rights Reserved.

ANS2802I Tracing has been disabled.
```

DSMTRACELISTEN Syntax:

DSMTRACEListen No | Yes

No:

The client will not start the trace listener thread, and dynamic tracing will not be available. This is the default.

Yes:

The client will start the trace listener thread, and dynamic tracing will be available.

On Windows, this option is specified in the client options file (usually dsm.opt). On Unix, this option is specified in the system options file (dsm.sys) in the stanza that the client uses. This option cannot be specified from the command line.

dsmtrace Syntax and Options:

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.

-PIPENameprefix=<*pipe name prefix*>

This option must be used when tracing API applications:

- On Unix, use prefix /tmp/TsmTraceTargetAPI
- On Windows, use prefix \\.\pipe\TsmTraceTargetAPI

-TRACEFIle=<*traceFileName*>

This must specify a valid file name to which the trace will be written. If tracing is already running, then this option has no effect.

-TRACEFLags=<*traceFlags*>

Specify one or more trace flags. Typically the trace flag SERVICE is used. Multiple trace flags are separated with a comma. Trace flags can also be turned off by prefixing the flag name with a minus sign. When combining trace flags that you want to turn on with trace flags that you want to turn off, put the flags 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=<*maximumTraceSize*>

This option limits the maximum trace file length to the specified value (by default the trace file will grow indefinitely). When the maximum length is reached, then the trace will wrap back to the beginning of the file. Specify a value in MB between 1 and 4095. If tracing is already running, then this option has no effect.

-TRACESegsize=<*maximumTraceSegmentSize*>

This option is used when you anticipate a large trace file, and you want the trace file to be written in smaller, more easily manageable segments. Each segment will be no larger than the specified size. When this option is used, a segment number will be appended to the trace file name for each segment. Specify a value in MB between 1 and 1000. If tracing is already running, then this option has no effect.

Notes:

- 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).

dsmtrace disable <*pid*> [<*options*>]

Use this command to stop tracing for a process. The trace file will be closed and the trace flags, maximum trace size, maximum trace segment size, and trace file name will all be cleared.

<*pid*>:

This is 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.

-PIPENameprefix=<*pipeNamePrefix*>

This option must be used when tracing API applications:

- On Unix, use prefix /tmp/TsmTraceTargetAPI
- On Windows, use prefix \\.\pipe\TsmTraceTargetAPI

dsmtrace help

This command displays basic syntax for dsmtrace.

dsmtrace query pids [-Filter=<*spec*>]

<*spec*>:

Client process name filter specification. This 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.

Note: When using the FILTER option on Unix, put the * symbol before and after the search text. This is necessary because the executable name often includes the path in front of it; and in some cases, the executable name may 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*>:

This is 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

-PIPENameprefix=<*pipeNamePrefix*>

This option must be used when tracing API applications:

- On Unix, use prefix /tmp/TsmTraceTargetAPI
- On Windows, use prefix \\.\pipe\TsmTraceTargetAPI

<*displayType*>:

This can be one of the following:

-ALl
: This 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
: This displays the names of the trace flags that are turned on. The information shown with the -BASIC display type is also included.

-OFf
: This displays the names of the trace flags that are turned off. The information shown with the -BASIC display type is also included.

-BASic
: This displays the name of the trace file and the maximum trace and trace segment sizes, and indicates whether tracing is enabled or disabled.

Other Examples

DSMTRACE QUERY PIDS

Display all running processes whose names are listed in the table in the Background section.

```
dsmtrace query pids -filter=*
```

Display all running processes.

```
dsmtrace query pids -filter=dsm*
```

Display all running processes whose names begin with "dsm".

dsmtrace query pids -filter=dsm?

Display all running processes whose names are "dsm" plus one other character.

dsmtrace enable 2132 -traceflags=service -tracefile=c:\trace.txt

Turn on SERVICE tracing for process 2132. Trace output will be written to file c:\trace.txt.

dsmtrace enable 2132 -traceflags=-extrc

Turn 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

Turn on FILEOPS tracing for process 4978. The trace will be written to files /tmp/dsmtrace.out.1, /tmp/dsmtrace.out.2, etc., with each file being no larger than 200 MB. Once 1000 MB have been written, tracing will wrap back to /tmp/dsmtrace.out.1.

dsmtrace query trace 4978 -on

Display basic trace information and list trace flags that are turned on for process 4978.

dsmtrace disable 4978

Disable tracing for process 4978.

dsmtrace disable 364 -pipenameprefix=/tmp/TsmTraceTargetAPI

Disable tracing for API application process 364.

Known Problems and Limitations

- If tracing is not currently active for a process, and dsmtrace is used only with the -TRACEFLAGS option, e.g. dsmtrace enable 2346 -traceflags=service, then you will still see the following message:

  ANS2805I Tracing has been enabled.

  In this case, the trace flags have been enabled (turned on), but tracing is not actually active until a trace file has been specified using the -TRACEFILE option.

- Do not use the **dsmtrace enable** command to start tracing the API for Data Protection applications if the Data Protection application is being run in a manner that does not cause it to connect to the Tivoli Storage Manager server. For example, The Data Protection for Lotus Domino command line interface has several such commands:

  - domdsmc help
  - domdsmc set
  - domdsmc query domino
  - domdsmc query pendingdbs
  - domdsmc query preferences

  Using dsmtrace to enable tracing for such commands can result in a hang of the dsmtrace process and (Unix only) a residual named pipe in the /tmp directory.

- On Windows systems, you must be logged in with a local administrative account in order to use **dsmtrace**.
- On Unix systems, you must be logged in as root to use **dsmtrace**.
- On Unix systems, if a client process crashes or is killed, it may leave a named pipe (Unix FIFO) in the /tmp directory. These FIFOs have names beginning with TsmTrace, and they include a process ID (PID) number. If this happens, and then a new client process is started whose PID happens to match that of the old residual FIFO, then the trace listener thread will probably not start. A fix for this is planned for a future release. Any old FIFOs with process numbers that do not match those of running Tivoli Storage Manager processes can be safely deleted. Do NOT delete the FIFO of a running process.
- The threading model for some versions of Linux is to run each thread as a separate process. This means that when you query process information, you may see several processes for each instance of the client. The process you need to identify is the dsmc parent process. Refer to the Dynamic Tracing Example section above for an example.
- When multiple instances of the same program are running, you need to identify the PID of the instance you want to trace. In such a situation, other clues, such as process information from the operating system, may be available to narrow down the desired 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 an invalid option and the client does not start, you might see some named pipe errors in the dsmerror.log file. These error messages may be safely ignored. A fix for this is planned for a future release.

Main Menu | Tracing | Enabling trace for the Tivoli Storage Manager device driver

Enabling trace for the Tivoli Storage Manager Device Driver

Tracing is available for the Tivoli Storage Manager device driver. These instructions apply to the Tivoli Storage Manager device driver on all platforms where the device driver is supported. The Tivoli Storage Manager device driver can be traced from the server console, an administrative client, or from a shell running on the system where the device driver is installed.

For devices that use device drivers other than the Tivoli Storage Manager device driver, the ability to trace and instructions on how to trace those device drivers is provided by the vendor for that device.

Tracing from the server console/admin client:

To trace the driver from the server, you must first issue the TRACE ENABLE and TRACE BEGIN commands (see Tracing for server or storage agent).

The TSM device driver actually consists of three drivers: one for library-autochanger devices, one for tape devices, and one for optical drives. You may choose which one you want to trace. The syntax for the command is:

DDTRACE START [LIBRARYDD | TAPEDD | OPTICALDD] [flags=EE | , FULL | , SYSLOG | BASE]
DDTRACE GET [LIBRARYDD | TAPEDD | OPTICALDD]
DDTRACE END [LIBRARYDD | TAPEDD | OPTICALDD]

The options are:

START
: Turns on tracing. This 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. This does not wipe out the contents of the buffer, so it is allowable to run END before running GET.

LIBRARYDD
: Traces the device driver that controls library-autochangers.

TAPEDD
: Traces the device driver that controls tape drives.

OPTICALDD
: Traces the device driver that controls optical drives.

For these options, it is permitted to specify any one device driver, or the library device driver and one of the other two. These are space delimited. For example:

DDTRACE START TAPEDD - Starts tracing the device driver that controls tape drives.
DDTRACE START OPTICALDD - Starts tracing the device driver that controls optical drives.
DDTRACE START LIBRARYDD - Starts tracing the library-autochanger.
DDTRACE START LIBRARYDD TAPEDD - Traces both the library and the tape drives.
DDTRACE START LIBRARYDD OPTICALDD - Traces both the library and the optical drives.

Whichever of these you use, you should specify the same ones for all commands in the start-get-end series.

The FLAGS parameter is optional and usually not required. The values for the FLAGS parameter are:

EE
: Traces all device driver routine entries and exits.

FULL
: Turns on more debug tracing. It provides more detail but because the memory buffer size is fixed, fewer events are traced. It also does not trace routine entry and exit points.

SYSLOG
: On some platforms, this directs the trace statements to be written to the system log in addition to the memory buffer. This is most useful in debugging kernel crashes or in circumstances where the trace wraps in the memory buffer.

BASE
: This is the default and cannot be specified with any other flags. It is only used to turn off the EE, FULL, and SYSLOG flags without turning off trace.

Tracing from a command shell (only valid for 5.2.2 and above) - All platforms

A stand-alone utility, ddtrace, is provided that exactly mimics the DDTRACE server commands for ease of remembering them. It 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 the stand-alone utility is primarily for cases when the driver needs to be traced during the Tivoli Storage Manager server initialization. It writes the memory buffer to the file "ddtrace.out" in the current directory. If the file exists, it appends to the file and does not overwrite it.

Main Menu | Tracing | Enabling trace for the server or storage agent

Enabling trace classes for the server or storage agent.

1. Trace commands can be issued from the following locations:
   - server console
   - storage agent console
   - administrative client connected to either the server or storage agent
   - server options file (dsmserv.opt)
   - storage agent options file (dsmsta.opt)

   Trace commands apply to the server or storage agent to which the command was submitted. Trace commands in the options files are used to trace during startup and initialization of the applications or to provide a default set of trace classes. Starting with Tivoli Storage Manager Version 5.3, there are two trace classes that are always enabled by default, whether they appear on the options file or not. These are ADDMSG and SYSTIME.
2. Determine the trace classes to enable. See Trace Classes for server or storage agent for possible trace classes.

3. To have trace messages issued for a given trace class, that trace class needs to be enabled either prior to beginning the trace or after the tracing has begun. To enable one or more trace classes, issue:

   TRACE ENABLE <*traceClassName*>

   Note that <*traceClassName*> may be a space-delimited list of trace classes. For example, this 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 additional times. For example, if you wanted to add the PVR trace class in addition to those that were already enabled, issue: TRACE ENABLE PVR
4. To stop having trace messages issued for a given trace class, that trace class needs to be disabled either prior to beginning the trace or after the tracing has begun. To disable one or more trace classes, issue:

   TRACE DISABLE <*traceClassName*>

   Note that <trace class name> may be a space delimited list of trace classes. For example, this could be entered as TRACE DISABLE TM SESSION. Additional trace classes can also be disabled by issuing TRACE DISABLE. For example, if you wanted to remove the PVR trace class in addition to those that were already disabled, issue: TRACE DISABLE PVR. By issuing TRACE DISABLE without specifying any trace classes, all currently enabled trace classes will be disabled.
5. Tracing can occur to the console or to a file.
   - For tracing to the console, issue:
     TRACE BEGIN
   - For tracing to a file with no size limitation, issue:
     TRACE BEGIN <*fileName*>
   - For tracing to a file with a size limitation, issue:
     TRACE BEGIN <*fileName* > MAXSIZE=<*maximum size in megabytes*>

   NOTE: the *fileName* can be a fully qualified path such as /opt/tmp or c:\temp. If a full path is not given, the trace file will be located in the same directory as the running executable.
6. Perform the operation that is causing the problem.
7. To stop trace messages from being issued, issue:

   TRACE END

   If tracing was being done to a file, ending the trace will write any remaining trace messages to the file and then close the file.

As a general recommendation, it is best to trace to a file. Typically, the tracing for the server or storage agent will generate a large amount of output.

As noted in item 1, it is possible to enable tracing and begin it using the server or storage agent options file. The commands and syntax discussed above 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

If the tracing was no longer needed after startup, the tracing could be ended by issuing TRACE END from the console or from an administrative client.

Main Menu | Tracing | Trace classes for the Administration Center

Trace classes for the Administration Center.

The Administration Center provides individual and aggregate trace classes. Aggregate classes are a shortcut for enabling many related trace classes by simply specifying the aggregate trace class name.

The trace classes listed are those that are most typically requested or used for diagnosing problems. This list does not include all possible trace classes that are available.

To find information about a trace class, click the trace class in the bulleted list below or scroll through the list of trace classes below the bulleted list. The trace class name should be used when turning tracing on using the Administration Center Support Utility. The information about a trace class is in the form:

TRACE_CLASS_NAME

Description of trace class

Recommendation for when to use

Other information about this trace class

List of trace classes

- ALL
- SNAPIN
- CONTROL
- VIEWS
- ACTIONS
- UTIL
- PS
- API
- HEALTH
- XMLVAL
- TAGLIB

ALL

Displays all processing information for the Administration Center. An aggregate that includes all of the trace classes listed in the bulleted list above.

Useful in most cases. Generally recommended when the nature of the problem is unknown. If the ALL trace flag is used it will not be necessary to specify any other trace flags because it already includes all of the trace classes.

Return to list of trace classes

SNAPIN

Displays processing information for the snapin portlet classes. An aggregate that includes most of the Administration Center trace classes listed in the bulleted list above.

Useful in many cases. Generally recommended for diagnosing problems with the Administration Center when the Health Monitor and XML files are working fine.

Aggregate of all trace classes except HEALTH, XMLVAL and TAGLIB.

Return to list of trace classes

CONTROL

Displays processing information for the connection manager, the portlet factory and other key control components.

Useful for problems related to control objects. This is useful for debugging problems that occur when performing actions that use internal control objects such as the connection manager.

Return to list of trace classes

VIEWS

Displays processing information for the view classes.

Useful for problems related to view objects. This is useful for debugging problems that occur while displaying a page.

Return to list of trace classes

ACTIONS

Displays processing information for the action classes.

Useful for problems related to action objects. This is useful for debugging problems that occur as an action before displaying a page.

Return to list of trace classes

UTIL

Displays processing information for the utility classes.

Useful for problems related to utility objects. This is useful for debugging problems that occur in various internal utility routines.

Return to list of trace classes

PS

Displays general information in the processing of Presentation Services object extensions. This traces all listener, validator and JavaServer Page (JSP) bean classes.

Useful for problems related to graphical user interface objects such as tables and buttons.

Return to list of trace classes

API

Displays information in the processing of the API interface. The API interface is used for communications between the Administration Center and the Tivoli Storage Manager server.

Useful for problems related to command execution on the Tivoli Storage Manager server.

Return to list of trace classes

HEALTH

Displays all information about the health monitor.

Useful for problems related to the Health Monitor.

Return to list of trace classes

XMLVAL

Displays all information in the processing of XML files.

Useful for problems related to the file controlling which servers the Administration Center knows about (tsmservers.xml) and the command file cache.

Return to list of trace classes

TAGLIB

Displays all information related to the JSP tag library extensions.

Useful for problems related to JSP pages.

Return to list of trace classes

Main Menu | Tracing | Trace classes for the client

Trace classes for the client.

The client provides individual and aggregate trace classes. Aggregate classes are a shortcut for enabling many related trace classes by simply specifying the aggregate trace class name. In the table below, there may 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 listed are those that are most typically requested or used for diagnosing problems. This list does not discuss all possible trace classes that are available.

To find information about a trace class, click the trace class in the table below or scroll through the list of trace classes below the table. The trace class name should be used with the TRACEFLAG options in DSM.OPT. The information about a trace class is in the form:

TRACE_CLASS_NAME

Description of trace class

Recommendation for when to use

Other information about this trace class

List of trace classes

- SERVICE
- SCHEDULER
- ALL_BACK
- ALL_FILE
- ALL_IMAGE
- ALL_JBB
- ALL_NAS
- ALL_SESS
- ALL_SNAPSHOT
- ALL_WAS
- AUDIT
- CLIENTTYPE
- COMPRESS
- DELTA
- DIROPS
- DOMAIN
- ENCRYPT
- ERROR
- FILEOPS
- FIOATTRIBS
- INCR
- INCLEXCL
- MEMORY
- NETWARE
- OPTIONS
- PASSWORD
- PID
- POLICY
- SESSION
- SESSVERB
- SMS_DEBUG
- SMS_DETAIL
- STAS
- THREAD_STATUS
- TID
- TXN
- VERBDETAIL
- VERBINFO
- WIN2K

SERVICE

Displays general processing information for the client. An aggregate that includes most of the client trace classes listed in this table.

Useful in many cases. Generally recommended when the nature of the problem is unknown. If the SERVICE trace flag is used it will generally not be necessary to specify any other trace flags because it already includes most of the basic trace classes.

Aggregate of all trace classes except MEMORY and *DETAIL classes. The SERVICE trace flag can generate a substantial amount of information. Consider using the TRACEMAX option in conjunction with the SERVICE trace flag.

Return to list of trace classes

SCHEDULER

Displays general processing information for the scheduler. An aggregate that includes most of the client trace classes listed in this table.

Useful in many cases. Generally recommended for diagnosing scheduler problems when the nature of the problem is unknown. If the SCHEDULER trace flag is used, it will generally not be necessary to specify any other trace flags because it already includes most of the basic trace classes.

Aggregate of all trace classes except MEMORY, THREAD_STATUS and *DETAIL classes.

Return to list of trace classes

ALL_BACK

Displays general backup processing information for the client.

Useful for problems related to selective or incremental backups.

Aggregate of TXN, INCR, POLICY, and PFM trace classes and implicitly included in the SERVICE trace class.

Return to list of trace classes

ALL_FILE

Displays general backup processing information for the client.

Useful for problems related to reading and writing of data and obtaining file attribute information.

Aggregate of DIROPS, FILEOPS, and FIOATTRIBS trace classes and implicitly included in the SERVICE trace class.

Return to list of trace classes

ALL_IMAGE

Displays image processing information for the client.

Useful for problems related to all aspects of volume image backup-and-restore operations.

Aggregate of several image-related trace classes and implicitly included in the SERVICE trace class.

Return to list of trace classes

ALL_JBB

Displays journal-based backup processing information for the client.

Useful for problems related to all aspects of journal-based backups.

Aggregate of several journal-based-backup related trace classes and implicitly included in the SERVICE trace class.

Return to list of trace classes

ALL_NAS

Displays NDMP processing information for the client.

Useful for problems related to all aspects of NDMP backup-and-restore operations.

Aggregate of several NDMP-related trace classes and implicitly included in the SERVICE trace class.

Return to list of trace classes

ALL_SESS

Displays all session and verb information sent between the client and the server.

Useful for problems related to the client and server session such as communication timeouts, protocol violations, and instances where the client appears to be hung waiting for the server or visa versa.

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 for VERBDETAIL

Return to list of trace classes

ALL_SNAPSHOT

Displays information relating to volume snapshot operations.

Useful to determine problems related to volume snapshots that are used in online image backup and open file-support operations.

Aggregate of several volume snapshot-related trace classes and implicitly included in the SERVICE trace class.

Return to list of trace classes

ALL_WAS

Displays Web Application Server processing information for the client.

Useful for problems related to all aspects of Web Application Server backup-and-restore operations.

Aggregate of several Web Application Server related trace classes and implicitly included in the SERVICE trace class.

Return to list of trace classes

AUDIT

Displays auditing information for backup-and-restore processing.

Useful to keep record of files processed, committed and restored in a file.

Part of the SERVICE trace aggregate.

Return to list of trace classes

CLIENTTYPE

Displays client type on each trace output line.

Useful for tracing situations when more then one client component is involved, such as the client acceptor daemon (CAD) and the file system agent.

None.

Return to list of trace classes

COMPRESS

Displays compression information.

Useful to determine how much data is compressed on a per-file basis.

Part of the SERVICE trace aggregate.

Return to list of trace classes

DELTA

Displays adaptive subfile backup processing information.

Useful to determine errors in adaptive subfile backup-and-restore operations.

Part of the SERVICE trace aggregate.

Return to list of trace classes

DIROPS

Displays directory read and write operations.

Useful when problem occurs in a directory read or write.

Part of the SERVICE and ALL_FILE trace aggregates.

Return to list of trace classes

DOMAIN

Displays incremental domain processing information.

Useful for determining how DOMAIN statements are resolved during backup processing such as problems in resolving the ALL-LOCAL domain.

Part of the SERVICE trace aggregates.

Return to list of trace classes

ENCRYPT

Displays data encryption information.

Useful to determine if a file is included for encryption processing.

Part of the SERVICE trace aggregate.

Return to list of trace classes

ERROR

Displays operating system specific error information.

Useful to determine error codes generated by the operating system.

Part of the SERVICE trace aggregate.

Return to list of trace classes

FILEOPS

Displays file read and write operations.

Useful when problem occurs in a file open, read, write, or close operation.

Part of the SERVICE and ALL_FILE trace aggregates.

Return to list of trace classes

FIOATTRIBS

Displays comparisons of file attributes between the local client version and the active version on the Tivoli Storage Manager server.

Useful in determining why a file was backed up during an incremental backup.

Part of the SERVICE, ALL_BACK, and ALL_FILE trace aggregates.

Return to list of trace classes

INCR

Displays incremental list processing comparisons between the client and server.

Useful for determining if files are candidates for incremental backup, especially in conjunction with the FIOATTRIBS trace class.

Part of the SERVICE and ALL_BACK trace aggregates.

Return to list of trace classes

INCLEXCL

Displays include/exclude status for the object being processed. This flag is also used for the Preview function.

Useful for determining which object (usually file or directory) is included or excluded during backup/archive/preview.

Return to list of trace classes

MEMORY

Displays memory allocation and free requests.

Useful for determining memory leaks, memory spikes, and other memory-related problems.

This trace class writes a large amount of information into the trace file and is not included in any aggregate classes.

Return to list of trace classes

NETWARE

Displays detailed information relating to the Novell SMS backup API.

Useful for isolating NetWare client problems between the backup application and operating system.

Aggregate of the SMS_DEBUG, SMS_DETAIL TXN and TID trace classes.

Return to list of trace classes

OPTIONS

Displays current processing options.

Useful for determining which options are in effect for the current session and problems in accepting processing options from server client-options sets.

Part of the SERVICE trace aggregate.

Return to list of trace classes

PASSWORD

Displays password file-access information (does not show passwords).

Useful for determining problems with reading Tivoli Storage Manager server passwords from local storage, for example, PASSWORDACCESS=GENERATE errors.

Part of the SERVICE trace aggregate.

Return to list of trace classes

PID

Displays process ID on each trace statement.

Useful when diagnosing problems that may involve multiple processes.

Part of the SERVICE trace aggregate.

Return to list of trace classes

POLICY

Displays policy information available to the Backup/Archive client.

Useful to see which policies are available during a backup or archive operation.

Part of the SERVICE and ALL_BACK trace aggregates.

Return to list of trace classes

SESSION

Displays minimal session information between the client and the server.

Useful for giving session context to general processing errors or in conjunction with one of the VERB* trace classes to determine session problems such as session timeouts and protocol violations.

Part of the SERVICE and ALL_SESS trace aggregates.

Return to list of trace classes

SESSVERB

Displays additional session information between the client and the server.

Useful for giving session context to general processing errors or in conjunction with one of the VERB* trace classes to determine session problems such as session timeouts and protocol violations.

Part of the SERVICE and ALL_SESS trace aggregates.

Return to list of trace classes

SMS_DETAIL

Displays detailed information relating to the Novell SMS backup API.

Useful for isolating NetWare client problems between the backup application and operating system.

Part of the SERVICE and NETWARE trace aggregates.

Return to list of trace classes

SMS_DEBUG

Displays detailed information relating to the Novell SMS backup API.

Useful for isolating NetWare client problems between the backup application and operating system.

Part of the SERVICE and NETWARE trace aggregates.

Return to list of trace classes

STATS

Displays final processing statistics in the trace file.

Useful for collecting final processing statistics into a file.

Part of the SERVICE trace aggregate.

Return to list of trace classes

THREAD_STATUS

Displays thread status.

Useful when diagnosing problems related to threading.

Part of the SERVICE trace aggregate.

Return to list of trace classes

TID

Displays thread ID on each trace statement.

Useful when diagnosing problems related to threading and to distinguish different processing threads in the trace file.

Part of the SERVICE and NETWARE trace aggregates.

Return to list of trace classes

TXN

Displays transaction processing information.

Useful when diagnosing problems related to transaction processing problems on the server such as transaction aborts and retries.

Part of the SERVICE, ALL_BACK and NETWARE trace aggregates.

Return to list of trace classes

VERBDETAIL

Displays detailed verb information pertinent to client-server sessions.

Useful to determine contents of verbs sent between the client and server.

Part of the ALL_SESS trace aggregates.

Return to list of trace classes

VERBINFO

Displays verb information pertinent to client-server sessions.

Useful in conjunction with the SESSION traceflag to give session context to general processing errors or to determine session problems like session timeouts and protocol violations.

Part of the SERVICE and ALL_SESS trace aggregates.

Return to list of trace classes

WIN2K

Displays Windows(R) system object or system state processing.

Useful for determining errors with backup or restore of the system state information.

Part of the SERVICE trace aggregates. Only valid on the Windows Backup/Archive client.

Return to list of trace classes

Main Menu | Tracing | Trace classes for the server or storage agent

Trace classes for server or storage agent.

The server and storage agent provide aggregate trace classes. These are a shortcut for enabling many related trace classes by simply specifying the aggregate trace class name for the TRACE ENABLE command. In the table below, there may be references to trace classes that are enabled as part of an aggregate trace class, but they are not explicitly discussed on their own. In this case, these trace classes are not typically used for problem diagnosis.

The trace classes listed are those that are most typically requested or used for diagnosing problems. This list does not discuss all possible trace classes that are available.

To find information about a trace class, click the trace class in the table below or scroll through the list of trace classes below the table. The trace class name should be used with the TRACE ENABLE and TRACE DISABLE commands. The information about a trace class is in the form:

TRACE_CLASS_NAME

Description of trace class

Recommendation for when to use

Other information about this trace class.

List of trace classes

- ACCNT
- ADDMSG
- ADMCMD
- ADMNODE
- ADMPROC
- ADMRES
- AF
- AFCOPY
- AFCREATE
- AFLOCK
- AFMOVE
- AFTXN
- ARCH
- ARCHD
- ARCHUTIL
- AS
- ASALLOC
- ASDEALLOC
- ASMOUNT
- ASRTRV
- ASTXN
- ASVOL
- ASSD
- AUDIT
- BACK
- BACKD
- BF
- BFAGGR
- BFCREATE
- BFLOCK
- BFREMOTE
- BFRTRV
- BFSALVAGE
- BFTRG
- BITVECTOR
- BKSET
- BLKDISK
- BRNODE
- CA
- CAACOMAX
- CAACOMIN
- CAENG
- CAMEM
- CASCSI
- CRCVAL
- CRYPTO
- DB
- DBALLOC
- DBBACKUP
- DBBUF
- DBBUFPF
- DBBUFWR
- DBCKPT
- DBDISPCOLS
- DBDUMP
- DBLOG
- DBNETDB
- DBNETENCODE
- DBSERVER
- DBTXN
- DEADLOCK
- DELTA
- DF
- DFCOPY
- DFCREATE
- DFLOCK
- DFMOVE
- DFRTRV
- DFTXN
- DRM
- DRMAPPEND
- DRMLICCOMP
- DRMMACH
- DRMMACHCHAR
- DRMPR
- DRMPRDBBACK
- DRMPRDEVCLASS
- DRMPRLOGDBVOL
- DRMPRMDRM
- DRMPRREQPOOL
- DRMPRST
- DRMPRSTGPOOL
- DRMRECOVINST
- DRMRECOVMEDIA
- DRMSETGA
- DS
- DSALLOC
- DSDEALLOC
- IC
- ICBACK
- ICBITS
- ICDISK
- ICHDRPG
- ICLOCK
- ICLOG
- ICREST
- ICSTREAM
- ICTRIGGER
- ICVOLHST
- ID
- IMDEL
- IMEXP
- IMFS
- IMLOCK
- IMSORT
- INSTR
- INTERACTION
- LANFREE
- LANG
- LM
- LNFDATA
- LNFENTRY
- LNFMEM
- LNFVERB
- LOCK
- LOG
- LOGCOMP
- LOGPARTW
- LOGREC
- LOGRESIZE
- LOGSEG
- LVM
- MMS
- MPCOUNT
- NA
- NACONFIG
- NADISK
- NAHBAAPI
- NALOCK
- NAMOVER
- NAPATH
- NQR
- PERFSTAT
- PM
- PMCACHE
- PMCG
- PMLOCK
- PMMC
- PROXYNODE
- PVR
- REDIRECT
- RETPROT
- SCHED
- SCHEDWAIT
- SESSION
- SESSIOND
- SESSREMOTE
- SETOBJRANGE
- SFACO
- SFPERF
- SPI
- SPID
- SS
- SSBUFQ
- SSCLASS
- SSFRAME
- SSLOCK
- SSPOOL
- SSSESS
- SSSWINIT
- SSSWTHR
- SSTRANS
- SSVOL
- SW
- STATXN
- TAPE
- TBPREFETCH
- TCP
- TCPDATA
- TCPINFO
- TEC
- TOC
- TOCBUILD
- TOCLOAD
- TOCREAD

- CC
- CCLOCK
- CENTERA
- COLLOCATE
- COLLOCATEDETAIL
- CRC
- CRCDATA
- CRCPROTO
- DSRTRV
- DSVOL
- EXTERNAL
- FORMAT
- GROUP
- HSMPFR
- HT
- NQRDETAIL
- NQRSEARCH
- OBJSET
- ORM
- ORMCMD
- ORMDELETE
- ORMUPDATE
- TOCUTIL
- TXN
- UNDERSIZE
- UNICODE
- UNICODEDETAIL
- XI
- XIDETAIL

ADDMSG

Issues console messages (ANR, ANE, and so on) to the trace file.

Valuable for correlating server messages to trace messages and preserving the timing for when each was issued.

Available in V5.2 or V5.1.7.0 or later.

Return to list of trace classes.

ADMCMD

Traces related to command processing.

Use to debug the command interpreter, including the PARALLEL and SERIAL command handling.

None.

Return to list of trace classes.

AF

Displays information about user data stored on sequential media devices.

Used to diagnose problems reading or writing user files to sequential media volumes.

This is an aggregate trace class that enables: AFCREATE, AFMOVE, AFLOCK, AFTXN, and AFCOPY. Typically recommended to issue TRACE DISABLE AFLOCK unless the locking information is explicitly requested or needed.

Return to list of trace classes.

AFCREATE

Displays information about storing user data on sequential media volumes.

Used to diagnose writing user data on sequential media volumes.

None.

Return to list of trace classes.

AFMOVE

Displays operations that move user data using sequential media volumes. Move operations are performed by MIGRATION, RECLAMATION, MOVE DATA, and MOVE NODEDATA server processes.

Used to diagnose problems with the data movement server processes.

None.

Return to list of trace classes.

AS

Displays information volume selection and assignment, coordination of drives (mount points), and management of data placement on volumes.

Used to diagnose many different problems relating to volumes, mount points, or data read and write operations.

This is an aggregate trace class that enables: ASALLOC, ASRTRV, ASDEALLOC, ASMOUNT, ASVOL, ASTXN, and ASSD. Typically recommended to issue TRACE DISABLE ASTXN unless the locking information is explicitly requested or needed.

Return to list of trace classes.

ASALLOC

Displays information about reserving and allocating space on sequential media volumes for the purpose of storing data. This may be storing data on behalf of a client session or for server data movement operations such as MIGRATION, RECLAMATION, MOVE DATA, or MOVE NODEDATA.

Used to diagnose problems where the server or storage agent report no space available but there should be space available in the storage hierarchy.

None.

Return to list of trace classes.

ASRTRV

Displays information about reading data from sequential media volumes.

Used to diagnose problems reading data such as RESTORE or RETRIEVE client by the client or MIGRATION, RECLAMATION, STORAGE POOL BACKUP, AUDIT VOLUME, GENERATE BACKUPSET, EXPORT, MOVE DATA, or MOVE NODEDATA by the server.

None.

Return to list of trace classes.

ASDEALLOC

Displays information about releasing and deallocating space on sequential media volumes for the purpose of storing data. Typical de-allocation operations on the server are EXPIRATION, MIGRATION, RECLAMATION, MOVE DATA, MOVE NODEDATA, AUDIT VOLUME, DELETE VOLUME, and DELETE FILESPACE.

Used to diagnose during the deletion of data.

None.

Return to list of trace classes.

ASMOUNT

Displays information about drive (mount point) selection and assignment for sequential media devices.

Used to diagnose situations where sessions or processes are waiting for mount points or cases where an operation fails because no mount point is available. Also useful in cases where a mount point is preempted.

None.

Return to list of trace classes.

ASVOL

Displays information about volume selection and assignment for sequential media volumes.

Used to diagnose situations where sessions or processes are waiting for volumes or cases where an operation fails because no volume is available. Also useful in cases where volume access is preempted.

Return to list of trace classes.

ASTXN

Displays information about transactions used to make database updates relating to information for sequential media volumes, storage pools, device classes, and other attributes.

Used to diagnose hangs, database operations failures reported for sequential media operations, or general data storage problems.

Return to list of trace classes.

ASSD

Displays information about sequential stream data operations. These are operations that use sequential media device classes, volumes, or mount points but do not store data in the storage hierarchy. Server processes that perform sequential stream data operations are BACKUP DB, EXPORT/IMPORT, and GENERATE BACKUPSET.

Used to diagnose server processes that perform sequential stream data operations.

None.

Return to list of trace classes.

BF

Information about user data (files) stored in the storage hierarchy.

Used to diagnose general data read or write problems for client operations and server processes.

This is an aggregate trace class that enables: BFCREATE, BFRTRV, BFSALVAGE, BFLOCK, BFAGGR, BFREMOTE, and BFTRG.

Return to list of trace classes.

BFAGGR

Displays information about server aggregation of user data. The server will aggregate many smaller user files into a larger file in the storage hierarchy to optimize performance for data movement operations such as MIGRATION, MOVE DATA, and MOVE NODEDATA.

Used to diagnose general data read or write problems for client operations and server processes or both.

Return to list of trace classes.

BFCREATE

Displays information about client operations that store data in the storage hierarchy. Typically, these are BACKUP, ARCHIVE, or SPACE MANAGE operations by the client.

Used to diagnose failures or other problems while a client is storing data.

None.

Return to list of trace classes.

BFRTRV

Displays information about client operations that read data from the storage hierarchy.

Used to diagnose failures or other problems while a client is reading data.

None.

Return to list of trace classes.

BKSET/OBJSET

Trace class for backupset functions.

Use this trace class to debug problems in the GENERATE BACKUPSET command, or during a client restore from a backupset.

The BKSET and OBJSET trace classes are synonymous.

Return to list of trace classes.

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.

None.

Return to list of trace classes.

BRNODE

Trace class for the BACKUP and RESTORE NODE commands, used during NDMP operations.

Use this trace class to debug problems in the BACKUP and RESTORE NODE commands.

Only available in 4.2.1.0 and up.

Return to list of trace classes.

COLLOCATE

Displays information about collocation processing on storage pools.

Used to diagnose problems with collocation processing.

COLLOCATEDETAIL trace class can also be used to get more detailed information about the collocation processing such as files being processed for a collocation group. This can cause a large amount of output trace statements.

Return to list of trace classes.

CRC

Displays information about generating and managing CRCs on the server or storage agent.

Used to diagnose data corruption issues where CRC processing did not report data corruption.

This is an aggregate trace class that enables: CRCDATA, CRCPROTO, and CRCVAL.

Return to list of trace classes.

CRCDATA

Displays information about generating and managing CRCs for data stored in storage pools with CRCDATA=YES set.

Used to diagnose data corruption issues where CRC processing did not report data corruption.

None.

Return to list of trace classes.

CRCPROTO

Displays information about generating and managing CRCs for data exchanged between the client and either the server or storage agent where this node is configured with VALIDATEPROTOCOL=ALL or VALIDATEPROTOCOL=DATAOnly on the server.

Used to diagnose data corruption issues where CRC processing did not report data corruption.

None.

Return to list of trace classes.

CRCVAL

Displays information about generating and comparing CRC values.

Informational for displaying CRC values during processing.

None.

Return to list of trace classes.

CRYPTO

Displays information about AES encryption operations and some general encryption settings.

Used to isolate and identify encryption related problems.

None.

Return to list of trace classes.

DBNETDB

Displays information about LAN-free operations and the negotiation and management of information between the server and storage agent.

Used to diagnose problems with LAN-free data movement.

Typically applies to server and storage agent prior to Version 5.2. After Version 5.2, this still displays information but has been superceded by the LANFREE trace class. Typically best to enable this trace class on both the server and storage agent concurrently.

Return to list of trace classes.

DELTA/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, including client backup and restore, expiration, deletion (DELETE FILESPACE, DELETE VOLUME), export/import, backupset generation and restore, no-query restore, db audit, etc.

The DELTA and GROUP trace classes are synonymous.

Return to list of trace classes.

DF

Displays information about user data stored on DISK volumes.

Used to diagnose problems reading or writing user files to DISK volumes.

This is an aggregate trace class that enables: DFCREATE, DFRTRV, DFMOVE, DFLOCK, DFTXN, and DFCOPY. Typically recommended to issue TRACE DISABLE DFLOCK unless the locking information is explicitly requested or needed.

Return to list of trace classes.

DFCREATE

Displays information about storing user data on DISK volumes.

Used to diagnose writing user data on DISK volumes.

None.

Return to list of trace classes.

DFRTRV

Displays information about reading user data from DISK volumes.

Used to diagnose reading user data from DISK volumes.

None.

Return to list of trace classes.

DFMOVE

Displays operations that move user data using DISK volumes. Move operations are performed by MIGRATION, MOVE DATA, and MOVE NODEDATA server processes.

Used to diagnose problems with the data movement server processes.

None.

Return to list of trace classes.

DS

Displays information volume selection, space reservation, assignment, and management of data placement on DISK volumes.

Used to diagnose many different problems relating to DISK volume data read and write operations.

This is an aggregate trace class that enables: DSALLOC DSRTRV DSDEALLOC DSVOL. Typically recommended to issue TRACE DISABLE DSTXN unless the locking information is explicitly requested or needed.

Return to list of trace classes.

DSALLOC

Displays information about reserving and allocating space on DISK volumes for the purpose of storing data. This may be storing data on behalf of a client session or for server data movement operations such as MIGRATION, MOVE DATA, or MOVE NODEDATA.

Used to diagnose problems where the server or storage agent report no space available but there should be space available in the storage hierarchy.

None.

Return to list of trace classes.

DSDEALLOC

Displays information about releasing and deallocating space on DISK volumes. Typical deallocation operations on the server are

EXPIRATION, MIGRATION, MOVE DATA, MOVE NODEDATA, AUDIT VOLUME, DELETE VOLUME, and DELETE FILESPACE.

Used to diagnose during the deletion of data.

None.

Return to list of trace classes.

DSRTRV

Displays information about reading data from DISK volumes.

Used to diagnose problems reading data such as RESTORE or RETRIEVE client by the client or MIGRATION, STORAGE POOL BACKUP, AUDIT VOLUME, GENERATE BACKUPSET, EXPORT, MOVE DATA, or MOVE NODEDATA by the server.

None.

Return to list of trace classes.

DSVOL

Displays information about volume selection and assignment for DISK volumes.

Used to diagnose situations where sessions or processes are waiting for volumes or cases where an operation fails because no volume is available.

Return to list of trace classes.

IMFS

Trace class for filespace functions.

Use this trace class to debug problems related to inventory filespaces (e.g. during DELETE FILESPACE).

None.

Return to list of trace classes.

ICVOLHST

Trace class for volume history functions.

Use this trace class to debug problems with creating volume history entries (e.g. during EXPORT, BACKUP DB, or GENERATE BACKUPSET) or deleting volume history entries (e.g. during DELETE VOLHISTORY).

None.

Return to list of trace classes.

LANFREE

Displays general information about LAN-free operations on either the server or storage agent. Also displays error information for LAN-free-related operations.

Any LAN-free failure.

Available only in Version 5.2 or later. Best if traced on both the server and storage agent concurrently. This is an aggregate trace class that enables: LNFVERB, LNFMEM, LNFENTRY, and LNFDATA.

Return to list of trace classes.

MMS

Displays information about tape libraries and the server or storage agent use of these.

Used to diagnose problems with tape libraries, library volume inventories, or other general library issues.

Not Available for z/OS(R). This is an aggregate trace class that enables: MMSBASE, MMSTXN, MMSLIB, MMSDRIVE, MMSOP, MMSMAN, MMSSCSI, MMSFLAG, MMSACSLS, and MMSSHARE. It is usually recommended to include NA and PVR trace classes when tracing MMS.

Return to list of trace classes.

NA

Displays information about path information for the server or storage agent. This relates to the DEFINE PATH, UPDATE PATH, DELETE PATH, and QUERY PATH commands.

Used to diagnose problems with paths to devices.

This is an aggregate trace class that enables: NALOCK, NAPATH, NAMOVER, NADISK, and NACONFIG. Usually best to include MMS and PVR trace classes when tracing NA.

Return to list of trace classes.

PROXYNODE

Displays information about proxynode sessions and the commands related to proxynode associations (GRANT, REVOKE, QUERY PROXYNODE).

Used to diagnose problems with proxynode sessions and related commands.

Usually best to include SESSION trace when analyzing proxynode session problems.

Return to list of trace classes.

PVR

Displays information about sequential media devices and the server or storage agent use of these devices.

Used to diagnose problems with tape drives, failures reading or writing tape volumes, or other tape-volume-related issues.

This is an aggregate trace class that enables: PVRVOL, PVRCLASS, PVRMP, PVRREMOTE. Additionally, this aggregate enables the following classes for open platforms only: PVR8MMBASE, PVR4MMBASE, PVRQICBASE, PVRDLTBASE, PVRECARTBASE, PVRDTFBASE, PVRTAPIBASE, PVRGTS, PVRFILE, PVRTAPE, PVRODSK, PVRNTP, PVRSHARE, and PVRNAS. It is often useful to include MMS and NA trace classes on open platforms and TAPE on z/OS when tracing PVR.

Return to list of trace classes.

RETPROT

Trace class for the archive retention protection functions.

Use this trace class to debug problems using the RETINIT and RETMIN parameters in the archive copy group, or problems using the VB_SignalObject verb (only supported via the client API) to signal an object's event or to hold or release an object, or problems during expiration or deletion of retention protected objects.

Only available in 5.2.2.0 and up.

Return to list of trace classes.

SCHED

Trace class for the central scheduler functions.

Use this trace class to debug problems related to schedule commands like DEFINE/UPDATE/QUERY SCHEDULE or DEFINE ASSOCIATION. Also use this trace class to debug problems related to the central scheduler background processes such as the schedule manager and schedule prompter.

Applies to classic and enhanced schedules equally.

Return to list of trace classes.

SESSION

Displays information about sessions connected to the server, including all verbs sent and received by the server.

Useful in many cases. Generally recommended for protocol violations, transaction processing errors, or in cases where the client is hung and not responding.

None

Return to list of trace classes.

TAPE

Displays information about sequential media devices and the server or storage agent use of these devices.

Used to diagnose problems with tape drives, failures reading or writing tape volumes, or other tape-volume-related issues.

This is only available on z/OS. This is an aggregate trace class that enables: CART, CARTDATA, CARTINFO, REEL, REELDATA, FLAT, and FLATDATA. It is often useful to include the PVR trace class when tracing TAPE.

Return to list of trace classes.

TCP

Collects information regarding TCP/IP used between the client and either server or storage agent.

To debug session connection errors or data corruption issues that may be caused by the network.

This is an aggregate trace class. It enables TCPINFO and TCPERROR.

Return to list of trace classes.

TCPDATA

Detailed TCP/IP trace used to display byte-level information about data that is sent or received.

To debug session data corruption issues that may be caused by the network.

None.

Return to list of trace classes.

TCPINFO

General TCP/IP trace used to display setup and characteristics of TCP/IP on the server or storage agent.

To debug session data corruption issues that may be caused by the network.

None.

Return to list of trace classes.

TEC

Provides information regarding events sent to a TEC server. Corresponds to the 'tivoli' event receiver.

To debug connection issues encountered with TEC event logging.

None.

Return to list of trace classes.

TOC

General trace class for the Table Of Contents (TOC) component, used during file-level NDMP operations.

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.

Only available in 5.2.0.0 and up. This is an aggregate trace class that enables: TOCBUILD, TOCLOAD, TOCREAD, and TOCUTIL.

Return to list of trace classes.

TOCBUILD

Table Of Contents (TOC) build functions.

Use this trace class to debug problems during an NDMP backup with the TOC=YES parameter.

Only available in 5.2.0.0 and up.

Return to list of trace classes.

TOCLOAD

Table Of Contents (TOC) load functions.

Use this trace class to debug problems while displaying files and directories on the client GUI.

Only available in 5.2.0.0 and up.

Return to list of trace classes.

TOCREAD

Table Of Contents (TOC) read functions.

Use this trace class to debug problems during a QUERY TOC command or while trying to load a TOC for displaying files and directories on the client GUI.

Only available in 5.2.0.0 and up.

Return to list of trace classes.

TOCUTIL

Table Of Contents (TOC) utility functions.

Use this trace class to debug problems related to TOC component initialization or TOC retention.

Only available in 5.2.0.0 and up.

Return to list of trace classes.

UNICODE

Displays information about code page conversions and Unicode filespace operations

Use this trace class to debug problems related to code page conversion problems or unicode filespace problems.

Return to list of trace classes.

XI

Displays general information for the IMPORT and EXPORT commands.

Use this trace class to debug problems related to IMPORT and EXPORT commands.

Return to list of trace classes.

Main Menu | Installing and running the tsmdiag utility

Installing and running the tsmdiag utility.

The purpose of the Tivoli Storage Manager Diagnostic utility (tsmdiag) is to ease and speed up the process for gathering data that is considered valuable to assist in diagnosing a problem caused by a server component. There is a set of data that tsmdiag collects by default. This data is specified in a configuration file called configfile.txt. If at any point more data is identified as valuable to diagnose any particular problem from a Tivoli Storage Manager component, users will be instructed to modify the configuration file used by tsmdiag so that such data is collected by the tsmdiag utility.

It is strongly recommended that users execute tsmdiag on the system where the Tivoli Storage Manager component with a problem is

installed. This should be done before submitting a problem to Tivoli Storage Manager Customer Support. The data collected by tsmdiag should be submitted along with the problem report.

Supported Operating environments and Tivoli Storage Manager Components.

Supported OS Environments and Components	
Operating System Environment	**Tivoli Storage Manager Component**
AIX	Server Storage Agent Device Driver
HP-UX	Server Storage Agent Device Driver
Linux	Server Storage Agent Device Driver
Solaris	Server Storage Agent Device Driver
Windows	Server Storage Agent Device Driver Management Console

NOTE: tsmdiag utility is not installed with the Tivoli Storage Manager device driver, or with the Tivoli Storage Manager Management Console, but if these components are installed on the system where tsmdiag is executed, tsmdiag will collect data for these components.

Where to find and how to use the Tivoli Storage Manager diagnostic utility.

The tsmdiag utility is installed with supported Tivoli Storage Manager components, under the component's install directory. For the Tivoli Storage Manager Server V5.3 on Linux, for instance, tsmdiag utility can be found under opt/tivoli/tsm/server/bin/tsmdiag/. For any other version of Tivoli Storage Manager server and storage agent, you can download tsmdiag from ftp://ftp.software.ibm.com/storage/tivoli-storage-management/tsmdiag/

Usage:

tsmdiag [options]

Where options can be any combination of the following:

Supported OS Environments and Components

Option	Description	Default Value
-id *adminName*	Tivoli Storage Manager server administrator id	admin
-pa *adminPwd*	Tivoli Storage Manager server administrator password	admin
-tcpserveraddress *ipAddress*	Tivoli Storage Manager server tcp/ip name or address	localhost
-tcpport *portNumber*	Tivoli Storage Manager server tcp/ip port	1500
-results *resultsDir*	Directory for created files	($cwd)/results
-v	verbosely list activity processed	non-verbose
-?	Display usage information	N/A

Examples:

./tsmdiag
./tsmdiag -v
./tsmdiag -id tsmadmin -pa pwd4u -results /home/tsmdiag/results_oct29

NOTE: For Tivoli Storage Manager server environments on Unix, if you're running multiple servers on the same system, you must set the DSMSERV_DIR and DSMSERV_CONFIG environment variables. Failure to do this will prevent tsmdiag from finding data for the desired server.

Where to find the data collected by tsmdiag

Once tsmdiag has completed execution, go to the either the default results directory, $CWD/results, or the directory that you specified via the *-results* command line option. There will be a file there called **tsmdiag_results.tar** for Unix environments and **tsmdiag_results.zip** for Windows environments. This is the file that needs to be submitted with the PMR.

Default data collected by tsmdiag

The default data collected by tsmdiag is specified in the configuration file named configfile.txt. More data can be added to this file as needed.

Each line in the configuration file consists of five fields. Each field is separated by comma. A line that begins with an asterisk(*) is considered a comment and is ignored by the tsmdiag utility.

The first field of a line in the configuration file denotes the operating system where tsmdiag will be executed. Valid platforms are Linux, AIX, HP, SUN, and Windows. The second field in the line denotes an action. Valid actions are as follows:

Valid Actions	
Action Name	**Action Performed**
SYSCOMMAND	Run a command on the system.
COPY	Copy a specified file.
COMPCOMMAND	Run a command on a specified Tivoli Storage Manager component.
REGISTRY	Get the value for a given registry key. (Windows only.)

The third field denotes the Tivoli Storage Manager component (i.e. component for which tsmdiag will collect data).

Valid components are:

SERVER
STAGENT
BACLIENT (Unix only)
MGMTCONSOLE (Windows only)
DEVDRIVER
SYSTEM
CLIENAPI (Unix only)

The fourth field denotes the data that the action will act upon. The data field can be a file name, a system command, a command for a Tivoli Storage Manager component, a registry key, etc, depending on the action.

The last field is optional, and it denotes the location of where the data collected is to be stored.

Main Menu | Understanding TSM messages

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

- Messages that begin with prefix **ANE** and are in range 4000-4999 originate from the backup-archive client. These messages (or events) are sent to the server for distribution to various event-logging receivers.
- The client may send statistics to the server providing information about a backup or restore. These statistics are informational messages that may be enabled or disabled to the various event-logging receivers. These messages are not published in this manual.
- Messages that begin with prefix **ANR** originate from the server or storage agent.
- Messages that begin with prefix **ANS** are from one of the following clients:
  - Administrative clients
  - Application program interface clients
  - Backup-archive clients
  - Space Manager (HSM) clients
- Messages that begin with prefix **ACD** are from Data Protection for Lotus Domino(R).
- Messages that begin with prefix **ACN** are from Data Protection for Microsoft(R) Exchange Server.
- Messages that begin with prefix **ACO** are from Data Protection for Microsoft SQL Server.
- Messages that begin with prefix **ANU** are from Data Protection for Oracle.
- Messages that begin with prefix **BKI** are from Data Protection for mySAP for DB2(R) UDB and Data Protection for mySAP for Oracle.
- Messages that begin with prefix **DKP** and are in range 0001-9999 are from Data Protection for WebSphere(R) Application Server.
- Messages that begin with prefix **EEO** and are in range 0000-9999 are from Data Protection for IBM Enterprise Storage Server(R) (ESS) for Oracle.
- Messages that begin with prefix **EEP** and are in range 0000-9999 are from Data Protection for Enterprise Storage Server (ESS) for DB2 UDB.
- Messages that begin with prefix **IDS** and are in range 1000-1999 are from Data Protection for Enterprise Storage Server (ESS) for mySAP.

Message Format>

The following examples describe the Tivoli Storage Manager message format:

The callouts to the right identify each part of the format.

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

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

How to Read a Return Code Message

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

Example One for QUERY EVENT Command

Example Two for DEFINE VOLUME Command

Main Menu | Authority problems for an ISC User

Authority problems for an ISC User.

The Settings tab of the Integrated Solutions Console (ISC) is used to perform the following tasks:

- Create, delete, and view information for users and groups.
- Assign groups or users permissions to access resources.

You can use the pages on the Settings tab to view settings for access permissions, users, and user groups. You can also define new users and assign the users to groups that are authorized to access console resources. ISC provides help for the Settings tab. This help will assist you when navigating the Settings tab and its pages.

ISC access permissions are role-based. Each role represents a set of actions that an authorized user can perform on a specific resource. A role is always defined in relationship to a resource. In ISC a resource can be a page, a portlet on a page, or an organizational node in the navigation tree. To simplify assigning access permissions, ISC supports propagating and inheriting roles.

ISC users can access only the pages for which they are authorized. In addition, the user must have permission to access one or more portlets on the page. Otherwise, the page will display to the user as a blank page. If a page contains multiple portlets, the user can be given access to some or all of the portlets. Access permissions can be set for each individual portlet or for the concrete portlet application to which the portlet belongs.

The installation of the Administration Center creates a new ISC user group for the Tivoli Storage Manager Administration Center. The new user group is called TSM_AdminCenter. New ISC users should be created in the group and they will have access to the Administration Center.

The Administration Center resource is composed of several concrete portlet applications and all of its portlets are associated with one of these applications. The TSM_AdminCenter group is assigned the Administrator role for these concrete portlet applications. Administrator is the highest level of authority for a resource. Users and groups that have Administrator authority can perform all of the permitted actions.

The following steps describe how to create an ISC user that has access to the Tivoli Storage Manager Administration Center:

- Log into the ISC console.
- Select the Settings Tab.
- Select User and Group Management.
    - Search for the phrase: User Groups.
    - Search for the phrase: All Available.
    - Click the Search button.
    - The table should now include the TSM_AdminCenter group.
    - Click on TSM_AdminCenter.
        - Click on the New User button.
        - Fill in the form and click the OK button.

If an ISC user does not have access to the Administration Center, the following is a list of items that should be verified.

- Ensure that group to which the user belongs has the necessary permissions to the resource. Ideally the group is TSM_AdminCenter, but that is not a requirement.
    - Select the Settings Tab.
    - Select User and Group Permissions.
    - Select User Groups.
        - Search for the phrase: All Available.
        - Click the Search button.
        - The table should now include the User Group of the given user.
        - Click on the Select Resource Type button of that group.
        - Select Portlet Applications and verify the permissions of the user.
        - Select Portlets and verify the permissions of the user.
        - Select Pages and verify the permissions of the user.
- Ensure that the TSM_AdminCenter group does exist.
- Ensure that TSM_AdminCenter group still has the necessary permissions to all the concrete portlet applications of the Administration Center
- Ensure that the user has access to all the portlets and pages of the Administration Center.
    - Select the Settings Tab.
    - Select User and Group Permissions.
    - Select Users.
        - Search for keyword: uid.
        - Search for keyword: user_id.
        - Click the Search button.
        - The table should now include the specified user.
        - Click on Select Resource Type.
        - Select Portlet Applications and verify the permissions of the user.
        - Select Portlets and verify the permissions of the user.
        - Select Pages and verify the permissions of the user.

Main Menu | Integrated Solutions Console Server crash

Integrated Solutions Console Server crash.

Integrated Solutions Console (ISC) is built on top of a WebSphere Application Server and WebSphere Portal. If the Websphere Portal server were to crash, ISC would no longer be accessible from a browser. See Unable to access from a Web browser for additional information.

Before restarting the ISC server, you should gather information about the crash so that the problem can be reported.

WebSphere includes a number of troubleshooting tools that are designed to help you isolate the source of problems. Many of these tools are designed to generate information to be used by IBM Support, and their output might not be understandable by the customer.

This section only discusses tools that are bundled with ISC. A wide range of tools which address a variety of problems are available from the WebSphere Technical Support website.

Selecting a diagnostic tool.

Select the appropriate tool for the task.

- For more information on the capacities of the supplied troubleshooting tools, see the relevant articles in this section.
- Run the tool as described in the relevant article.
- Contact Websphere Support for assistance in deciphering the output of the tool.

Collector tool.

The collector tool gathers information about your WebSphere server installation and packages it in a Java archive (JAR) file that you can send to Websphere Customer Support to assist in determining and analyzing your problem. Information in the JAR file includes logs, property files, configuration files, operating system and Java data, and the presence and level of each software prerequisite.

There are two ways to run the collector tool. The collector tool can be executed to collect summary data or to traverse the system to gather relevant files and command results. The collector tool produces a Java archive (JAR) file of information needed to determine and solve a problem. The collector summary option produces a lightweight collection of version and other information that is useful when first reporting the problem to Websphere Support.

There are two phases of using the collector tool. The first phase runs the collector tool on your WebSphere server and produces a Java archive (JAR) file . The Websphere Support team performs the second phase, which is analyzing the Java archive (JAR) file that the collector program produces.

The collector program runs to completion as it creates the JAR file, despite any errors that it might find. Errors might include missing files or commands. The collector tool collects as much data in the JAR file as possible.

Running the collector tool

1. Log on to the system as root (or Administrator on Windows).
2. Verify that Java 1.2.2 or higher is available in the path.
3. The collector program requires Java to run. It also collects data about the Java Development Kit (JDK) in which it runs. If there are multiple JDKs on the system, verify that the JDK that the WebSphere server uses is the one in the path for the collector program.
4. Verify that all necessary information is in the path being used by the collector program and that you are not running the program from within the ISC installation root directory.
    1. If this system is a Linux or UNIX-based platform, verify that the path contains:
        - /bin
        - /sbin
        - /usr/bin
        - /usr/sbin
    2. If this system is a Windows platform, include regedit in the path.
5. The collector tool is located in the directory: isc_root/AppServer/bin
6. Create a work directory where you can start the collector program.
7. Make the work directory the current directory.
   The collector program writes its output JAR file to the current directory. The program also creates and deletes a number of temporary files in the current directory. Creating a work directory to run the collector program avoids naming collisions and makes cleanup easier. You cannot run the collector tool in a directory under the ISC installation directory.
8. Run the collector program by entering the command: collector from the command line.
   For example: c:\work> collector

Using the collector command with no additional parameters gathers one copy of the node data and data from each server in the node, and stores them in a single JAR output file. To gather data from a specific server in the node, use the command collector servername, where servername is the name of the problem server.The name of the ISC server is ISC_Portal

Results.

The collector program creates a log file, Collector.log, and an output JAR file in the current directory.

The name of the JAR file is based on the hostname and package of the Websphere server, in the format: hostname-ND-WASenv.jar or hostname-Base-WASenv.jar. For example, if you run the collector tool on the server ws-laceweb within a Network Deployment cell, the filename is ws-laceweb-ND-WASenv.jar.

The Collector.log log file is one of the files collected in the hostname-ND-WASenv.jar or hostname-Base-WASenv.jar file.

Log Analyzer tool (showlog)

The Log Analyzer takes one or more service or activity logs, merges all of the data, and displays the entries. Based on its symptom database, the tool analyzes and interprets the event or error conditions in the log entries to help you diagnose problems. The Log Analyzer has a special feature enabling it to download the latest symptom database from the IBM website.

About the Service or Activity Log.

The Websphere server creates the service or activity log file from the activity of the various WebSphere server components. The Log Analyzer is used to view the service or activity log file. The Log Analyzer can merge service or activity log files into one log file. The service or activity log file, activity.log, is a binary file in the logs directory of the install_root.

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

Using the Log Analyzer.

The Websphere server on which ISC runs does not include the Java administrative console, so there is no graphical interface available for viewing a service or activity log file. The alternate viewing tool, **showlog**, must be used to view the service or activity log file:

1. Change to the directory *your_isc_root/*AppServer/bin
   where *your_isc_root* is the root directory for your ISC installation.
2. Run the **showlog** tool with no parameters to display usage instructions:
   - On Windows systems, run showlog.bat
   - On UNIX systems, run showlog.sh

To direct the service or activity log (activity.log) contents to stdout, use the **showlog** activity.log command.

To dump the service or activity log to a text file for viewing with a text editor, use the **showlog** activity.log textFileName command.

Main Menu | Integrated Solutions Console Server has excessive memory consumption

Integrated Solutions Console Server has excessive memory consumption.

Integrated Solutions Console (ISC) is built on top of a WebSphere application server and WebSphere Portal. If it is determined that the Websphere server is using a large amount of memory, there are a few actions that can be taken.

To determine if a larger amount of memory is in use, the operating system tools that provide such information should be used. For example, on Windows, the Task Manager shows memory use.

- You should log off from ISC on a regular basis, or if you determine that the memory usage of the Websphere Server is high. Logging out at the end of each day should help minimize the memory use.

- Logging out instead of only closing the browser is recommended.
- The default value of Session Timeout is 30 minutes. Reducing this value to a lower number can help reduce memory consumption requirements. The value can be changed with the Administration Center Support Utility.