Tivoli® Storage Manager for UNIX and Linux Backup-Archive Clients
Version 6.2.2

Installation and User's Guide

IBM

Tivoli® Storage Manager for UNIX and Linux Backup-Archive Clients
Version 6.2.2

Installation and User's Guide

IBM

Note:

Before using this information and the product it supports, read the information in "Notices" on page 563.

Edition notice

This edition applies to Version 6.2.2 of IBM Tivoli Storage Manager (product numbers 5608-E01, 5608-E02, and 5608-E07), and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters. This edition replaces SC23-9791-03.

Contents

Tables

Preface

IBM® Tivoli® Storage Manager is a client-server licensed product that provides storage management services in a multiplatform computer environment. The backup-archive client program permits users to back up and archive files from their workstations or file servers to storage, and restore and retrieve backup versions and archived copies of files to their local workstations.

In addition to the backup-archive client, Tivoli Storage Manager includes the following components available on a variety of platforms:

- A *server program* that permits systems to perform as a backup and archive server for distributed workstations and file servers.

  The server program also supplies hierarchical storage management (HSM) services, and permits systems to perform as a migration server.

- An *administrative client program* that you can access from a Web browser or the command line. The program permits a Tivoli Storage Manager administrator to control and monitor server activities, define storage management policies for backup, archive and space management services, and set up schedules to perform those services at regular intervals.

- An *application program interface (API)* that permits you to enhance an existing application with storage management services. When an application is registered with a server as a client node, the application can back up, restore, archive, and retrieve objects from storage.

- A *Web backup-archive client* that permits an authorized administrator, help desk person, or end user to perform backup, restore, archive, and retrieve services using a Web browser on a remote system.

Associated with Tivoli Storage Manager, but sold separately, is the *Tivoli Storage Manager for Space Management* client program which was previously a feature of ADSM known as *Hierarchical Storage Manager (HSM)*. Tivoli Space Manager automatically migrates eligible files to storage to maintain specific levels of free space on local file systems and automatically recalls migrated files when they are accessed. It also permits users to migrate and recall specific files. This client program runs only on AIX®, HP-UX, Linux x86/x86_64, and Solaris (SPARC) operating systems.

The terms *hierarchical storage management* and *space management* have the same meaning throughout this publication.

Who should read this publication

This publication provides instructions for a user to install, configure, and use the Tivoli Storage Manager client.

Related information

Chapter 1, "Installing the Tivoli Storage Manager clients," on page 1

Chapter 2, "Configure the Tivoli Storage Manager client," on page 47

"UNIX and Linux client root and authorized user tasks" on page 47

Publications

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

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

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

Go to Tivoli Documentation Central to find information centers that contain official product documentation for current and previous versions of Tivoli products, including Tivoli Storage Manager products at http://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Tivoli+Storage+Manager.

You can also order some related publications from the IBM Publications Center Web site. The Web site provides information about ordering publications from countries other than the United States. In the United States, you can order publications by calling 1-800-879-2755.

Tivoli Storage Manager publications

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

Table 1. IBM Tivoli Storage Manager troubleshooting and tuning publications

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

Table 2. Tivoli Storage Manager server publications

Publication title	Order number
IBM Tivoli Storage Manager for AIX Installation Guide	GC23-9781
IBM Tivoli Storage Manager for AIX Administrator's Guide	SC23-9769
IBM Tivoli Storage Manager for AIX Administrator's Reference	SC23-9775
IBM Tivoli Storage Manager for HP-UX Installation Guide	GC23-9782
IBM Tivoli Storage Manager for HP-UX Administrator's Guide	SC23-9770
IBM Tivoli Storage Manager for HP-UX Administrator's Reference	SC23-9776
IBM Tivoli Storage Manager for Linux Installation Guide	GC23-9783
IBM Tivoli Storage Manager for Linux Administrator's Guide	SC23-9771
IBM Tivoli Storage Manager for Linux Administrator's Reference	SC23-9777
IBM Tivoli Storage Manager for Sun Solaris Installation Guide	GC23-9784
IBM Tivoli Storage Manager for Sun Solaris Administrator's Guide	SC23-9772
IBM Tivoli Storage Manager for Sun Solaris Administrator's Reference	SC23-9778
IBM Tivoli Storage Manager for Windows Installation Guide	GC23-9785

Table 2. Tivoli Storage Manager server publications (continued)

Publication title	Order number
IBM Tivoli Storage Manager for Windows Administrator's Guide	SC23-9773
IBM Tivoli Storage Manager for Windows Administrator's Reference	SC23-9779
IBM Tivoli Storage Manager Server Upgrade Guide	SC23-9554
IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack	SC27-2828

Table 3. Tivoli Storage Manager storage agent publications

Publication title	Order number
IBM Tivoli Storage Manager for SAN for AIX Storage Agent User's Guide	SC23-9797
IBM Tivoli Storage Manager for SAN for HP-UX Storage Agent User's Guide	SC23-9798
IBM Tivoli Storage Manager for SAN for Linux Storage Agent User's Guide	SC23-9799
IBM Tivoli Storage Manager for SAN for Sun Solaris Storage Agent User's Guide	SC23-9800
IBM Tivoli Storage Manager for SAN for Windows Storage Agent User's Guide	SC23-9553

Table 4. Tivoli Storage Manager client publications

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

Table 5. Tivoli Storage Manager Data Protection publications

Publication title	Order number
IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for DB2	SC33-6341
IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for Oracle	SC33-6340

Support information

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

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

Getting technical training

Information about Tivoli technical training courses is available online.

Visit the following Web sites for training information:

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

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

Searching knowledge bases

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

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

Searching the Internet

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

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

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

From the Support Web site, you can search various resources including:
- IBM technotes
- IBM downloads
- IBM Redbooks® publications
- IBM Authorized Program Analysis Reports (APARs)

Select the product and click **Downloads** to search the APAR list.

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

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

To share your experiences and learn from others in the Tivoli Storage Manager user community, go to the Tivoli Storage Manager wiki at http://www.ibm.com/ developerworks/wikis/display/tivolistoragemanager.

Using IBM Support Assistant

IBM Support Assistant is a complimentary software product that helps you with problem determination. You can install the stand-alone IBM Support Assistant application on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use.

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

The product-specific plug-in modules provide you with the following resources:
* Support links
* Education links
* Ability to submit problem management reports

Find add-ons for specific products here: http://www.ibm.com/support/ docview.wss?&uid=swg27012689.

Finding product fixes

A product fix to resolve your problem might be available from the IBM Software Support Web site.

You can determine what fixes are available by checking the IBM Software Support Web site at http://www.ibm.com/support/entry/portal/.
* If you previously customized the site based on your product usage:
  1. Click the link for your IBM Tivoli Storage Manager product, or one of the other Tivoli Storage Manager components for which you want to find a fix.
  2. Click **Downloads**, and then click **Fixes by version**.
* If you have not customized the site based on your product usage, click **Downloads** and search for your product.

Receiving notification of product fixes

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

To sign up to receive notifications about IBM products, follow these steps:
1. From the support page at http://www.ibm.com/support/entry/portal/, click **My notifications** in the notifications module.
2. Sign in using your IBM ID and password. If you do not have an ID and password, click **register now** above the IBM ID and password.
3. Click the **Subscribe** tab to select your product family and click **Continue**.
4. Select the type of information that you want to receive, and add your personal preferences. You can specify how you want to be notified, how often, and you can also optionally select a folder for the notifications.
5. Click **Submit**.
6. For notifications for other products, repeat steps 4 and 5.

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

Contacting IBM Software Support

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

Before you contact IBM Software Support, follow these steps:

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

Then see "Submitting the problem to IBM Software Support" on page xvii for information on contacting IBM Software Support.

Setting up a subscription and support contract

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

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

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

Determining the business impact

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

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

Describing the problem and gather background information

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

To save time, know the answers to these questions:

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

Submitting the problem to IBM Software Support

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

Online

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

By phone

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

Conventions used in this publication

This publication uses the following typographical conventions:

Example	Description
autoexec.ncf hsmgui.exe	A series of lowercase letters with an extension indicates program file names.
DSMI_DIR	A series of uppercase letters indicates return codes and other values.
dsmQuerySessInfo	Boldface type indicates a command that you type on a command line, the name of a function call, the name of a structure, a field within a structure, or a parameter.
timeformat	Boldface italic type indicates a Tivoli Storage Manager option. The bold type is used to introduce the option, or used in an example.
dateformat	Italic type indicates an option, the value of an option, a new term, a placeholder for information you provide, or for special emphasis in the text.
`maxcmdretries`	Monospace type indicates fragments of a program or information as it might appear on a display screen, such a command example.
plus sign (+)	A plus sign between two keys indicates that you press both keys at the same time.

Reading syntax diagrams

To read a syntax diagram for entering a command, follow the path of the line.
Read from left to right and from top to bottom.

- The ►►── symbol indicates the beginning of a syntax diagram.
- The ──► symbol at the end of a line indicates that the syntax diagram continues on the next line.

- The ►── symbol at the beginning of a line indicates that a syntax diagram continues from the previous line.
- The ──►◄ symbol indicates the end of a syntax diagram.

Syntax items, such as a keyword or a variable, can be:
- On the line (required element)
- Above the line (default element)
- Below the line (optional element)

Symbols

Enter these symbols *exactly* as they appear in the syntax diagram.
- * Asterisk
- *{ }* Braces
- *:* Colon
- *,* Comma
- = Equal Sign
- - Hyphen
- *()* Parentheses
- . Period
- Space
- ″ quotation mark
- ′single quotation mark

Variables

Italicized lowercase items such as *<var_name>* indicate variables. In this example, you can specify a *<var_name>* when you enter the **cmd_name** command.

►►──cmd_name──*<var_name>*──────────────────────────────────►◄

Repetition

An arrow returning to the left means that the item can be repeated. A character within the arrow means that you must separate repeated items with that character.

```
        ┌─,─────┐
►►──────▼─repeat─┴──────────────────────────────────────────►◄
```

A footnote (1) by the arrow refers to a limit that tells how many times the item can be repeated.

```
        ┌─,───────┐
        │ (1)     │
►►──────▼───repeat─┴─────────────────────────────────────────►◄
```

Notes:

1   Specify *repeat* up to 5 times.

Required choices

When two or more items are in a stack and one of them is on the line, you *must* specify one item.

In this example, you must choose A, B, or C.

```
►►─cmd_name─┬─A─┬─────────────────────────────────────────────►◄
            ├─B─┤
            └─C─┘
```
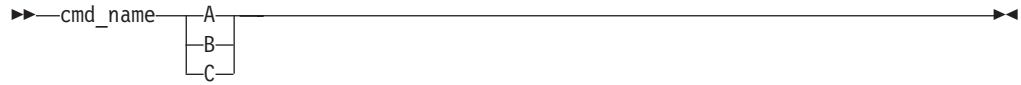
Optional choices

When an item is *below* the line, that item is optional. In the first example, you can select A or nothing at all.

```
►►─cmd_name─┬───┬─────────────────────────────────────────────►◄
            └─A─┘
```

When two or more items are in a stack below the line, all of them are optional. In the second example, you can choose A, B, C, or nothing at all.

```
►►─cmd_name─┬───┬─────────────────────────────────────────────►◄
            ├─A─┤
            ├─B─┤
            └─C─┘
```
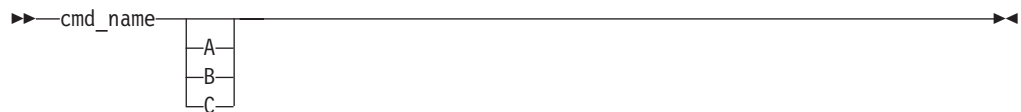
Repeatable choices

A stack of items followed by an arrow returning to the left indicates that you can select more than one item, or in some cases, repeat a single item.

In this example, you can select any combination of A, B, or C.

```
              ┌─,─┐
►►─cmd_name─▼─┬─A─┬─┴──────────────────────────────────────────►◄
              ├─B─┤
              └─C─┘
```

Defaults

Defaults are above the line. The default is selected unless you override it, or you can select the default explicitly. To override the default, include an option from the stack below the line.

In this example, A is the default. Select either B or C to override A.

```
            ┌─A─┐
►►─cmd_name─┼───┼─────────────────────────────────────────────►◄
            ├─B─┤
            └─C─┘
```

New for Version 6.2

Tivoli Storage Manager backup-archive client version 6.2 contains many new features and changes.

New for Version 6.2.2

This section summarizes changes that were made to Tivoli Storage Manager Version 6.2.2.

The following features are new for IBM Tivoli Storage Manager in Version 6.2.2:

Unicode support for NetApp snapshot-assisted progressive incremental backup

> If you run snapshot difference incremental backups, some files might not be backed up because their names contain non 7 bit ASCII characters. Support for non 7 bit ASCII file names has been added to NetApp Data ONTAP version 7.3.3 (or later) or 8.1 (or later). If you migrated to a file server version that supports non 7 bit ASCII file names, use the option createnewbase to create a new base snapshot and use it as a source to run a full incremental to ensure that the backup of any files that might have been skipped previously.

> See "Createnewbase" on page 287 for more information.

Related reference

"Createnewbase" on page 287

New for IBM Tivoli Storage Manager Version 6.2

This section summarizes changes that were made to Tivoli Storage Manager Version 6.2.

Several features in IBM Tivoli Storage Manager Version 6.2 are new for previous Tivoli Storage Manager users.

The following features are new for IBM Tivoli Storage Manager in Version 6.2:

Extended Secure Sockets Layer (SSL) support

> You can now use SSL with HP-UX, Linux, and Solaris operating systems.

> You can also use certificates that are signed by an independent software vendor. In addition to self-signed IBM Tivoli Storage Manager certificates, certificates signed by a Certificate Authority can also be used.

> See "Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer" on page 65 for more information.

Client-side data deduplication support

> Data deduplication reduces storage needs by eliminating redundant data. Client-side data deduplication reduces the amount of data sent over the network. The processing that is required to remove duplicate data on the server is eliminated.

> See these options for more information:
> - "Deduplication" on page 293
> - "Dedupcachepath" on page 291

- "Dedupcachesize" on page 292
- "Enablededupcache" on page 312
- "Exclude options" on page 321
- "Ieobjtype" on page 345
- "Include options" on page 351

z/OS® UNIX System Services support

The Tivoli Storage Manager z/OS UNIX System Services client in Version 6.2 is at the Version 6.1 functional level.

For information about the z/OS UNIX System Services client, see http://publib.boulder.ibm.com/infocenter/tsminfo/v6.

Tivoli Storage Manager FastBack client support

The integration of Tivoli Storage Manager and the Tivoli Storage Manager Linux FastBack client provides short-term retention for backing up and archiving the latest snapshots.

See these options and commands for more information:
- "Backmc" on page 277
- "Fbbranch" on page 326
- "Fbclientname" on page 326
- "Fbpolicyname" on page 328
- "Fbreposlocation" on page 329
- "Fbserver" on page 330
- "Fbvolumename" on page 332
- "Archive FastBack" on page 467
- "Backup FastBack" on page 469
- "Set Password" on page 555

Mac OS X 10.6

The Mac client now supports Mac OS X 10.6, Snow Leopard. Compressed system files are fully supported.

Txnbytelimit option limit is increased

The limit on the txnbytelimit option has been increased to 32 GB to improve tape drive performance.

See the "Txnbytelimit" on page 446 option for more information.

Related concepts

"Client-side data deduplication" on page 74

"Deduplicate data attribute" on page 234

Related tasks

"Configuring the client for data deduplication" on page 77

"Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer" on page 65

"Excluding files from data deduplication" on page 79

Chapter 1. Installing the Tivoli Storage Manager clients

The Tivoli Storage Manager backup-archive client helps you protect information on your workstations.

You can maintain backup versions of your files that you can restore if the original files are damaged or lost. You can also archive files that you do not currently need, preserve them in their current state, and retrieve them when necessary.

The Tivoli Storage Manager clients work in conjunction with the Tivoli Storage Manager server. Contact your Tivoli Storage Manager server administrator to obtain backup or archive access to the server, or refer to the server publications to install and configure a Tivoli Storage Manager server.

Upgrading from earlier versions

The following sections explain what you need to do if you are upgrading to Tivoli Storage Manager Version 6.2.2 from a previous version.

Upgrade path for clients and servers

As part of a migration plan from Tivoli Storage Manager Version 5.5 or 6.1 to Tivoli Storage Manager Version 6.2, Tivoli Storage Manager clients and servers can be upgraded at different times. This configuration is supported as long as the older version is supported.

To help prevent disruption to your backup and archive activities during the migration, follow these guidelines:

- A Tivoli Storage Manager Version 5.5 or 6.1 client can perform backup, restore, archive, retrieve, and query functions to a Tivoli Storage Manager Version 6.2 server.
- A Tivoli Storage Manager Version 6.2.2 client can perform backup, restore, archive, retrieve, and query functions to a Tivoli Storage Manager Version 5.5 or 6.1 server.
- If you back up or archive data from a Tivoli Storage Manager Version 5.5 or 6.1 client to any Tivoli Storage Manager server, you can restore or retrieve that data using a Tivoli Storage Manager Version 6.2.2 client.
- If you back up or archive data from a Tivoli Storage Manager Version 6.2.2 client, you cannot restore or retrieve that data using a Tivoli Storage Manager Version 6.1 or earlier client.
- All administrative clients can administer Tivoli Storage Manager Version 5.5, 6.1, and 6.2 servers.
- NAS backups performed on a Tivoli Storage Manager Version 6.2 server using the **BACKUP NODE** server command can only be restored using the **RESTORE NODE** server command or a Tivoli Storage Manager Version 6.2.2 client.

Refer to "Cluster environment configuration and use" on page 81 for details about using the Tivoli Storage Manager backup-archive client in the cluster environment, including the following information:

- Migrating your current AIX HACMP™ setups

- Configuring the Tivoli Storage Manager backup-archive client in a cluster environment
- Enabling the Tivoli Storage Manager Web client in a cluster environment

Related reference

Additional upgrade information

This section explains some additional information that you need to know when upgrading your Tivoli Storage Manager client.

Determine which AIX package you need based on the following information:
- If you have a 64-bit system and you do not use HSM, upgrade to the 64-bit backup-archive client.
- If you have a 32-bit system or you use enhanced journaled file system (JFS2) hierarchical storage management (HSM), install the JFS2 backup-archive client.
- If you use GPFS™ HSM, install the GPFS backup-archive client

The Version 6.2 HP Itanium and Solaris SPARC backup-archive client packages do not contain HSM clients. The HSM clients are in a separate package at the Version 6.1 functional level for HP and Solaris.

When you install the Web client, you must install the Web-client language files that correspond to those languages you want to use.

To view the non-English online help from the Web Client applet, you must install the language versions of the help files on the agent, the system where the Tivoli Storage Manager backup-archive client was installed. If the language versions are not installed or are not available, the online help is displayed in English.

See the `client_message.chg` file in the client package for a list of new and changed messages since the previous Tivoli Storage Manager release.

Client environment requirements

This section contains Tivoli Storage Manager client environment information, components, and hardware and software requirements.

The following list shows the location of the environment prerequisites for each supported platform.
- "AIX client environment" on page 3
- "HP-UX Itanium 2 client environment" on page 5
- "Linux on POWER client environment" on page 6
- "Linux x86/x86_64 client environment" on page 8
- "Linux on System z client environment" on page 10
- "Mac OS X client environment" on page 11
- "Solaris client environment" on page 13
- "NDMP support requirements (Extended Edition only)" on page 14

For current information concerning the client environment prerequisites for all Tivoli Storage Manager supported client platforms, go to the Web site at: http://www.ibm.com/software/tivoli/products/storage-mgr/product-links.html

AIX client environment

This section contains client environment information, Tivoli Storage Manager client components, and hardware and software requirements for the AIX platform.

AIX client installable components

The Tivoli Storage Manager command-line, Java GUI, Web backup-archive, API, and Hierarchical Storage Management comprise the AIX backup-archive client installable components.

You can install the following components with Tivoli Storage Manager Version 6.2:
- Tivoli Storage Manager backup-archive client
- Tivoli Storage Manager administrative client
- Tivoli Storage Manager backup-archive Java Graphical User Interface (Java GUI)
- Tivoli Storage Manager Web backup-archive client

  Important: Some components are bundled together and cannot be installed individually.
- Tivoli Storage Manager API (32-bit and 64-bit)
- Tivoli Storage Manager XOpen API (32-bit)
- Tivoli Storage Manager Hierarchical Storage Management client

Important: The IBM Tivoli Space Manager HSM Client for AIX is available on the client DVD. However, it is a separate product, and therefore subject to the fees and licensing required by Tivoli and IBM.

System requirements for AIX clients

The Tivoli Storage Manager AIX client requires a minimum amount of hardware, disk space, memory, and software.

Hardware requirements

Table 6 describes the minimum hardware requirements that are needed for your AIX client.

Table 6. AIX client hardware requirements

Type of hardware	Hardware requirements
Hardware	- A RISC System/6000, IBM System p®, System i®, or compatible hardware, as supported by AIX - For HACMP: One additional network adapter for each workstation - A minimum screen resolution of 800 by 600 pixels is required to display the Tivoli Storage Manager Backup-Archive Client Java GUI and Web GUI.

Table 6. AIX client hardware requirements (continued)

Type of hardware	Hardware requirements
Disk space	• Backup-archive client with 32-bit API and 32-bit GSKit: 92 MB, used with the HSM clients • Backup-archive client with 64-bit API and 64-bit GSKit: 92 MB • 64-bit API with 64-bit GSKit: 29 MB • Web client: 32 MB • Non-English language packs: 97 MB • JFS2 HSM client: 202 MB • GPFS HSM client: 170 MB • Filepath Kernel extension: 3 MB
Memory	512 MB

Software requirements

Table 7 describes the minimum software requirements that are needed for your AIX client.

Table 7. AIX client software requirements

Type of software	Minimum software requirements
Backup-archive client	• AIX 6.1 (does not have a 32-bit kernel) • AIX 5.3 (32-bit or 64-bit kernel mode) • AIX GPFS support: – GPFS 3.2, 3.2.1, and 3.3
Applications	Java JRE 5 or 6 for the Java GUI
Web browser	A Firefox 3.0.14 or higher browser for the Web client and to access online help and documentation

AIX client communication methods

The TCP/IP and shared memory communication methods are available for the AIX backup-archive client.

You can use the following communication methods with the Tivoli Storage Manager Version 6.2 AIX client:

Table 8. AIX client communication methods

To use this communication method:	Install this software:	To connect to these Tivoli Storage Manager servers:
TCP/IP	TCP/IP (Standard with supported AIX platforms)	AIX, HP-UX, Linux, Solaris, Windows
Shared Memory	TCP/IP (Standard with supported AIX platforms)	AIX

Backup-archive client features that are available on AIX platforms

This topic lists the features that are supported on the different AIX platforms.

Table 9. Supported features on AIX platforms

Features	AIX 32-bit JFS2	AIX 32-bit GPFS	AIX 64-bit
Backup-archive command-line and GUI	yes	yes	yes
HSM with JFS2	yes	no	no
HSM with GPFS	no	yes	no
Journal-based backup	yes	no	yes
NetApp Snapshot Difference (snapdiff option)	no	no	yes
LAN-free operations	yes	yes	yes
Online image backup	yes	yes	yes
Offline image backup	yes	yes	yes

HP-UX Itanium 2 client environment

This section contains client environment information, Tivoli Storage Manager client components, and hardware and software requirements for the HP-UX Itanium 2 platform.

HP-UX Itanium 2 client installable components

The Tivoli Storage Manager command-line, Java GUI, Web backup-archive, API, and Hierarchical Storage Management comprise the HP-UX Itanium 2 backup-archive client installable components.

You can install the following components with Tivoli Storage Manager Version 6.2:
- Tivoli Storage Manager backup-archive client
- Tivoli Storage Manager administrative client
- Tivoli Storage Manager backup-archive Java Graphical User Interface (Java GUI)
- Tivoli Storage Manager Web backup-archive client
- Tivoli Storage Manager API (32-bit and 64-bit)
- Tivoli Storage Manager XOpen API (32-bit)
- Tivoli Storage Manager Hierarchical Storage Management client

Important: The IBM Tivoli Space Manager HSM Client for HP-UX Itanium 2 is available on the client DVD. However, it is a separate product, and therefore subject to the fees and licensing required by Tivoli and IBM.

System requirements for HP-UX Itanium 2 clients

The Tivoli Storage Manager HP-UX Itanium 2 client requires a minimum amount of hardware, disk space, memory, and software.

Hardware requirements

Table 10 on page 6 describes the minimum hardware requirements that are needed for your HP-UX Itanium 2 client.

Table 10. HP-UX Itanium 2 client hardware requirements

Type of hardware	Hardware requirements
Hardware	• An HP 9000 Series 700 or 800 workstation or server • A minimum screen resolution of 800 by 600 pixels is required to display the Tivoli Storage Manager Backup-Archive Client Java GUI and Web GUI
Disk space	610 MB
Memory	512 MB

Software requirements

Table 11 describes the minimum software requirements that are needed for your HP-UX Itanium 2 client.

Table 11. HP-UX Itanium 2 client software requirements

Type of software	Minimum software requirements
One of the following:	• HP-UX 11i v2 (32-bit or 64-bit) • HP-UX 11i v3
Applications	Java JRE 5 or 6 for the Java GUI: http://www.hp.com/products1/unix/java/index.html
Web browser	A Firefox 3.0.14 or higher browser for the Web client and to access online help and documentation

HP-UX Itanium 2 client communication methods

The TCP/IP and shared memory communication methods are available for the HP-UX Itanium 2 backup-archive client.

You can use the following communication methods with the Tivoli Storage Manager Version 6.2 HP-UX Itanium 2 client:

Table 12. HP-UX Itanium 2 client communication methods

To use this communication method:	Install this software:	To connect to these Tivoli Storage Manager servers:
TCP/IP	TCP/IP (Standard with HP-UX)	AIX, HP-UX, Linux, Solaris, Windows
Shared Memory	TCP/IP (Standard with HP-UX)	HP-UX Itanium 2

Linux on POWER client environment

This section contains client environment information, Tivoli Storage Manager client components, and hardware and software requirements for the Linux on POWER® client platforms.

Linux on POWER client installable components

The Tivoli Storage Manager command-line, Java GUI, Web backup-archive, and API comprise the Linux on POWER backup-archive client installable components.

You can install the following components with Tivoli Storage Manager Version 6.2:

- Tivoli Storage Manager backup-archive client
- Tivoli Storage Manager administrative client
- Tivoli Storage Manager backup-archive Java Graphical User Interface (Java GUI)
- Tivoli Storage Manager Web backup-archive client
- Tivoli Storage Manager API (32-bit and 64-bit)

System requirements for Linux on POWER clients

The Tivoli Storage Manager Linux on POWER clients require a minimum amount of hardware, disk space, memory, and software.

Hardware requirements

Table 13 describes the minimum hardware requirements that are needed for your Linux on POWER client.

Table 13. Hardware requirements

Type of hardware	Hardware requirements
Hardware	The following: • Linux on POWER (64–bit mode) • IBM System p or IBM System i, with POWER processors, as supported by Linux distribution • OpenPower®
Disk space	For the backup-archive client: • 80 MB
Disk space	For the API: • 20 MB for the API 32-bit • 20 MB for the API 64-bit
Memory	For the backup-archive client: • 64 MB
Memory	For the API: • 64 MB

The GUI interface requires at least a 800x600x16 resolution to display properly on the local console.

Software requirements

Table 14 describes the minimum software requirements that are needed for your Linux on POWER system.

Table 14. Software requirements

Type of software	Minimum software requirements
One of the following:	• SUSE Linux Enterprise Server 10 and 11 • Red Hat Enterprise Linux 5.0
Applications	Java 5 SR11 or Java 6 for the Java GUI **Important:** Java 5 versions below SR11 are not supported.
Web browser	A Firefox 3.0.14 or higher browser for the Web client and to access online help and documentation

For EXT2/EXT3/XFS ACL support, you need the acl-2.0.19-17 package.

The Tivoli Storage Manager client uses the libacl.so library, so it is searched for in the following locations:
- A colon-separated list of directories in the LD_LIBRARY_PATH user environment variable
- The list of libraries cached in /etc/ld.so.cache. /usr/lib, followed by /lib

The image backup and restore functions for LVM2 volumes require libdevmapper.so Version 1.01 or higher. If the library is at a lower level than required or it is not available, image function for LVM2 volumes fail. The libdevmapper.so library is distributed with the LVM2 packages.

Linux on POWER client communication methods
The TCP/IP and shared memory communication methods are available for the Linux on POWER backup-archive client.

You can use the following communication methods with the Tivoli Storage Manager Version 6.2 Linux on POWER client:

Table 15. Linux on POWER client communication methods

To use this communication method:	Install this software:	To connect to these Tivoli Storage Manager servers:
TCP/IP	TCP/IP (Standard with Linux)	AIX, HP-UX, Linux, Solaris, Windows
Shared Memory	TCP/IP (Standard with Linux)	Linux on POWER

Linux x86/x86_64 client environment

This section contains client environment information, Tivoli Storage Manager client components, and hardware and software requirements for the Linux on Intel x86 (Linux x86/x86_64) platform.

Linux x86/x86_64 client installable components
The Tivoli Storage Manager command-line, Java GUI, Web backup-archive, API, and Hierarchical Storage Management comprise the Linux on Intel x86 (Linux x86/x86_64) backup-archive client installable components.

You can install the following components with Tivoli Storage Manager Version 6.2:
- Tivoli Storage Manager backup-archive client
- Tivoli Storage Manager administrative client
- Tivoli Storage Manager backup-archive Java Graphical User Interface (Java GUI)
- Tivoli Storage Manager Web backup-archive client
- Tivoli Storage Manager API (32-bit and 64-bit)
- Tivoli Storage Manager Hierarchical Storage Management client

Important: The IBM Tivoli Space Manager HSM Client for HP-UX PA-RISC is available on the client DVD. However, it is a separate product, and therefore subject to the fees and licensing required by Tivoli and IBM.

System requirements for Linux x86/x86_64 clients

The Tivoli Storage Manager Linux x86/x86_64 clients require a minimum amount of hardware, disk space, memory, and software.

Hardware requirements

Table 16 describes the minimum hardware requirements that are needed for your Linux x86/x86_64 client.

Table 16. Hardware requirements

Type of hardware	Hardware requirements
Hardware	x86-based PC (Pentium or newer) or AMD64/EM64T processor
Disk space	The following minimum disk space: • 8 MB for the API 32-bit • 17 MB for the API 64-bit (this includes the API 32-bit) • 29 MB for the Backup-Archive client • 143 MB for the HSM client
Memory	512 MB

A minimum screen resolution of 800 by 600 pixels is required to display the Tivoli Storage Manager Backup-Archive Client Java GUI and Web GUI.

Software requirements

Table 17 describes the minimum software requirements that are needed for your Linux x86/x86_64 system.

Table 17. Software requirements

Type of software	Minimum software requirements
One of the following:	• Red Hat Enterprise Linux 5.0 • SUSE Linux Enterprise Server 10 and 11 (including Novell OES distributions) For SuSE Linux Enterprise Server 10 or 11 64-bit support, you need to install the **glibc-locale-32bit** package if you run the client on a locale other than en_US or EN_US.
Linux x86/x86_64 GPFS support:	GPFS 3.2, 3.2.1, and 3.3
Applications	Java JRE 5 or 6
Web browser	A Firefox 3.0.14 or higher browser for the Web client and to access online help and documentation

For GPFS ACL support on Linux x86/x86_64, the Tivoli Storage Manager client uses libgpfs.so library (which comes in the standard GPFS package). For Ext2/Ext3/XFS ACL support on Linux x86/x86_64, the Tivoli Storage Manager client uses libacl.so (which comes in the libacl package).

For EXT2/EXT3/XFS ACL support, you also need the acl-2.0.19-17 package.

The image backup and restore functions for LVM2 volumes require
`libdevmapper.so` Version 1.01 or later. If the library is at an earlier level than
required or it is not available, image function for LVM2 volumes fail. The
`libdevmapper.so` library is distributed with the LVM2 packages. For Linux 64-bit
support you need to install the 32-bit `libdevmapper.so` library.

Linux x86/x86_64 client communication methods

The TCP/IP and shared memory communication methods are available for the
Linux on Intel x86 (Linux x86/x86_64) backup-archive client.

You can use the following communication methods with the Tivoli Storage
Manager Version 6.2 Linux on Intel x86 (Linux x86/x86_64) client:

Table 18. Linux on Intel x86 client communication methods

To use this communication method:	Install this software:	To connect to these Tivoli Storage Manager servers:
TCP/IP	TCP/IP (Standard with Linux)	AIX, HP-UX, Linux, Solaris, Windows
Shared Memory	TCP/IP (Standard with Linux)	Linux x86/x86_64

Linux on System z client environment

This section contains client environment information, Tivoli Storage Manager client
components, and hardware and software requirements for the Linux on System z[®]
platform.

Linux on System z client installable components

The Tivoli Storage Manager command-line, Java GUI, Web backup-archive, and
API comprise the Linux on System z backup-archive client installable components.

You can install the following components with Tivoli Storage Manager Version 6.2:
*   Tivoli Storage Manager backup-archive client
*   Tivoli Storage Manager administrative client
*   Tivoli Storage Manager Web backup-archive client
*   Tivoli Storage Manager API (32-bit and 64-bit)

System requirements for Linux on System z clients

The Tivoli Storage Manager Linux System z clients require a minimum amount of
hardware, disk space, memory, and software.

Hardware requirements

Table 19 describes the minimum hardware requirements that are needed for your
Linux System z client.

Table 19. Hardware requirements

Type of hardware	Hardware requirements
Hardware	A Linux on System z 800 or 900 (64-bit mode)
Disk space	The following minimum disk space: • 6 MB for the API 32-bit • 11 MB for the API 64-bit (this includes the API 32-bit) • 20 MB for the Backup-Archive client

Table 19. Hardware requirements (continued)

Type of hardware	Hardware requirements
Memory	128 MB

A minimum screen resolution of 800 by 600 pixels is required to display the Tivoli Storage Manager Backup-Archive Client Web GUI.

Software requirements

Table 20 describes the minimum software requirements that are needed for your Linux System z system.

Table 20. Software requirements

Type of software	Minimum software requirements
One of the following:	• SUSE Linux Enterprise Server 10 or 11 Linux on System z. For SUSE Linux Enterprise Server 10 or 11 64-bit support, you must install the **glibc-locale-32bit** package, if you run the client on a locale other than en_US or EN_US. • Red Hat Enterprise Linux 5.0 Linux on System z
Applications	Java JRE 5 or 6 for the Java GUI
Web browser	A Firefox 3.0.14 or higher browser for the Web client and to access online help and documentation

For EXT2/EXT3/XFS ACL support, you need the acl-2.0.19-17 package.

The image backup and restore functions for LVM2 volumes require libdevmapper.so Version 1.01 or higher. If the library is at a lower level than required or it is not available, images for LVM2 volumes fail. The libdevmapper.so library is distributed with the LVM2 packages.

Linux on System z client communication methods

The TCP/IP and shared memory communication methods are available for the Linux on System z backup-archive client.

You can use the following communication methods with the Tivoli Storage Manager Version 6.2 Linux on System z client:

Table 21. Linux on System z client communication methods

To use this communication method:	Install this software:	To connect to these Tivoli Storage Manager servers:
TCP/IP	TCP/IP (Standard with Linux)	AIX, HP-UX, Linux, Solaris, Windows
Shared Memory	TCP/IP (Standard with Linux)	Linux on System z

Mac OS X client environment

This section contains client environment information, Tivoli Storage Manager client components, and hardware and software requirements for the Mac OS X client.

Mac OS X client installable components

The Tivoli Storage Manager command-line, Java GUI, Web backup-archive, and API comprise the Mac OS X backup-archive client installable components.

You can install the following components with Tivoli Storage Manager Version 6.2:
- Tivoli Storage Manager backup-archive client
- Tivoli Storage Manager administrative client
- Tivoli Storage Manager Web backup-archive client
- Tivoli Storage Manager API
- Tivoli Storage Manager Backup-Archive Java Graphical User Interface (GUI)

Tip: The dsmj shell script file for the Java GUI is installed in the following location:

`/Library/Application Support/tivoli/tsm/client/ba/bin`

System requirements for Mac OS X clients

The Tivoli Storage Manager Mac OS X clients require a minimum amount of hardware, disk space, memory, and software.

Hardware requirements

Table 22 describes the minimum hardware requirements that are needed for your Mac OS X client.

Table 22. Hardware requirements

Type of hardware	Hardware requirements
Hardware	A PowerPC® or Intel-based Macintosh
Disk space	100 MB
Memory	512 MB

A minimum screen resolution of 800 by 600 pixels is required to display the Tivoli Storage Manager Backup-Archive Client Java GUI and Web GUI.

Software requirements

Table 23 describes the minimum software requirements that are needed for your Mac OS X system.

Table 23. Software requirements

Type of software	Minimum software requirements
Backup-archive client	One of the following, where $x >= 0$: • Mac OS X, version 10.5.x • Mac OS X, version 10.6.x

Mac OS X client communication methods

The TCP/IP communication methods are available for the Mac OS X backup-archive client.

You can use the following communication methods with the Tivoli Storage Manager Version 6.2 Mac OS X client:

Table 24. Mac OS X client communication methods

To use this communication method:	Install this software:	To connect to these Tivoli Storage Manager servers:
TCP/IP	TCP/IP (standard with Mac OS X)	AIX, HP-UX, Linux, Solaris, Windows

Solaris client environment

This section contains client environment information, Tivoli Storage Manager client components, and hardware and software requirements for the Sun Solaris platform.

Sun Solaris client installable components

The Tivoli Storage Manager command-line, Java GUI, Web backup-archive, and API comprise the Sun Solaris backup-archive client installable components.

You can install the following components with Tivoli Storage Manager Version 6.2:
- Tivoli Storage Manager backup-archive client
- Tivoli Storage Manager administrative client
- Tivoli Storage Manager backup-archive Java Graphical User Interface (Java GUI)
- Tivoli Storage Manager Web backup-archive client
- Tivoli Storage Manager API (32-bit and 64-bit)
- Tivoli Storage Manager X/Open API (SPARC only)

System requirements for Sun Solaris clients

The Tivoli Storage Manager Sun Solaris clients require a minimum amount of hardware, disk space, memory, and software.

Hardware requirements

Table 25 describes the minimum hardware requirements that are needed for your Solaris client.

Table 25. Hardware requirements

Type of hardware	Hardware requirements
Hardware	One of the following: • A SPARCstation based on sun4u architecture or compatible workstation • x86 32-bit and 64-bit (AMD64/EM64T) based system that is explicitly compatible and fully capable of running the Solaris 10 x86 operating system
Disk space	219 MB **Note:** The installation procedure requires up to 200 MB of temporary disk space in the /var file system.
Memory	512 MB

A minimum screen resolution of 800 by 600 pixels is required to display the Tivoli Storage Manager Backup-Archive Client Java GUI and Web GUI.

Software requirements

Table 26 describes the minimum software requirements that are needed for your Solaris system.

Table 26. Software requirements

Type of software	Minimum software requirements
Backup-archive client	Sun Solaris 10 (update 1 or higher) 32bit and 64bit kernel mode (SPARC or x86)
Applications	The Java Runtime Environment Java JRE 5 or 6 that you can download from the following Website: http://java.sun.com/ **Note:** The "java" executable needs to be found in the system PATH. To verify that you have the right JRE installed you can use the "java -version" command from a UNIX system window.
Web browser	A Firefox 3.0.14 or higher browser for the Web client and to access online help and documentation

Sun Solaris client communication methods

The TCP/IP and shared memory communication methods are available for the Sun Solaris backup-archive client.

You can use the following communication methods with the Tivoli Storage Manager Version 6.2 Sun Solaris client:

Table 27. Sun Solaris client communication methods

To use this communication method:	Install this software:	To connect to these Tivoli Storage Manager servers:
TCP/IP	TCP/IP (Standard with Solaris)	AIX, HP-UX, Linux, Solaris, Windows
Shared Memory	TCP/IP (Standard with Solaris)	Solaris

NDMP support requirements (Extended Edition only)

You can use the Network Data Management Protocol (NDMP) to back up and restore network attached storage (NAS) file systems to tape drives or libraries that are locally attached to Network Appliance and EMC Celerra NAS file servers.

NDMP support is available only on IBM Tivoli Storage Manager Extended Edition.

NDMP support requires the following hardware and software:
- Tivoli Storage Manager Extended Edition
- Tape drive and tape library. For supported combinations, go to: http://www.ibm.com/software/tivoli/products/storage-mgr/product-links.html

Installation requirements for backing up and archiving Tivoli Storage Manager FastBack client data

Before you can back up or archive your FastBack client data, you must install the required software.

You must install the following software:
- Tivoli Storage Manager FastBack Version 6.1
- Tivoli Storage Manager client Version 6.1.3.x (where x is 1 or higher) or Version 6.2 or later
- Tivoli Storage Manager server Version 6.1.3 or higher
- Tivoli Storage Manager Administration Center Version 6.1.3
  - Required only if you want to use integrated Tivoli Storage Manager FastBack - administration.

The Tivoli Storage Manager FastBack environment must be running. For information about installing and setting up Tivoli Storage Manager FastBack, see *Tivoli Storage Manager FastBack Installation and User's Guide*, at: http://publib.boulder.ibm.com/infocenter/tsmfbinf/v6

For information about integrating Tivoli Storage Manager and Tivoli Storage Manager FastBack, see *IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack*.

You can install Tivoli Storage Manager client in one of the following ways:
- Install Tivoli Storage Manager client on a workstation where the FastBack DR Hub is installed. In this case, the prerequisites are: the FastBack Disaster Recovery Hub setup, and the FastBack shell.
- Install Tivoli Storage Manager client on a workstation where neither the FastBack server or the FastBack Disaster Recovery Hub is installed. In this case, the FastBack shell is still required.

Related concepts

"Configuring the client to back up and archive Tivoli Storage Manager FastBack data" on page 80

Installing from the Tivoli Storage Manager DVDs

The locations and instructions to install Tivoli Storage Manager from the DVDs are listed in this topic.

IBM Tivoli Storage Manager AIX Clients
Contains the AIX client images. The AIX client images are in the `/usr/sys/inst.images` directory

IBM Tivoli Storage Manager HP Clients
Contains the HP-UX client images. The images reside in the `tsmcli/'platform'/` directory structure, where `'platform'` is **hp11ia64**.

IBM Tivoli Storage Manager Linux Clients
Contains the Linux x86/x86_64, Linux on System z, and LinuxPPC client images. The images reside in the `tsmcli/'platform'/` directory structure, where `'platform'` is one of the following platform designations: **Linux x86/x86_64, Linux on System z, LinuxPPC (for Linux on POWER client)**.

IBM Tivoli Storage Manager Mac OS X Client
Contains the Mac OS X client images. The images reside in the `tsmcli/mac` directory.

IBM Tivoli Storage Manager Solaris Clients (SPARC and x86)
Contains the Solaris client images. The images reside in the `tsmcli/'platform'/` directory structure, where `'platform'` is one of the following platform designations: **solaris, solaris_ x86**.

See the following Web site for the program directories with information about available UNIX and Linux client installation media: http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager

Install the UNIX and Linux backup-archive clients

This section provides instructions to install and set up Tivoli Storage Manager UNIX and Linux clients.

Note: A root user *must* install Tivoli Storage Manager on a UNIX or Linux workstation.

The supported UNIX and Linux clients and the location of the installation instructions for each client are listed here.
- "Installing the Tivoli Storage Manager AIX client"
- "Installing the Tivoli Storage Manager HP-UX Itanium 2 client" on page 21
- "Installing the Tivoli Storage Manager Linux on POWER client" on page 24
- "Installing the Tivoli Storage Manager Linux x86/x86_64 client" on page 27
- "Installing the Tivoli Storage Manager Linux on System z client" on page 31
- "Installing the Tivoli Storage Manager Mac OS X client" on page 34
- "Installing the Tivoli Storage Manager Solaris client" on page 39

Related concepts
Chapter 2, "Configure the Tivoli Storage Manager client," on page 47

Installing the Tivoli Storage Manager AIX client

You can install the Tivoli Storage Manager AIX client from the product DVD, or other installation media.

You can install the Tivoli Storage Manager AIX client on both 32-bit clients and 64-bit clients. However, only one client can be installed at a time. If you install both the 32-bit and 64-bit clients at the same time, you receive an error message. If you are moving from a 32-bit client to a 64-bit client, first remove the 32-bit client using SMIT (tivoli.tsm.client.ba). Then install the 64-bit client (tivoli.tsm.client.ba.64bit).

If you use the 64-bit client, you must also install the 64-bit version of the API (tivoli.tsm.client.api.64bit). The 32-bit API (tivoli.tsm.client.api.32bit) is installed at the same time, and can coexist with the 64-bit API.

If the 64-bit client is installed but you want to use the 32-bit client instead, you must remove the 64-bit client. Use SMIT to remove the 64-bit client (tivoli.tsm.client.ba.64bit), and then install the 32-bit client. You do not need to uninstall the 64-bit API.

If you are upgrading from a previous version of the Tivoli Storage Manager AIX client, use SMIT to remove any national language filesets that might be installed, and install the national language filesets provided with Tivoli Storage Manager 6.2.

You cannot perform an upgrade install from the Tivoli Storage Manager Version 5.5 or earlier client to the Tivoli Storage Manager Version 6.2 64-bit client. You must remove the older 32-bit client before installing the Tivoli Storage Manager Version 6.2 64-bit client.

All of the needed packages are included in the Tivoli Storage Manager AIX client package, and they overwrite any older runtime applications on your system during installation. LibC (C Set ++) runtime is required.

When using the **installp** command to install this client, do not change the default field values for the following two choices:
- "AUTOMATICALLY install requisite software?"
- "OVERWRITE same or newer versions?"

Disabling or changing the values allow a lower level Tivoli Storage Manager component to install over a currently higher installed component. Under such circumstances, function calls between components at different levels might not be valid any longer.

You can install only *one* of these clients:
- tivoli.tsm.client.ba (32-bit client, includes the HSM JFS2 client)
- tivoli.tsm.client.ba.gpfs (32-bit client, required when using the GPFS HSM client)
- tivoli.tsm.client.ba64 (64-bit client; no HSM)

You can install these APIs:
- tivoli.tsm.client.api.32bit
- tivoli.tsm.client.api.64bit

Note:
1. All of these packages are available on the installation media.
2. HSM installation information is in the *Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide*. You need an HSM license to install and use those file sets.
3. You need an Extended Edition license to use the NAS client.

AIX client

  tivoli.tsm.client.ba
     Installs the following Tivoli Storage Manager 32-bit files:
- Backup-archive base files
- Backup-archive common files
- Backup-archive Web client
- Image backup client
- NAS backup client

  tivoli.tsm.client.ba64
     Installs the following Tivoli Storage Manager 64-bit files:
- Backup-archive base files
- Backup-archive common files

- Backup-archive Web client
- Image backup client

tivoli.tsm.client.ba.gpfs
Installs the following Tivoli Storage Manager 32-bit files:
- Backup-archive base files
- Backup-archive common files
- Backup-archive Web client
- Image backup client

tivoli.tsm.client.hsm.gpfs
Installs the Tivoli Storage Manager Hierarchical Storage Management GPFS client.

tivoli.tsm.client.hsm
Installs the Tivoli Storage Manager Hierarchical Storage Management JFS2 client.

tivoli.tsm.client.jbb
Installs the Tivoli Storage Manager journal-based backup component.

tivoli.tsm.client.api.32bit
Installs the 32-bit API.

tivoli.tsm.client.api.64bit
Installs the 64-bit API.

tivoli.tsm.client.msg._lang_
Installs the NL messages for API, where _lang_ is the appropriate language code from Table 28 on page 19.

tivoli.tsm.filepath_aix
Installs the file path kernel extension required for journal-based backup.

GSKit8.gskcrypt32.ppc.rte and GSKit8.gskssl32.ppc.rte
SSL Toolkit 32-bit (required by the backup-archive client and its components).

GSKit8.gskcrypt64.ppc.rte and GSKit8.gskssl64.ppc.rte
SSL Toolkit 64-bit (required by the 64-bit client API).

Default installation directories

Here are the default directories where some files are stored as part of the client installation:
- The Tivoli Storage Manager backup-archive, Web client, and administrative client (dsmadmc) files are installed in the `/usr/tivoli/tsm/client/ba/bin` directory.
- The Tivoli Storage Manager backup-archive, Web client, and administrative client (dsmadmc) 64-bit files are installed in the `/usr/tivoli/tsm/client/ba/bin64` directory.
- The Tivoli Storage Manager API files are installed in the `/usr/tivoli/tsm/client/api/bin` directory.
- The Tivoli Storage Manager API 64 files are installed in the `/usr/tivoli/tsm/client/api/bin64` directory.
- The sample system-options file, `dsm.sys.smp`, is placed in the installation directory.

Table 28. AIX client: Language codes for installation packages

Language	Tivoli Storage Manager Language code
Simplified Chinese	ZH_CN
Traditional Chinese	ZH_TW
Czech	CS_CZ
French	FR_FR
German	DE_DE
Hungarian	HU_HU
Italian	IT_IT
Japanese	JA_JP
Korean	KO_KR
Polish	PL_PL
Brazilian Portuguese	PT_BR
Russian	RU_RU
Spanish	ES_ES

This installation procedure is designed to install directly from the DVD using a local or remote-mounted DVD drive.

If you are copying the client files into a local directory first, a .toc file is automatically created by the **installp** command. You can create a .toc file manually by running `/usr/sbin/inutoc` in the local directory to which you copied the Tivoli Storage Manager image. From the AIX command line, enter:

```
/usr/sbin/inutoc /usr/sys/inst.images
```

A **.toc** file is created in that directory.

To install Tivoli Storage Manager from the DVD:
1. Log in as the root user, insert the DVD into the DVD drive device. You must manually mount the DVD device for remote or NFS installations.
2. From the AIX command line, type **smitty install** and press Enter.
3. Select **Install and Update Software** and press Enter.
4. Select **Install and Update From ALL Available Software** and press Enter.
5. At the **INPUT device/directory for software** prompt, press the F4 key and select the DVD device containing the installation DVD or specify the directory containing the installation images, and press Enter.
6. At the **SOFTWARE to install** prompt, press the F4 key. Select the Tivoli Storage Manager file sets you want to install by pressing the F7 key. Then press Enter.
7. On the **Install and Update From ALL Available Software** panel, press the F4 key to change any entry fields, or use the default fields. Press Enter twice to begin the installation.
8. After the installation completes, select F10 to exit.

Note:
1. When installing the clients in a non-English locale, the corresponding local message file sets are required in addition to the clients from the **Install and**

Update From ALL Available Software panel. For example, when installing the Tivoli Storage Manager client in ES_ES locale, the following message fileset is also required:

```
TSM Client Messages - Spanish          msg.ES_ES
```

2. When file sets are installed, the file sets are automatically committed on the system. The previous version of Tivoli Storage Manager software is completely replaced by the newly installed version.
3. The 32-bit Application Programming Interface (tivoli.tsm.client.api.32bit) is an install prerequisite to the 64-bit Application Programming Interface (tivoli.tsm.client.api.64bit).

The Tivoli Storage Manager files are installed in the `/usr/tivoli/tsm/client/ba/ bin` directory. If you move the Tivoli Storage Manager files to another directory, you must perform the following steps:
1. Make sure the permissions of the installed files have not changed.
2. Update the symbolic links for the installed files in the following directories:
   - The `/usr/bin` directory
   - The `/usr/lib` directory for Tivoli Storage Manager libraries
   - The directory that contains symbolic links for each language package you install (for example, `/usr/lib/nls/msg/en_US`)
3. Ensure that every user of Tivoli Storage Manager sets the DSM_DIR environment variable to the newly installed directory.

After the installation completes, see Chapter 2, "Configure the Tivoli Storage Manager client," on page 47 for required and optional tasks to perform before using Tivoli Storage Manager.

Note:
- For AIX workload partitions (WPAR), installation in a global WPAR *is* supported, and installation in local WPARs is *not* supported.
- On AIX Version 6.1, if you are using encrypted file systems (EFS) with Tivoli Storage Manager, and if the EFS user keystore password is different from the user login password, the EFS keystore is not automatically opened when you log on. If the EFS keystore is not open when you log on, the Tivoli Storage Manager client might not restore a non-EFS file into an EFS file system. You can prevent the EFS file system restore problem one of the following ways:
  - Launch the Tivoli Storage Manager client using the **efskeymgr -o** command. For example:

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

  - Synchronize the keystore password with the user login password using the **efskeymgr -n** command. For example:

    ```
    efskeymgr -n
    ```

Uninstalling the Tivoli Storage Manager AIX client

You can use the following procedures to uninstall the Tivoli Storage Manager AIX client.
1. From the AIX command line, enter the following command: `smitty remove`.
2. Hit the ENTER key.
3. In the SOFTWARE name field, press F4 to list the Tivoli Storage Manager file sets that you want to uninstall; hit the ENTER key.
4. Select the Tivoli Storage Manager file sets that you want to uninstall; press the ENTER key.

5. In the PREVIEW only? field (remove operation will NOT occur), select No; hit the ENTER key.

Note: Tivoli Storage Manager client modules and components are tightly integrated and installed file sets are automatically committed. There is no option for rollbacks.

Installing the Tivoli Storage Manager HP-UX Itanium 2 client

You can install the Tivoli Storage Manager HP-UX Itanium 2 client from the product DVD, or other installation media.

The following source packages are available on the installation media:

`tsmcli/hp11ia64/gskcrypt64-8.x.x.x.hpux.ia64.tar.Z` and `tsmcli/hp11ia64/gskssl64-8.x.x.x.hpux.ia64.tar.Z`
Contains the GSKit. If you have a previous version of the GSKit, uninstall it before installing the new version.

`tsmcli/hp11ia64/TIVsmC`
In this package the software selection name used by swlist for the top level product name is `TIVsm64`. The components under `TIVsm64` are `TIVsm.CLIENT` and `TIVsm.CLIENT_API64`.

`TIVsm.CLIENT`
Contains the backup-archive client (command-line and Java GUI), administrative client (command-line), and the Web client with the English message catalogs.

`TIVsm.CLIENT_API64`
Contains the 64-bit API with the English message catalogs.

Note: Additional language support is available under the top level product name of **tsmcli/hp11/TIVsmC.msg.lang**. In this package, the additional language support is available for all components under the top level product name `TIVsm`. The component name is `TIVsm.CLIENT_msg_lang`. Replace `lang` with the appropriate language code from Table 29 on page 22.

`tsmcli/hp11ia64/TIVsmCapi64`
In this package the software selection name used by swlist for the top level product name is `TIVsm64`. The component under `TIVsm64` is `TIVsm.CLIENT_API64`.

Default installation directories

Here are the default directories where some files are stored as part of the client installation:
- The Tivoli Storage Manager backup-archive, Web client, and administrative client (`dsmadmc`) files are installed in the `/opt/tivoli/tsm/client/ba/bin` directory.
- The Tivoli Storage Manager API 64 files are installed in the `/opt/tivoli/tsm/client/api/bin64` directory.
- The sample system-options file, `dsm.sys.smp`, is placed in the installation directory.
- If DSM_DIR is not set, the executable file dsmc, the resource files, and the dsm.sys file are stored in the default installation directory.

- If DSM_CONFIG is not set, the client user-options file must reside in the default installation directory. If DSM_CONFIG is set, then the file must reside in the directory specified by DSM_DIR.
- If you do not define DSM_LOG, Tivoli Storage Manager writes messages to the dsmerror.log and dsmsched.log files in the current working directory.

Table 29. HP-UX Itanium 2 client: Language codes for installation packages

Language	Tivoli Storage Manager Language code
Simplified Chinese	ZH_CN
Traditional Chinese	ZH_TW
Czech	CS_CZ
French	FR_FR
German	DE_DE
Hungarian	HU_HU
Italian	IT_IT
Japanese	JA_JP
Korean	KO_KR
Polish	PL_PL
Brazilian Portuguese	PT_BR
Russian	RU_RU
Spanish	ES_ES

To remove previous Tivoli Storage Manager client versions, log in as the root user and enter the following command:

```
/usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm64
```

If you installed additional languages, execute the following command to remove them:

```
/usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm64.CLIENT_msg_lang
```

Replace *lang* with the appropriate language code from Table 29.

To install from the DVD follow these steps:
1. Log in as the root user
2. Mount the DVD to `/dvd`
3. Change the directory to `tsmcli/hp11ia64`.
4. To install GSKit: If you have a previous version of GSKit installed, remove it before installing the new version. Extract the contents of these files: `gskcrypt64-8.x.x.x.hpux.ia64.tar.Z` and `gskssl64-8.x.x.x.hpux.ia64.tar.Z` to a directory on your hard drive. Enter the following commands to install the packages: `/usr/sbin/swinstall -x mount_all_filesystems=false -v -s ` `pwd` `/gskcrypt64 gskcrypt64` and `/usr/sbin/swinstall -x mount_all_filesystems=false -v -s ` `pwd` `/gskssl64 gskssl64`
5. If you downloaded from ftp, go to the directory where the installable image is located. Enter the following command: `/usr/sbin/swinstall -x mount_all_filesystems=false -v -s ` `pwd` `/TIVsmC TIVsm64`

Note: `pwd` can be used instead of the absolute name of the current directory.

To install only the API, omit the last `TIVsm64` from the command above, and mark only the fileset for installation in the swinstall user interface you want to install:

`CLIENT` for the backup-archive client and API, or
`CLIENT_API64` for the API only

Note: The Client needs the API for Raw Logical Volume backup. Therefore if you mark **CLIENT** for installation the API is also installed.

To install additional languages, execute the following command:

```
/usr/sbin/swinstall -x mount_all_filesystems=false -v -s `pwd`
/TIVsmC.msg.lang TIVsm64.CLIENT_msg_lang
```

Note:

1. The *pwd* parameter can be used instead of the absolute name of the current directory.
2. Replace *lang* with the appropriate language code from Table 29 on page 22.

Related concepts

Chapter 2, "Configure the Tivoli Storage Manager client," on page 47

Increasing the default limit of the data segment size

The default limit of the data segment size of a process in HP-UX 11i v2 is 64 MB. When backing up large file systems, the Tivoli Storage Manager client might exceed this limit and run out of memory.

To increase this limit you can modify the kernel as follows:

1. As root user, start **sam**.
2. Select **Kernel Configuration**.
3. Select **Configurable Parameters**.
4. Locate **maxdsize** and increase its value through the menu entry **Actions/Modify Configurable Parameter...** (e.g. set maxdsize to 268435456 for a 256 MB max size of the data segment.
5. The kernel is rebuilt by **sam** after this change. You must reboot for the new setting to take effect.

Uninstalling the Tivoli Storage Manager HP-UX Itanium 2 client

You can use the following procedures to uninstall the Tivoli Storage Manager HP-UX Itanium 2 client.

Follow these steps to uninstall all of the packages related to Tivoli Storage Manager, including the command line, GUI, Web GUI, and administrative Client components.

Important: Make sure that you uninstall the packages in the given order.

1. Check the *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide* for instructions on removing HSM and installing a new version.
2. Enter the following command to uninstall the Tivoli Storage Manager Backup-Archive Client: `/usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm64.CLIENT`. This uninstalls all of the components of the Backup-Archive Client (command-line, Web GUI and the administrative client). You cannot uninstall a single component of this package (for example, the Web GUI) without uninstalling the complete package.
3. If one or more Tivoli Storage Manager language messages packages have been installed, you must remove these before removing the API package. To remove

the language messages packages, enter the following command as root user: /usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm64.CLIENT_msg_xx_XX, where xx_XX is the language you want to remove. The current supported languages are: cs_CZ, de_DE, es_ES, it_IT, fr_FR, hu_HU, ja_JP, ko_KR, pl_PL, pt_BR, ru_RU, zh_CN and zh_TW

4. To remove the CLIENT_API file set, enter the following command: /usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm64.CLIENT_API64

Installing the Tivoli Storage Manager Linux on POWER client

You can install the IBM Tivoli Storage Manager Linux on POWER client from the product DVD, or other installation media.

If a previous version of Tivoli Storage Manager is installed, remove it and other dependent software programs before installing a new version.

The following installation options are available in uncompressed packages on the DVD.

Table 30. Package Names and Directories

This software package	Is installed in this directory
TIVsm-API.ppc64.rpm	/opt/tivoli/tsm/tsmcli/api/bin
TIVsm-BA.ppc64.rpm	/opt/tivoli/tsm/tsmcli/linuxppc
gskcrypt32-8.x.x.x.linux.ppc.rpm (32-bit) gskssl32-8.x.x.x.linux.ppc.rpm (32-bit)	/usr/local/ibm/gsk8
gskcrypt64-8.x.x.x.linux.ppc.rpm (64-bit) gskssl64-8.x.x.x.linux.ppc.rpm (64-bit)	/usr/local/ibm/gsk8_64
TIVsm-msg.<XX_XX>.ppc64.rpm	/opt/tivoli/tsm/tsmcli/lang/<XX_XX> Where **<XX_XX>** defines the language installed.

Use the following procedure to install the Tivoli Storage Manager clients:
1. Log on using the root user ID and mount the DVD to /dvd.
2. Enter the following directory path where the installation packages reside on the DVD: /dvd/tsmcli/linuxppc
3. Enter the following commands to install the Global Security Toolkit (GSKit) 32-bit packages:

   rpm -U gskcrypt32-8.x.x.x.linux.ppc.rpm

   gskssl32-8.x.x.x.linux.ppc.rpm

   Option -U installs the package or upgrades, if an older version is already installed.
4. Enter the following commands to install the GSKit 64-bit packages:

   rpm -U gskcrypt64-8.x.x.x.linux.ppc.rpm

   gskssl64-8.x.x.x.linux.ppc.rpm

   Option -U installs the package or upgrades, if an older version is already installed.

5. Enter the following commands to install the backup-archive client (command-line and 64-bit API), the administrative client (command-line), and the Web client.

   a. `rpm -i TIVsm-API.ppc64.rpm` (required only for backup-archive client language packages and for a 32-bit independent software vendor) or `rpm -i TIVsm-API64.ppc64.rpm` (required for TIVsm-BA.ppc64.rpm).

   b. `rpm -i TIVsm-BA.ppc64.rpm`.

6. To circumvent the dependence check, use the --nodeps option. Ensure that you check the dependencies manually using the following commands:

   a. `rpm -i --nodeps TIVsm-API.ppc64.rpm` (required only for backup-archive client language packages and for a 32-bit independent software vendor) or `rpm -i --nodeps TIVsm-API64.ppc64.rpm` (required for TIVsm-BA.ppc64.rpm).

   b. `rpm -i --nodeps TIVsm-BA.ppc64.rpm`.

   The backup-archive client requires the API package to perform image backups.

7. To install additional language support (the default is American English), enter the following command:

   a. `rpm -i TIVsm-msg.`*lang*`.ppc64.rpm`. Replace *lang* with the appropriate language code from Table 31.

To install this language	Enter this command
NLS French messages	**rpm -i TIVsm-msg.FR_FR.ppc64.rpm**
NLS German messages	**rpm -i TIVsm-msg.DE_DE.ppc64.rpm**
NLS Italian messages	**rpm -i TIVsm-msg.IT_IT.ppc64.rpm**
NLS Spanish messages	**rpm -i TIVsm-msg.ES_ES.ppc646.rpm**
NLS Portuguese messages	**rpm -i TIVsm-msg.PT_BR.ppc64.rpm**
NLS Czech messages	**rpm -i TIVsm-msg.CS_CZ.ppc64.rpm**
NLS Hungarian messages	**rpm -i TIVsm-msg.HU_HU.ppc64.rpm**
NLS Polish messages	**rpm -i TIVsm-msg.PL_PL.ppc64.rpm**
NLS Russian messages	**rpm -i TIVsm-msg.RU_RU.ppc64.rpm**
NLS Japanese messages	**rpm -i TIVsm-msg.JA_JP.ppc64.rpm**
NLS Korean messages	**rpm -i TIVsm-msg.KO_KR.ppc64.rpm**
NLS Traditional Chinese messages (EUC)	**rpm -i TIVsm-msg.ZH_CN.ppc64.rpm**
NLS Traditional Chinese Big5 messages	**rpm -i TIVsm-msg.ZH_TW.ppc64.rpm**

Table 31. Linux on POWER client: Language codes for installation packages

Language	Tivoli Storage Manager Language code
Simplified Chinese	ZH_CN
Traditional Chinese	ZH_TW
Czech	CS_CZ
French	FR_FR
German	DE_DE
Hungarian	HU_HU
Italian	IT_IT
Japanese	JA_JP
Korean	KO_KR

Table 31. Linux on POWER client: Language codes for installation packages (continued)

Language	Tivoli Storage Manager Language code
Polish	PL_PL
Brazilian Portuguese	PT_BR
Russian	RU_RU
Spanish	ES_ES

Default installation directories

Here are the default directories where some files are stored as part of the client installation:

- The Tivoli Storage Manager backup-archive, Web client, Java GUI, and administrative client (dsmadmc) files are installed in the `/opt/tivoli/tsm/client/ba/bin` directory.
- The Tivoli Storage Manager API files are installed in the `/opt/tivoli/tsm/client/api/bin` directory.
- The Tivoli Storage Manager documentation files are installed in the `/opt/tivoli/tsm/client/books/html/tsmunix` and `/opt/tivoli/tsm/client/books/pdf/tsmunix` directories.
- The sample system-options file, `dsm.sys.smp`, is placed in the installation directory.
- If DSM_DIR is not set, the executable file dsmc, the resource files, and the `dsm.sys` file are stored in the default installation directory.
- If DSM_CONFIG is not set, the client user-options file must reside in the default installation directory. If DSM_CONFIG is set, then the file must reside in the directory specified by DSM_DIR.
- If you do not define DSM_LOG, Tivoli Storage Manager writes messages to the `dsmerror.log` and `dsmsched.log` files in the current working directory.

Related concepts

Chapter 2, "Configure the Tivoli Storage Manager client," on page 47

Uninstalling the Tivoli Storage Manager Linux on POWER client

You can use the following procedures to uninstall Tivoli Storage Manager Linux on POWER client.

1. To uninstall previously installed Tivoli Storage Manager client packages, log in as root and enter:

   ```
   rpm -e TIVsm-BA
   ```

   .

2. To uninstall the API, enter this command:

   ```
   rpm -e TIVsm-API
   ```

   If other API-dependent programs (such as Tivoli Storage Manager Data Protection products) are installed, then you must either uninstall the API-dependent program, or remove the API package.

   a. To remove the 64-bit API package, enter the following command as root user:

      ```
      rpm -e TIVsm-API64 --nodeps
      ```

b. To remove the 32-bit API package, enter the following command as root
      user:

      ```
      rpm -e TIVsm-API --nodeps
      ```

   If you uninstalled the API-dependent product, you must reinstall it after
   installing a newer version of the backup-archive client and API. Consult the
   documentation of the dependent product for any considerations when
   uninstalling and reinstalling that product to prevent any configuration loss.

3. If an additional language package is installed, uninstall it by entering these
   commands, replacing XX_XX with the language code from Table 31 on page 25.

   a. `rpm -e TIVsm-BA.msg.XX_XX`

   b. `rpm -e TIVsm-API.msg.XX_XX`

Tip: The version number of the packages is not needed for uninstall.

Installing the Tivoli Storage Manager Linux x86/x86_64 client

You can install the IBM Tivoli Storage Manager Linux x86/x86_64 client from the
product DVD, or other installation media.

If a previous version of Tivoli Storage Manager is installed, remove it and other
dependent software programs before installing a new version.

This installation procedure is designed to install directly from the Tivoli Storage
Manager UNIX client DVD from a local or remote-mounted DVD drive.

Table 32. Package names and directories

This software package	Is installed in this directory
`TIVsm-API.i386.rpm`	`/opt/tivoli/tsm/client/api/bin`
`TIVsm-API64.i386.rpm`	`/opt/tivoli/tsm/client/api/bin64`
`TIVsm-BA.i386.rpm`	`/opt/tivoli/tsm/client/ba`
`gskcrypt32-8.x.x.x.linux.x86.rpm` (32-bit) `gskssl32-8.x.x.x.linux.x86.rpm` (32-bit)	`/usr/local/ibm/gsk8`
`gskcrypt64-8.x.x.x.linux.x86_64.rpm` (64-bit) `gskssl64-8.x.x.x.linux.x86_64.rpm` (64-bit)	`/usr/local/ibm/gsk8_64`
`TIVsm-msg.<XX_XX>.i386.rpm`	`/opt/tivoli/tsm/client/lang/<XX_XX>` where **<XX_XX>** defines the language that is installed
`TIVsm-BAhdw.i386.rpm`	`/opt/tivoli/tsm/client/ba/bin/plugins`
`libzephyr.so`	`/opt/tivoli/tsm/client/ba/bin`
`TIVsm-BAhdw.i386.rpm`	`/opt/tivoli/tsm/client/ba/plugins`

Follow these steps to install the Tivoli Storage Manager backup-archive client:

1. Log on as the root user and mount the Tivoli Storage Manager UNIX client
   DVD to /dvd.

2. Change to the directory where the packages are stored:

   ```
    /dvd/tsmcli/linux86
   ```

If the Tivoli Storage Manager UNIX client DVD is not mounted to /dvd, or if the packages are stored in a different directory (e.g. downloaded by ftp), change to the correct directory.

Make sure that you install these packages in the given order. This is required because the Tivoli Storage Manager API package is a prerequisite of the backup-archive client package.

The Web client is an integral part of the backup-archive client package and cannot be installed without it.

3. Enter the following commands to install the Global Security Toolkit (GSKit) 32-bit packages:

   ```
   rpm -U gskcrypt32-8.x.x.x.linux.x86.rpm gskssl32-8.x.x.x.linux.x86.rpm
   ```

   Option -U installs the package or upgrades, if an older version is already installed.

4. Enter the following commands to install the Tivoli Storage Manager 32-bit API and 32-bit backup-archive client. This installs the command-line client and the administrative client:

   ```
   rpm -i TIVsm-API.i386.rpm
   ```

   ```
   rpm -i TIVsm-BA.i386.rpm
   ```

5. Install the snapshot difference incremental backup support for NetApp and N-Series file servers by entering the following command:

   ```
   rpm -i TIVsm-BAhdw.i386.rpm
   ```

Follow these steps to install the Tivoli Storage Manager API only:

1. Log on as the root user and mount the Tivoli Storage Manager UNIX client DVD to /dvd.

2. Change to the directory where the packages are stored:

   ```
    /dvd/tsmcli/linux86
   ```

   If the Tivoli Storage Manager UNIX client DVD is not mounted to /dvd, or if the packages are stored in a different directory (e.g. downloaded by ftp), change to the correct directory.

   Make sure that you install these packages in the given order. This is required because the 32-bit GSKit and 32-bit API are prerequisites of the 64-bit GSKit and 64-bit API.

3. Enter the following commands to install the Global Security Toolkit (GSKit) 32-bit packages:

   ```
   rpm -U gskcrypt32-8.x.x.x.linux.x86.rpm gskssl32-8.x.x.x.linux.x86.rpm
   ```

   Option -U installs the package or upgrades, if an older version is already installed.

4. Optional: Do this to install the 64-bit API.

   Enter the following commands to install the GSKit 64-bit packages:

   ```
   rpm -U gskcrypt64-8.x.x.x.linux.x86_64.rpm gskssl64-8.x.x.x.linux.x86_64.rpm
   ```

   Option -U installs the package or upgrades, if an older version is already installed.

5. Enter the following commands to install the Tivoli Storage Manager 32-bit API.

   ```
   rpm -i TIVsm-API.i386.rpm
   ```

6. Optional: Do this to install the 64-bit API.

   Enter the following commands to install the Tivoli Storage Manager 64-bit API.

   ```
   rpm -i TIVsm-API64.i386.rpm
   ```

The Tivoli Storage Manager default language is English. Install additional languages (optional) by entering the commands listed for each additional language.

To install this language	Enter this command
NLS French messages	**rpm -i TIVsm-msg.FR_FR.i386.rpm**
NLS German messages	**rpm -i TIVsm-msg.DE_DE.i386.rpm**
NLS Italian messages	**rpm -i TIVsm-msg.IT_IT.i386.rpm**
NLS Spanish messages	**rpm -i TIVsm-msg.ES_ES.i386.rpm**
NLS Portuguese messages	**rpm -i TIVsm-msg.PT_BR.i386.rpm**
NLS Czech messages	**rpm -i TIVsm-msg.CS_CZ.i386.rpm**
NLS Hungarian messages	**rpm -i TIVsm-msg.HU_HU.i386.rpm**
NLS Polish messages	**rpm -i TIVsm-msg.PL_PL.i386.rpm**
NLS Russian messages	**rpm -i TIVsm-msg.RU_RU.i386.rpm**
NLS Japanese messages	**rpm -i TIVsm-msg.JA_JP.i386.rpm**
NLS Korean messages	**rpm -i TIVsm-msg.KO_KR.i386.rpm**
NLS Traditional Chinese messages (EUC)	**rpm -i TIVsm-msg.ZH_CN.i386.rpm**
NLS Traditional Chinese Big5 messages	**rpm -i TIVsm-msg.ZH_TW.i386.rpm**

Table 33. Linux x86/x86_64 client: Language codes for installation packages

Language	Tivoli Storage Manager Language code
Simplified Chinese	ZH_CN
Traditional Chinese	ZH_TW
Czech	CS_CZ
French	FR_FR
German	DE_DE
Hungarian	HU_HU
Italian	IT_IT
Japanese	JA_JP
Korean	KO_KR
Polish	PL_PL
Brazilian Portuguese	PT_BR
Russian	RU_RU
Spanish	ES_ES

Default installation directories

Here are the default directories where some files are stored as part of the client installation:

- The sample system-options file, `dsm.sys.smp`, is placed in the installation directory.
- If DSM_DIR is not set, the executable file dsmc, the resource files, and the dsm.sys file are stored in the default installation directory.
- If DSM_CONFIG is not set, the client user-options file must reside in the default installation directory. If DSM_CONFIG is set, then the file must reside in the directory specified by DSM_DIR.

- If you do not define DSM_LOG, Tivoli Storage Manager writes messages to the dsmerror.log and dsmsched.log files in the current working directory.

Related concepts

Chapter 2, "Configure the Tivoli Storage Manager client," on page 47

Uninstalling the Tivoli Storage Manager Linux x86/x86_64 client

You can use the following procedures to uninstall the Tivoli Storage Manager Linux x86/x86_64 client.

Follow these steps to uninstall all the packages related to Tivoli Storage Manager and HSM, including the command line, GUI, Web GUI and administrative client components.

Important: Make sure that you uninstall the packages in the given order.

1. Check the *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide* for instructions on removing HSM and installing a new version.

2. Enter the following command to uninstall the Tivoli Storage Manager backup-archive client. This uninstalls all of the components of the backup-archive client (command-line, GUI, Web GUI and the administrative client):

   ```
   rpm -e TIVsm-BA
   ```

   You cannot uninstall a single component of this package (for example, the Web GUI) without uninstalling the complete package.

3. If one or more language messages packages are installed, they must be removed before removing the API package. Enter the following command as root user, where XX_XX is the language you want to remove:

   ```
   rpm -e TIVsm-msg.XX_XX
   ```

4. To remove the 64-bit API package enter the following command as root user:

   ```
   rpm -e TIVsm-API64
   ```

   If other API-dependent programs (such as Tivoli Storage Manager Data Protection products) are installed, then you must do *one* of the following steps:

   a. Uninstall the API-dependent program

   b. Remove the 64-bit API package using the following command as root user:

      ```
      rpm -e TIVsm-API64 --nodeps
      ```

5. To remove the 32-bit API package enter the following command as root user:

   ```
   rpm -e TIVsm-API
   ```

   If other API-dependent programs (such as Tivoli Storage Manager Data Protection products) are installed, then you must do *one* of the following steps:

   a. Uninstall the API-dependent program

   b. Remove the 32-bit API package using the following command as root user:

      ```
      rpm -e TIVsm-API --nodeps
      ```

If you uninstalled the API-dependent product, you must reinstall it after installing a newer version of the backup-archive client and API. Consult the documentation of the dependent product for any considerations when uninstalling and reinstalling that product to prevent any configuration loss.

Related information
Table 33 on page 29

Installing the Tivoli Storage Manager Linux on System z client

You can install the IBM Tivoli Storage Manager Linux on System z client from the product DVD, or other installation media.

If a previous version of Tivoli Storage Manager is installed, remove it and other dependent software programs before installing a new version.

This installation procedure is designed to install directly from the Tivoli Storage Manager UNIX client DVD from a local or remote-mounted DVD drive.

Table 34. Package Names and Directories

This software package	Is installed in this directory
`TIVsm-API.s390.rpm`	`/opt/tivoli/tsm/client/api/bin`
`TIVsm-API64.s390.rpm`	`/opt/tivoli/tsm/client/api/bin64`
`TIVsm-BA.s390.rpm`	`/opt/tivoli/tsm/client/ba`
`gskcrypt31-8.x.x.x.linux.s390.rpm` (32-bit) `gskssl31-8.x.x.x.linux.s390.rpm` (32-bit)	`/usr/local/ibm/gsk8`
`gskcrypt64-8.x.x.x.linux.s390x.rpm` (64-bit) `gskssl64-8.x.x.x.linux.s390x.rpm` (64-bit)	`/usr/local/ibm/gsk8_64`
`TIVsm-msg.<XX_XX>.s390.rpm`	`/opt/tivoli/tsm/client/lang/<XX_XX>` Where **<XX_XX>** defines the language installed

Follow these steps to install Tivoli Storage Manager, including the command-line, the Web GUI and the administrative client.
1. Login as the root user and mount the Tivoli Storage Manager UNIX client DVD to `/dvd`.
2. Change to the directory where the packages are stored: `/dvd/tsmcli/linux390`. If the Tivoli Storage Manager UNIX client DVD is not mounted to `/dvd`, or if the packages are stored in a different directory (e.g. downloaded by ftp), change to the correct directory.
3. Enter the following commands to install the Global Security Toolkit (GSKit) 32-bit packages: (1) `rpm -U gskcrypt31-8.x.x.x.linux.s390.rpm` , (2) `gskssl31-8.x.x.x.linux.s390.rpm`. Option -U installs the package or upgrades, if an older version is already installed.
4. Enter the following commands to install the GSKit 64-bit packages: (1) `rpm -U gskcrypt64-8.x.x.x.linux.s390x.rpm` , (2) `gskssl64-8.x.x.x.linux.s390x.rpm`. Option -U installs the package or upgrades, if an older version is already installed.
5. To install the 32 Bit API , enter the following command: `rpm -i TIVsm-API.s390.rpm`.

6. To install the 64 Bit API and Backup-Archive client, enter the following commands: (1) `rpm -i TIVsm-API.s390.rpm`, (2) `rpm -i TIVsm-API64.s390.rpm`, and (3) `rpm -i TIVsm-BA.s390.rpm`

- Ensure that you install these packages in the given order. This is required because the Tivoli Storage Manager API package is a prerequisite of the backup-archive client package.
- The Web client is an integral part of the backup-archive client package and cannot be installed without it.

To install this language	Enter this command
NLS French messages	**rpm -i TIVsm-msg.FR_FR.s390.rpm**
NLS German messages	**rpm -i TIVsm-msg.DE_DE.s390.rpm**
NLS Italian messages	**rpm -i TIVsm-msg.IT_IT.s390.rpm**
NLS Spanish messages	**rpm -i TIVsm-msg.ES_ES.s390.rpm**
NLS Portuguese messages	**rpm -i TIVsm-msg.PT_BR.s390.rpm**
NLS Czech messages	**rpm -i TIVsm-msg.CS_CZ.s390.rpm**
NLS Hungarian messages	**rpm -i TIVsm-msg.HU_HU.s390.rpm**
NLS Polish messages	**rpm -i TIVsm-msg.PL_PL.s390.rpm**
NLS Russian messages	**rpm -i TIVsm-msg.RU_RU.s390.rpm**
NLS Japanese messages	**rpm -i TIVsm-msg.JA_JP.s390.rpm**
NLS Korean messages	**rpm -i TIVsm-msg.KO_KR.s390.rpm**
NLS Traditional Chinese messages (EUC)	**rpm -i TIVsm-msg.ZH_CN.s390.rpm**
NLS Traditional Chinese Big5 messages	**rpm -i TIVsm-msg.ZH_TW.s390.rpm**

Table 35. Linux on System z client: Language codes for installation packages

Language	Tivoli Storage Manager Language code
Simplified Chinese	ZH_CN
Traditional Chinese	ZH_TW
Czech	CS_CZ
French	FR_FR
German	DE_DE
Hungarian	HU_HU
Italian	IT_IT
Japanese	JA_JP
Korean	KO_KR
Polish	PL_PL
Brazilian Portuguese	PT_BR
Russian	RU_RU
Spanish	ES_ES

Default installation directories

Here are the default directories where some files are stored as part of the client installation:

- The sample system-options file, `dsm.sys.smp`, is placed in the installation directory.
- If DSM_DIR is not set, the executable file dsmc, the resource files, and the dsm.sys file are stored in the default installation directory.
- If DSM_CONFIG is not set, the client user-options file must reside in the default installation directory. If DSM_CONFIG is set, then the file must reside in the directory specified by DSM_DIR.
- If you do not define DSM_LOG, Tivoli Storage Manager writes messages to the `/var/ibm/tivoli/tsm/client/ba/log` or `/var/ibm/tivoli/tsm/client/api/bin/log` files.

Related concepts

Chapter 2, "Configure the Tivoli Storage Manager client," on page 47

Uninstalling the Tivoli Storage Manager Linux on System z client

You can use the following procedures to uninstall the Tivoli Storage ManagerLinux on System z client.

Follow these steps to uninstall all the packages related to Tivoli Storage Manager and HSM, including the command line, GUI, Web GUI and administrative client components.

Important: Make sure that you uninstall the packages in the given order.

1. Enter the following command to uninstall the Tivoli Storage Manager backup-archive client. This uninstalls all of the components of the backup-archive client (command-line, GUI, Web GUI and the administrative client):

   `rpm -e TIVsm-BA`

   You cannot uninstall a single component of this package (for example, the Web GUI) without uninstalling the complete package.

2. If one or more language messages packages are installed, remove them before removing the API package. Enter the following command using the root user ID, where `XX_XX` is the language you want to remove:

   `rpm -e TIVsm-msg.XX_XX`

3. To remove the 32-bit API package, enter the following command as root user:

   `rpm -e TIVsm-API`

   If other API-dependent programs (such as Tivoli Storage Manager Data Protection products) are installed, then you must do *one* of the following steps:

   a. Uninstall the API-dependent program

   b. Remove the 32-bit API package using the following command as root user:

      `rpm -e TIVsm-API --nodeps`

4. To remove the 64-bit API package, enter the following command as root user:

   `rpm -e TIVsm-API64`

   If other API-dependent programs (such as Tivoli Storage Manager Data Protection products) are installed, then you must do *one* of the following steps:

   a. Uninstall the API-dependent program

   b. Remove the 64-bit API package using the following command as root user:

      `rpm -e TIVsm-API64 --nodeps`

If you uninstalled the API-dependent product, you must reinstall it after installing a newer version of the backup-archive client and API. Consult the documentation of the dependent product for any considerations when uninstalling and reinstalling that product to prevent any configuration loss.

Related information

Table 35 on page 32

Installing the Tivoli Storage Manager Mac OS X client

You can install the Tivoli Storage Manager Mac OS X client from the product DVD, or other installation media.

Requirement: *You must be a system administrator to install Tivoli Storage Manager.*

To install the Mac OS X client from the Tivoli Storage Manager DVD, follow these steps:

1. Go to the `tsmcli/mac` directory, for the Mac OS X client installer.
2. Double-click the `6.2.0.0-TIV-TSMBAC-Mac.dmg` file, to mount the disk image.
3. Double-click the Tivoli Storage Manager installation package icon.
4. Follow the prompts to complete the installation.

Installing the backup-archive client using the command line is an optional method of installation that is useful in situations where the installation needs to be done on multiple systems or on systems with no console attached. The installation is done from the command line, can be scripted, and requires no user interaction.

To install the Tivoli Storage Manager backup-archive client from the command line, follow these steps:

1. Change directories to where the Tivoli Storage Manager installer is located.
2. Install the custom install package using the following command:

   ```
   /usr/sbin/installer -pkg "./Tivoli Storage Manager.mpkg"/Contents/
     Packages/TsmBaDsmadmc.pkg -target /
   ```
3. To do a typical installation, issue this command:

   ```
   /usr/sbin/installer -pkg "./Tivoli Storage Manager.mpkg" -target /
   ```
4. To install the Tivoli Storage Manager API, issue this command:

   ```
   /usr/sbin/installer -pkg "/Volumes/
     Tivoli Storage Manager/Tivoli Storage Manager.mpkg/
     Contents/Packages/TsmApi.pkg" -target /
   ```
5. To install the Tivoli Storage Manager API Samples, issue this command:

   ```
   /usr/sbin/installer -pkg "/Volumes/
     Tivoli Storage Manager/Tivoli Storage Manager.mpkg/
     Contents/Packages/TsmApiSample.pkg" -target /
   ```
6. To install a language pack, issue this command, where <xx_XX> is the language code to install, listed in the following table.

   ```
   /usr/sbin/installer -pkg "/Volumes/
     TSM Language Pack -- <xx_XX>/TsmBa<xx_XX>.pkg" -target
   ```

Table 36. Mac OS X client: Language codes for installation packages

Language	Tivoli Storage Manager Language code
English	EN_US
Simplified Chinese	ZH_CN
Traditional Chinese	ZH_TW

Table 36. Mac OS X client: Language codes for installation packages (continued)

Language	Tivoli Storage Manager Language code
Czech	CS_CZ
French	FR_FR
German	DE_DE
Hungarian	HU_HU
Italian	IT_IT
Japanese	JA_JP
Korean	KO_KR
Polish	PL_PL
Brazilian Portuguese	PT_BR
Russian	RU_RU
Spanish	ES_ES

The packages listed in Table 37 are part of Tivoli Storage Manager and are used by the Mac OS X Install application. Use the information in the table to assist you if you need to do a fully customized command-line installation.

Table 37. Mac OS X installation packages

Package	Package Name	Package Description
Tivoli Storage Manager Backup-Archive Java Bundle	TsmBaJavaBundle.pkg	Contains the following component: • Tivoli Storage Manager Java client GUI **Note:** Requires that the following packages be installed first: – Tivoli Storage Manager Backup-Archive Java GUI Files – Tivoli Storage Manager Backup-Archive Client Common Java Files – Tivoli Storage Manager Backup-Archive Client Common Files – Tivoli Storage Manager Backup-Archive IBM Crypto for C Library – Tivoli Storage Manager Backup-Archive Help and Messages

Table 37. Mac OS X installation packages (continued)

Package	Package Name	Package Description

Tivoli Storage Manager Backup-Archive Web Applet	TsmBaWebApplet.pkg	Contains the following component:
		• Tivoli Storage Manager Web Client
		Note: Requires that the following packages be installed first:
		  – Tivoli Storage Manager Backup-Archive Java GUI Files
		  – Tivoli Storage Manager Backup-Archive Client Common Java Files
		  – Tivoli Storage Manager Backup-Archive Client Common Files
		  – Tivoli Storage Manager Backup-Archive IBM Crypto for C Library
		  – Tivoli Storage Manager Backup-Archive Help and Messages

Tivoli Storage Manager Administrative Client	TsmBaDsmadmc.pkg	Contains the Tivoli Storage Manager Administrative Client **Note:** Requires that the following packages be installed first: • Tivoli Storage Manager Backup-Archive Client Common Files • Tivoli Storage Manager Backup-Archive IBM Crypto for C Library • Tivoli Storage Manager Backup-Archive Help and Messages
Tivoli Storage Manager Sample Option Files	TsmBaSmp.pkg	Contains the Tivoli Storage Manager sample option files
Tivoli Storage Manager Backup-Archive Client Common Files	TsmBaClientCommon.pkg	Contains the files that are used by all of the Tivoli Storage Manager Backup-Archive client components.
Tivoli Storage Manager Backup-Archive Client Common Java GUI Files	TsmBaJavaCommon.pkg	Contains the files that are used by all of the Tivoli Storage Manager Backup-Archive client Java components.
Tivoli Storage Manager Backup-Archive Java GUI Files	TsmBaJavaGui.pkg	Contains the files used by the Tivoli Storage Manager Backup-Archive Java GUI component.

Table 37. Mac OS X installation packages (continued)

Package	Package Name	Package Description
Tivoli Storage Manager Backup-Archive Help and Messages	TsmBaHelpCat.pkg	Contains the Tivoli Storage Manager messages and command-line help that is used by all of the Tivoli Storage Manager Backup-Archive client components.
Tivoli Storage Manager Backup-Archive IBM Crypto for C Library	TsmBaIcc.pkg	Contains the Tivoli Storage Manager IBM Crypto for C Library that is used by all Tivoli Storage Manager Backup-Archive client components.
Tivoli Storage Manager Application Programming Interface (API)	TsmApi.pkg	Contains the following component: Tivoli Storage Manager API.
Tivoli Storage Manager Language Pack – Czech	TsmBaCS_CZ.pkg	Contains Czech localized messages for the Tivoli Storage Manager API and Backup-Archive client component.
Tivoli Storage Manager Language Pack – German	TsmBaDE_DE.pkg	Contains German localized messages for the Tivoli Storage Manager API and Backup-Archive client component.
Tivoli Storage Manager Language Pack – Spanish	TsmBaES_ES.pkg	Contains Spanish localized messages for the Tivoli Storage Manager API and Backup-Archive client component.
Tivoli Storage Manager Language Pack – French	TsmBaFR_FR.pkg	Contains French localized messages for the Tivoli Storage Manager API and Backup-Archive client component.
Tivoli Storage Manager Language Pack – Hungarian	TsmBaHU_HU.pkg	Contains Hungarian localized messages for the Tivoli Storage Manager API and Backup-Archive client component.
Tivoli Storage Manager Language Pack – Italian	TsmBaIT_IT.pkg	Contains Italian localized messages for the Tivoli Storage Manager API and Backup-Archive client component.
Tivoli Storage Manager Language Pack – Japanese	TsmBaJA_JP.pkg	Contains Japanese localized messages for the Tivoli Storage Manager API and Backup-Archive client component.

Table 37. Mac OS X installation packages (continued)

Package	Package Name	Package Description
Tivoli Storage Manager Language Pack – Korean	TsmBaKO_KR.pkg	Contains Korean localized messages for the Tivoli Storage Manager API and Backup-Archive client component.
Tivoli Storage Manager Language Pack – Polish	TsmBaPL_PL.pkg	Contains Polish localized messages for the Tivoli Storage Manager API and Backup-Archive client component.
Tivoli Storage Manager Language Pack – Brazilian Portuguese	TsmBaPT_BR.pkg	Contains Brazilian Portuguese localized messages for the Tivoli Storage Manager API and Backup-Archive client component.
Tivoli Storage Manager Language Pack – Russian	TsmBaRU_RU.pkg	Contains Russian localized messages for the Tivoli Storage Manager API and Backup-Archive client component.
Tivoli Storage Manager Language Pack – Simplified Chinese	TsmBaZH_CN.pkg	Contains Simplified Chinese localized messages for the Tivoli Storage Manager API and Backup-Archive client component.
Tivoli Storage Manager Language Pack – Traditional Chinese	TsmBaZH_TW.pkg	Contains Traditional Chinese localized messages for the Tivoli Storage Manager API and Backup-Archive client component.

Default installation directories

Here are the default directories where some files are stored as part of the client installation:

- The sample system-options file, dsm.sys.smp, is placed in the installation directory.
- If DSM_DIR is not set, the executable file dsmc, the resource files, and the dsm.sys file are stored in the default installation directory. The default installation directory is /Library/Application Support/tivoli/tsm/client/ba/bin.
- If DSM_CONFIG is not set, the client user-options file must reside in the default installation directory. If DSM_CONFIG is set, then the file must reside in the directory specified by DSM_DIR.
- If you do not define DSM_LOG, Tivoli Storage Manager writes messages to the /Library/Logs or ~/Library/Logs files.

Uninstalling the Tivoli Storage Manager Mac OS X client

You can use the following procedures to uninstall the Tivoli Storage Manager Mac OS X client.

If the Tivoli Storage Manager scheduler is configured as a startup item, use the TSM Tools for Administrators function or the StopCad.sh shell script to stop and uninstall the scheduler before proceeding.

You can use the shell script to uninstall the backup-archive client on the Mac OS X platform. The shell script name is `uninstall.sh` and it is located in the default installation directory `/Library/Application Support/tivoli/tsm/client/ba/bin`. Use the `sudo` command to run the script.

Alternately, you can complete the following steps instead of using the script:
1. Move the following folders to the trash:
   - /Applications/Tivoli Storage Manager
   - /Library/Application Support/tivoli
2. Remove the following symbolic links:
   - /usr/bin/dsmc
   - /usr/bin/dsmtca
   - /usr/bin/dsmcad
   - /usr/bin/dsmadmc
   - /usr/bin/dsmtrace
   - /usr/bin/dsmagent
   - /usr/lib/libxmlutil-6.2.0.dylib
   - /usr/lib/libtsm620xerces-c1_6_0.dylib

Tivoli Storage Manager might have created log files in these locations:
1. /Library/Logs/tivoli
2. ~/Library/Logs/tivoli

The Tivoli Storage Manager Option files (dsm.opt and dsm.sys) are usually located here:
1. /Library/Preferences/Tivoli Storage Manager
2. ~/Library/Preferences/Tivoli Storage Manager

Removal of the log and option files is optional.

Installing the Tivoli Storage Manager Solaris client

You can install the Tivoli Storage Manager Solaris client from the product DVD, or other installation media.

If a previous version of Tivoli Storage Manager has been installed, remove it before installing a new version. Refer to "Uninstalling the Tivoli Storage Manager Solaris client" on page 45 for information about removing previous Tivoli Storage Manager packages.

This installation procedure is designed to install directly from the Tivoli Storage Manager UNIX client DVD from a local or remote-mounted DVD drive.

The following installation packages are available on the DVD in the following locations:
- /dvd/tsmcli/solaris (Solaris SPARC packages)
- /dvd/tsmcli/solaris_x86 (Solaris x86 packages)

A Tivoli Storage Manager installation administration file (tsmadmin) is used in place of the default administration file (/var/sadm/install/admin), so that you are not asked about setuid, setgid or superuser permission during installation. If you want to use the default administration file, remove the -a ./tsmadmin option from the commands listed, and answer the questions about setuid, setgid, or superuser permission during installation with *y*.

Table 38. Installation package names and descriptions

Package	Package Name	Package Description
IBM Global Security Toolkit (GSKit) 32 bit	gsk8cry32 and gsk8ssl32	Contains the IBM Global Security Kit (GSKit) that provides Secure Sockets Layer (SSL) 32 bit data encryption between the Tivoli Storage Manager client and server.
IBM Global Security Toolkit (GSKit) 64 bit	gsk8cry64 and gsk8ssl64	Contains the IBM Global Security Kit (GSKit) that provides Secure Sockets Layer (SSL) 64bit data encryption between the Tivoli Storage Manager client and server.
Tivoli Storage Manager Application Programming Interface (API)	TIVsmCapi.pkg	Contains the Tivoli Storage Manager API shared libraries and samples. **Note:** On Solaris SPARC and x86-based systems the installation process installs both the 32 bit and 64 bit Tivoli Storage Manager API.

Table 38. Installation package names and descriptions (continued)

Package	Package Name	Package Description
Tivoli Storage Manager Backup-Archive (B/A)	TIVsmCba.pkg	Contains the following components: • Tivoli Storage Manager Backup-Archive Client (command-line and GUI) • Tivoli Storage Manager Administrative Client (command-line) • Tivoli Storage Manager Web Backup-Archive Client **Note:** 1. The Tivoli Storage Manager Backup-Archive package includes the 32-bit Backup-Archive Client which is used in 32-bit and 64-bit kernel modes on all supported Solaris releases (see "System requirements for Sun Solaris clients" on page 13). 2. The 32 bit Backup-Archive Client works with 32-bit and 64-bit Tivoli Storage Manager servers. 3. TCP/IP and Shared memory are supported as communication methods. 4. The Web client is an integral part of the backup-archive client package and cannot be installed without it.
Tivoli Storage Manager French language	TIVsmClFr.pkg	Contains French localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager German language	TIVsmClDe.pkg	Contains German localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Italian language	TIVsmClIt.pkg	Contains Italian localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.

Table 38. Installation package names and descriptions (continued)

Package	Package Name	Package Description
Tivoli Storage Manager Portuguese language	TIVsmClPt.pkg	Contains Portuguese localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Spanish language	TIVsmClEs.pkg	Contains Spanish localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Czech language	TIVsmClCs.pkg	Contains Czech localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Hungarian language	TIVsmClHu.pkg	Contains Hungarian localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Polish language	TIVsmClPl.pkg	Contains Polish localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Russian language	TIVsmClRu.pkg	Contains Russian localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Japanese language	TIVsmClJa.pkg	Contains Japanese localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Korean language	TIVsmClKo.pkg	Contains Korean localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Simplified Chinese language	TIVsmClSc.pkg	Contains Simplified Chinese localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.

Table 38. Installation package names and descriptions (continued)

Package	Package Name	Package Description
Tivoli Storage Manager Traditional Chinese (EUC) language	TIVsmClTc.pkg	Contains Traditional Chinese (EUC) localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.

Make sure that you install these packages in the given order. This is required for the following reasons:

- The GSKit is a prerequisite of the API
- The API is a prerequisite of the backup-archive client package
- The backup-archive client is a prerequisite of the HSM client

Follow these steps to install Tivoli Storage Manager:

1. Login as the root user and mount the Tivoli Storage Manager UNIX client DVD to /dvd.
2. Change to the directory where the packages are stored: `cd /dvd/tsmcli/ solaris`. If the Tivoli Storage Manager UNIX client DVD is not mounted to /dvd, or if the packages are stored in a different directory (for example, downloaded by ftp), change to the correct directory.
3. Install the IBM Global Security Kit (GSKit) that is a prerequisite of the Tivoli Storage Manager API package.

   a. Copy the GSKit compressed files and the tsmadmin file from the DVD to a location on your hard drive (for example, /tmp). Then change to the directory where you copied the GSKit compressed files and the tsmadmin file. Extract these GSKit compressed files, for example:

   ```
   zcat gskcrypt32-8.x.x.x.sun.sparc.tar.Z | tar xvf -
   zcat gskssl32-8.x.x.x.sun.sparc.tar.Z | tar xvf -
   zcat gskcrypt64-8.x.x.x.sun.sparc.tar.Z | tar xvf -
   zcat gskssl64-8.x.x.x.sun.sparc.tar.Z | tar xvf -
   ```

   b. Install IBM Global Security Kit (GSKit), using these commands:

   ```
   pkgadd -a  ./tsmadmin -d . gsk8cry32
   pkgadd -a ./tsmadmin -d . gsk8ssl32
   pkgadd -a  ./tsmadmin -d . gsk8cry64
   pkgadd -a ./tsmadmin -d . gsk8ssl64
   ```

   These commands install the GSKit 32 bit and 64 bit in the global zone and in all running non-global zones. If your purpose is to install Tivoli Storage Manager in a sparse-root non-global zone only, the GSKit must first be installed in the global zone for a successful GSKit installation.
4. To install the Tivoli Storage Manager API client, go to this directory, where the packages are stored:

   ```
   cd /dvd/tsmcli/solaris
   ```

   Enter this command:

   ```
   pkgadd -a ./tsmadmin -d ./TIVsmCapi.pkg TIVsmCapi
   ```

   This command installs the Tivoli Storage Manager API in the global zone and in all running non-global zones. If you want to install it in the global zone only, use the -G parameter of the pkgadd command as follows:

   ```
   pkgadd -G -a ./tsmadmin -d ./TIVsmCapi.pkg TIVsmCapi
   ```
5. To install the Tivoli Storage Manager backup-archive (B/A) client, enter the following command:

```
pkgadd -a ./tsmadmin -d ./TIVsmCba.pkg TIVsmCba
```

This installs the command-line, the GUI, and the administrative client in the global zone and in all running non-global zones. If you want to install them in the global zone only, use the -G parameter of the pkgadd command as follows:

```
 pkgadd -G -a ./tsmadmin -d ./TIVsmCba.pkg TIVsmCba
```

6. Refer to *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide* for instructions to install the Tivoli Storage Manager HSM client.

7. The Tivoli Storage Manager default language is English. To install additional languages (optional), enter the commands listed for each additional language.

Table 39. Solaris client: Language codes for installation packages

Language	Tivoli Storage Manager Language code
Simplified Chinese	ZH_CN
Traditional Chinese	ZH_TW
Czech	CS_CZ
French	FR_FR
German	DE_DE
Hungarian	HU_HU
Italian	IT_IT
Japanese	JA_JP
Korean	KO_KR
Polish	PL_PL
Brazilian Portuguese	PT_BR
Russian	RU_RU
Spanish	ES_ES

To install this language	Enter this command
NLS French messages	**pkgadd -a ./tsmadmin -d ./TIVsmClFr.pkg TIVsmClFr**
NLS German messages	**pkgadd -a ./tsmadmin -d ./TIVsmClDe.pkg TIVsmClDe**
NLS Italian messages	**pkgadd -a ./tsmadmin -d ./TIVsmClIt.pkg TIVsmClIt**
NLS Spanish messages	**pkgadd -a ./tsmadmin -d ./TIVsmClEs.pkg TIVsmClEs**
NLS Portuguese messages	**pkgadd -a ./tsmadmin -d ./TIVsmClPt.pkg TIVsmClPt**
NLS Czech messages	**pkgadd -a ./tsmadmin -d ./TIVsmClCs.pkg TIVsmClCs**
NLS Hungarian messages	**pkgadd -a ./tsmadmin -d ./TIVsmClHu.pkg TIVsmClHu**
NLS Polish messages	**pkgadd -a ./tsmadmin -d ./TIVsmClPl.pkg TIVsmClPl**
NLS Russian messages	**pkgadd -a ./tsmadmin -d ./TIVsmClRu.pkg TIVsmClRu**
NLS Japanese messages	**pkgadd -a ./tsmadmin -d ./TIVsmClJa.pkg TIVsmClJa**

To install this language	Enter this command
NLS Korean messages	**pkgadd -a ./tsmadmin -d ./TIVsmClKo.pkg TIVsmClKo**
NLS Simplified Chinese messages	**pkgadd -a ./tsmadmin -d ./TIVsmClSc.pkg TIVsmClSc**
NLS Traditional Chinese messages (EUC)	**pkgadd -a ./tsmadmin -d ./TIVsmClTc.pkg TIVsmClTc**

This installs the localized messages for the Tivoli Storage Manager API, the Tivoli Storage Manager Backup-Archive Client, and the HSM Client. If you want to install them in the global zone only, use the -G parameter of the pkgadd command.

Important: For a Solaris 10 sparse root non-global zone, the /usr file system is normally mounted as read-only (LOFS) from the global zone, and the following conditions apply:

- If Tivoli Storage Manager is not installed in the global zone, a warning message appears at the end of the installation. The message asks the global administrator to create the required links that are provided as part of the warning messages.
- If Tivoli Storage Manager is already installed in the global zone, creation of these links is not necessary. The links are already present and they are pointing to the correct executable and libraries

Related concepts

Chapter 2, "Configure the Tivoli Storage Manager client," on page 47

Uninstalling the Tivoli Storage Manager Solaris client

Follow these steps to uninstall all the packages related to Tivoli Storage Manager and HSM, including the command line, GUI, Web GUI and administrative client components.

Important: Make sure that you uninstall the packages in the given order.

A Tivoli Storage Manager installation administration file (tsmadmin) is used in place of the default administration file (/var/sadm/install/admin), so that you will not be prompted for questions about setuid, setgid or superuser permission during installation. If you want to use the default administration file, remove the "-a ./tsmadmin" option from the commands listed below and answer the questions about setuid, setgid, or superuser permission during installation with 'y'.

1. See *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide* for information about uninstalling HSM.
2. Enter the following command to uninstall the Tivoli Storage Manager backup-archive client:

   ```
   pkgrm -n -a ./tsmadmin TIVsmCba
   ```

   This uninstalls all of the components of the backup-archive client (command-line, GUI, Web GUI, and the administrative client). You cannot uninstall a single component of this package (e.g. Web GUI) without uninstalling the complete package.

   Important: The backup-archive client cannot be removed if the HSM client is installed. The HSM client must be removed first.

3. If one or more language messages packages are installed, remove them before removing the API package. Enter the following command using the root user ID:

```
pkgrm -n -a ./tsmadmin TIVsmClCs TIVsmClDe TIVsmClEs TIVsmClFr \
                        TIVsmClHu TIVsmClIt TIVsmClJa TIVsmClKo \
                        TIVsmClPl TIVsmClPt TIVsmClRu TIVsmClSc TIVsmClTc
```

4. Enter the following command to uninstall the Tivoli Storage Manager API:

```
pkgrm -n -a ./tsmadmin TIVsmCapi
```

The API cannot be removed if the TSM backup-archive client is installed. The Tivoli Storage Manager backup-archive client must be removed first.

5. Enter the following commands to uninstall the GSKit:

```
pkgrm -n -a ./tsmadmin gsk8ssl64
pkgrm -n -a ./tsmadmin gsk8cry64
pkgrm -n -a ./tsmadmin gsk8ssl32
pkgrm -n -a ./tsmadmin gsk8cry32
```

Chapter 2. Configure the Tivoli Storage Manager client

After successfully installing the Tivoli Storage Manager client, you must configure the client before performing any operations.

If you are upgrading your Tivoli Storage Manager client, it is unnecessary to reconfigure the scheduler, Web client, or other configuration settings. If the dsm.opt and dsm.sys files used by the previous client installation are available in the default installation directory or the directory or file pointed to by the DSM_CONFIG and DSM_DIR environment variables, Tivoli Storage Manager accesses these files for configuration information.

There are some *required* and *optional* configuration tasks. The *required* configuration tasks include the following:
* "Creating and modifying the client system-options file" on page 49
* "Register your workstation with a server" on page 92

Optional configuration tasks include the following:
* "Creating a default client user-options file" on page 51
* "Creating a customized client user-options file" on page 52
* "Environment variables" on page 54
* "Configuring the Web client" on page 58
* "Configuring the client scheduler" on page 59
* "Creating an include-exclude list" on page 93

UNIX and Linux client root and authorized user tasks

The phrases *root user* and *authorized user* identify tasks that either a system administrator, a user that does not have administrator authority, or both can perform.

Tivoli Storage Manager commands and options that you can use to perform these tasks are also identified in this way.

An authorized user is any user running with a real user ID of 0 (root) or a user who owns the Tivoli Storage Manager executable file and whose owner execution permission bit is set to **s**. In the following example, the user **tivoli** is an authorized user:

```
-rwsr-xr-x  1 tivoli  dsmdev  2880479  Nov  5 13:42  dsmc*
```

The following table shows the tasks that can and cannot be performed by the root user, authorized user, and the user that does not have administrator authority.

Table 40. Tasks for root users, authorized users, and users with no administrator authority

Task	Root user	Authorized user	User with no administrator authority
Install the backup-archive client	Yes	No	No

Table 40. Tasks for root users, authorized users, and users with no administrator authority (continued)

Task	Root user	Authorized user	User with no administrator authority
Register new nodes with the Tivoli Storage Manager server	Yes	Yes	No, even when the registration is set to *open* on the server
Set or recreate the Tivoli Storage Manager password for client workstations	Yes, if the password file has been deleted	Yes, if the password file has been deleted and you have write permission to the file	No, not if the password file has been deleted
Create and modify dsm.sys	Yes	Yes, if you have write permission to the file	No
Create and modify the client user-options file (dsm.opt)	Yes	Yes, if you have write permission to the file	Yes, if you own the file
Create and modify an include-exclude list	Yes	Yes	No
Backup	Yes	Yes, if you have read permission, regardless of ownership	Yes, if you own the file
Restore	Yes; when restoring to a new location or the same location, file permission and ownership are preserved	Yes; however, the operating system prevents writing to the same location if the file has read only permission. When restoring to the same location, file permissions and ownership are preserved. When restoring to a different location, the permissions of the restored file are preserved but the ownership is changed to that of the authorized user.	Yes, if you own the file or you are granted access. The ownership of the restored file changes to the non-authorized user.
Archive	Yes	Yes, if you have read permission, regardless of ownership	Yes, if you have read permission, regardless of ownership
Retrieve	Yes. When retrieving to a new location or to the same location, file permissions and ownership are preserved.	Yes. However, the operating system prevents writing to the same location if the file has read only permission. Ownership of all retrieved objects is changed to authorized user. **Note:** Retrieve for an authorized user, using the dsmj GUI only works for individual files and not when selecting directories. Retrieve each file individually using the dsmj client or use the dsmc client to restore an entire package or directory structure.	Yes, if you archived the file. Ownership of all retrieved objects is changed to the non-authorized user.
Client scheduler	Yes	Yes	No

Task	Root user	Authorized user	User with no administrator authority
Grant user access to files on the Tivoli Storage Manager server	Yes	Yes	Yes, for files that you own on the Tivoli Storage Manager server
Delete Tivoli Storage Manager server file spaces	Yes, if you are granted backup or archive delete authority by a Tivoli Storage Manager server administrator	Yes, if you are granted backup or archive delete authority by a Tivoli Storage Manager server administrator	No

Mac OS X

A system administrator is any user that is allowed to administer the system. You can check your account type using the System Preferences **Accounts** tool. System Administrators have an account type of **Admin**.

The system administrator is responsible for configuring Tivoli Storage Manager so non-administrators can manage their own data. Non-administrators (or non-authorized users) meet the following criteria:
- They do not have a user ID of 0. They are not the root user.
- They have a user account that has not been configured as a system administrator.

When a task requires additional authority to complete, you must use the authorization application to start Tivoli Storage Manager. This allows Tivoli Storage Manager to run with sufficient system privileges to complete the task. The following table lists the authorization tools to use.

Table 41. Mac OS X authorization tools and associated Tivoli Storage Manager applications

Mac OS X authorization tool	Associated Tivoli Storage Manager application
TSM Tools For Administrators	Tivoli Storage Manager StartCad.sh StopCad.sh
sudo	dsmc

Enable non-administrators to manage their own data

To enable non-administrators to use Tivoli Storage Manager to manage their own data, the system administrator must do the following in addition to the normal configuration steps: (1) Set the passwordaccess option to *generate* in the system-options file (dsm.sys), and (2) Generate the initial password by using Tivoli Storage Manager one time.

Creating and modifying the client system-options file

The client system-options file is an editable file that identifies the server and communication method, and provides the configuration for backup, archiving, hierarchical storage management, and scheduling.

Creating and modifying the client system-options file (dsm.sys) is a required root user or authorized user task.

The Tivoli Storage Manager client GUI provides a Configuration Wizard that can be used to create basic configuration files and test the connection to the Tivoli Storage Manager server. The Configuration Wizard starts automatically if the configuration files are not found when the GUI starts. If you want to modify the configuration files after they are created, click on **Setup Wizard** from the **Tools** menu of the GUI.

If you do not use the Configuration Wizard, you can create and modify the client options file manually.

For Mac OS X, if you are a root user or authorized user, copy the dsm.sys.smp file to dsm.sys in one of the following locations. The default locations are listed in the order that they are searched.
1. A location identified by the DSM_DIR environment variable
2. /Library/Application Support/tivoli/tsm/client/ba/bin/
3. /Library/Preferences/Tivoli Storage Manager/

Tivoli Storage Manager uses the first options file that is found. You *must* use the name dsm.sys for this file. The dsm.sys file is controlled by the system administrator.

For systems other than Solaris, if you are a root user or authorized user, copy the dsm.sys.smp file to dsm.sys. You *must* use the name dsm.sys for this file. It is assumed that the dsm.sys file is controlled by the system administrator.

Important: If you are reinstalling and you want to keep your existing dsm.sys file intact, do not copy the dsm.sys.smp file to dsm.sys.

Use the dsm.sys file to specify one or more servers to contact for services, and communications options for each server. This file can also include authorization options, backup and archive processing options, and scheduling options.

Edit dsm.sys to include the server or servers to which you want to connect. The following is an example of a client system-options file stanza which contains the required options for a server you want users to contact. You can specify options for more than one server:

```
Servername              server_a
COMMMethod              TCPip
TCPPort                 1500
TCPServeraddress        node.domain.company.com
```

Important: If you want to use the Web client, you must also specify the passwordaccess=*generate* option, and log in with the client to save the password.

As the default, your client node contacts the first server identified in the dsm.sys file. You can specify a different server to contact by entering the servername option in your own client user-options file (dsm.opt), or by entering that option with a command.

You can also specify a default server and a migration server (if you have the HSM client installed on your workstation) in your dsm.sys file.

The dsm.sys file can also contain the following option categories:
* Communication options
* Backup and archive processing options
* Restore and retrieve processing options
* Scheduling options
* Authorization options
* Error processing options
* Transaction processing option
* Web client options

You can modify your dsm.sys file using one of the following methods:
* From the client Java GUI main window, select **Edit** → **Client Preferences**.
* Use your favorite text editor.

Important: For Mac OS X, the system-options file must be a plain text file. By default, TextEdit does not save files as plain text. Select **Format** → **Make PlainText** to save the user-options file as a plain text file. Do not add the .txt extension.

If you update the dsm.sys file while the client is running, you must restart the process to pick up the changes.

Related concepts

Chapter 10, "Processing options," on page 241

Related tasks

"Setting options in the client options file" on page 53

Related reference

"Defaultserver" on page 294

"Passwordaccess" on page 380

Creating a default client user-options file

During installation, a sample client user-options file called dsm.opt.smp is placed in the installation directory.

Creating a default client user-options file (dsm.opt) is an optional root user or authorized user task.

You can use the default client user-options file to define the following options:
* Backup and archive processing options
* Restore and retrieve processing options
* Scheduling options
* Format options
* Command processing options
* Authorization options
* Error processing options
* Transaction processing option
* Web client options

If you are a root user or a system administrator, you can create or modify a default client user-options file for all users on your workstation. From the UNIX command line:
1. Change to the directory containing the sample file.
2. For systems other than Solaris, copy the dsm.opt.smp file to dsm.opt or to a new file name of your choice.
3. For UNIX and Linux (other than Mac OS X), use the following information:

- **For the default client user-options file**: You can store your default client user-options file in the same directory as the dsm.sys.smp file, or in any directory for which you have write access. If you rename your client user-options file or place it in a different directory, you can set the DSM_CONFIG environment variable or use the optfile option to point to your new client user-options file.
4. For Mac OS X systems, use the following information:
   - **For the default client user-options file**: You can store your default client user-options file as `/Library/Preferences/Tivoli Storage Manager/dsm.opt`.
   - **For the client user-options file**: You can copy your client user-options file to `~/Library/Preferences/Tivoli Storage Manager/dsm.opt`, or in any directory for which you have write access. If you rename your client user-options file or place it in a different directory, you must set the DSM_CONFIG environment variable or use the optfile option to point to your new client user-options file.

   Important: The user-options file must be a plain text file. By default, TextEdit does not save files as plain text. Select **Format → Make PlainText** to save the user-options file as a plain text file. Do not add the `.txt` extension.

You can then edit your dsm.opt file as appropriate for your system. From the GUI, you can edit this file using the Preferences editor by opening the **Edit** menu and selecting **Preferences**. The Preferences editor updates the client configuration files, dsm.opt and dsm.sys, if any options have changed. If you update the dsm.opt file during a session, you must restart the session to pick up the changes.

The Preferences editor uses the environment variable DSM_DIR to locate the dsm.sys file and DSM_CONFIG to locate the client user-options file (default name dsm.opt). The Preferences editor queries the server for options on the server, but cannot change the server options file.

Related concepts

Chapter 10, "Processing options," on page 241

Creating a customized client user-options file

If you want to use different options than those specified in the default client user-options file (dsm.opt), you can create your own client user-options file.

You can set all of the options that can be set in the default user options file. Creating a customized client user-options file (dsm.opt) is an optional root user or authorized user task. To create or modify a client user-options file, use the following method:

1. Contact the Tivoli Storage Manager administrator on your workstation to determine the location of the sample client user-options file dsm.opt.smp, and to get the TCP/IP address of the backup server you are connecting to and the port it listens on.
2. Copy dsm.opt.smp to your home directory as dsm.opt, or a new file name of your choice. You can store your client user-options file in any directory to which you have write access.
3. Set the DSM_CONFIG environment variable to point to your new client user-options file.
4. Edit your dsm.opt file as appropriate for your system or use the Tivoli Storage Manager Preferences editor by selecting **Edit → Preferences** from the Tivoli Storage Manager Java GUI.

Once you have created an options file, you can use the following steps to edit your options file from the GUI.

1. Open the **Edit** menu and select **Preferences**.
2. Make any necessary changes, then click **OK** to save those changes.

Note: For Mac OS X, the user options file must be a plain text file. By default, TextEdit does not save files as plain text. Select **Format** → **Make PlainText** to save the user options file as a plain text file. Do not add the `.txt` extension.

Related concepts

"Environment variables" on page 54

Related tasks

"Setting options in the client options file"

Setting options in the client options file

This section describes how to set options in your client options file, and how to use options with commands.

To set an option in this file, enter the option name and one or more blank spaces, followed by the option value.

To view or modify an options file, click **Edit** → **Preferences** from the Tivoli Storage Manager client Java GUI. The Preferences editor updates the client system-options file or client user-options file.

Note: For Mac OS X, the user options file must be a plain text file. By default, TextEdit does not save files as plain text. Select **Format** → **Make PlainText** to save the user options file as a plain text file. Do not add the `.txt` extension.

You can use the **query options** command to display all or part of your options and their current settings. This command accepts an argument to specify a subset of options. The default is to display all options.

You can also edit an options file with your favorite text editor.

To set an option in your client options file, enter the option name and one or more blank spaces, followed by the option value. For example:

```
compression  yes
nodename     client_a
```

Some options consist of only the option name, such as verbose and quiet. You can enter the entire option name or its abbreviation. For example, you can specify the verbose option as either of the following:

```
verbose
ve
```

Follow these additional rules when entering options in your client user-options file (dsm.opt):

- Begin each comment with an asterisk (*) as the first character in a line.
- Do not enter comments on the same line as an option.
- Optionally, indent options with spaces or tabs.

- Enter each option on a separate line and enter all parameters for an option on the same line. For example, to specify a group of five different file systems as your default client domain, enter one of the following:

```
domain /home /mfg /planning /mrkting /mgmt
Domain / /Volumes/fs2 /Volumes/fs2 /Volumes/fs3 /Volumes/fs4
```

- Enter one or more blank spaces between parameters.
- The maximum number of characters for a file name is 255. The maximum combined length of the file name and path name is 1024 characters. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.

If you update the client user-options file while a session is active, you must restart the session to pick up the changes.

You can use the **query options** command to display all or part of your options and their current settings. The default is to display all options.

Related reference

"Query Options" on page 519

Environment variables

Generally, setting the environment variables is an optional task. Setting these variables makes it more convenient for you to use the command line.

Set language environment variables

The Tivoli Storage Manager client automatically detects the language of the system locale and displays in that language.

For example, a French operating system displays Tivoli Storage Manager in French by default. If Tivoli Storage Manager cannot load the French message catalog, it defaults to the English (United States) language pack. For example, if the client is running in an unsupported language/locale combination, such as French/Canada or Spanish/Mexico, Tivoli Storage Manager defaults to English (United States).

You can use the **LANG** environment variable to specify the language for the UNIX and Linux clients.

Note: The operating system locale, the terminal character set, and the file name character set encoding must match in order to display or enter file names correctly.

To set the **LANG** environment variable to French, type the following:

```
export LANG=fr_FR
```

Note:
- This section does not apply to Mac OS X.
- To display the Tivoli Storage Manager help browser menus in the language of your current locale, ensure that the NLSPATH environment variable in the /etc/profile file contains the following path:

```
NLSPATH=/usr/dt/lib/nls/msg/%L/%N.cat:$NLSPATH
export NLSPATH
```

If the locale of the Tivoli Storage Manager client is the same as the character encoding of the file names, all of those files are backed up or restored correctly. If

you are running in any single-byte character set (SBCS), then all file names are valid and are backed up or restored by the backup-archive client.

If you are running in a DBCS or UTF-8 locale, file names composed of characters that are not valid in the DBCS or UTF-8 locale cannot be entered on the TSM client command line. The files might be skipped when running a backup where a wild card ("*") specification is used. If files are skipped, here is an example of the error message issued:

```
ANS4042E Object name '/testData/en_US_files/file3?'
contains one or more unrecognized characters and is not valid.
```

If all directories and files are not created with the same locale, then run your scheduled backups using a single byte character set locale. This ensures that files are not skipped due to the file name containing characters not defined in the current locale. When you restore files, run in the same locale which matches the locale encoding of the file name.

For example, file names consisting of Japanese characters might contain invalid multibyte characters if they are displayed in a Chinese locale. These files are not backed up and are not shown by the graphical user interface. If such files are found during backup, the dsmerror.log file lists the skipped files.

Tip: When using the backup-archive client scheduling mode to back up a whole system, set the **LANG** environment variable to **en_US** (or some other SBCS language) to avoid skipped files.

Related information

http://www.ibm.com/software/sysmgmt/products/support/Field_Guides.html

Set processing environment variables

There are some circumstances where you must set the environment variables.

You must set the environment variables in either of the following circumstances:
* You want to invoke Tivoli Storage Manager from a directory other than the directory where Tivoli Storage Manager is installed
* You want to specify a different options file for the backup-archive client, the administrative client, or both.

Tip: You can also specify an alternate client options file for the command-line client (not the administrative client) using the optfile option.

There are three environment variables you can set which affect Tivoli Storage Manager processing:

PATH   Includes the directory where the executable file dsmc resides.

DSM_DIR
   Specifies the directory where the executable file dsmc, the resource files, and the dsm.sys file reside. You *cannot* specify the root (/) directory for DSM_DIR.

   Refer to the installation section for your operating system to find the default installation directory information.

   When you request an image backup, image restore, snapshot-based file backup, NAS backup, or NAS restore, Tivoli Storage Manager uses the

DSM_DIR environment variable to locate the corresponding plug-in library. If DSM_DIR is not set, the client looks for the plug-in library in the following directories:

AIX   /usr/tivoli/tsm/client/ba/bin/plugins

HP-UX, all Linux clients, and Solaris
        /opt/tivoli/tsm/client/ba/bin/plugins

DSM_CONFIG

Specifies the fully-qualified path and file name of the client user options file for users who create their own personalized options file. The root directory is *not* a valid value for DSM_CONFIG. Specify a directory other than the root directory. If DSM_CONFIG is *not* set, or the -optfile option is not used, the client user options file is expected to satisfy these requirements:
1. The options file must be named `dsm.opt`.
2. For UNIX clients other than Mac OS X, if DSM_DIR is *not* set, then the file must reside in the default installation directory. If DSM_DIR *is* set, then the file must reside in the directory specified by DSM_DIR.
3. For Mac OS X, the file can reside in any of the following locations. These directories are searched in order, and the first option file found is used. `~/Library Preferences/Tivoli Storage Manager`, `/Library Preferences/Tivoli Storage Manager`, or `/Library/Application Support/tivoli/tsm/client/ba/bin`.

Refer to the installation section for your operating system to find the default installation directory information.

DSM_LOG

Points to the directory where you want the dsmerror.log, dsmwebcl.log, and dsmsched.log files to reside. You *cannot* specify the root (/) directory for DSM_LOG. The error log file contains information about any errors that occur during processing. The client creates the error log to help the Tivoli Storage Manager technical support team diagnose severe errors.

Refer to the installation section for your operating system to find the default installation directory information.

Important: Set the DSM_LOG environment variable to name a directory where read-write permissions allow the required write access for the user to create and write to the log file. This prevents log write failures and process termination. Use the *chmod* or *setacl* commands to give the files permissions that allow all expected client user IDs to read and write them. If the log names are the default names, just set the DSM_LOG environment variable to point to the directory where they reside. When Tivoli Storage Manager cannot write to the log file, an error message is written to stderror and to the syslog daemon. The syslog daemon must be running and configured to process messages with a priority of LOG_ERR for the error message to appear in the system log. Starting and configuring the syslog daemon is system specific. See "man syslogd" for information on starting the syslog daemon, and "man syslog.conf" for information on configuring the syslog daemon.

Note:
1. The errorlogname and schedlogname options override DSM_LOG.
2. The log files cannot be symbolic links. Tivoli Storage Manager detects any such links, delete the links, then exit the operation. This prevents

Tivoli Storage Manager from overwriting protected data. The affected logs are created as files in a subsequent operation.

To set the Java GUI JAVA_BIN_DIR path to find modules, add the java binary directory your PATH variable, where JAVA_BIN_DIR is the path to the "java" executable:

```
export PATH=$PATH:<JAVA_BIN_DIR>
```

Related reference

"Optfile" on page 378

Set Bourne and Korn shell variables

Enter the environment variables in the `.profile` file (Korn shell) or `.bash_profile` file (Bourne shell) in your $HOME directory.

The following is an example, where /home/davehil/dsm.opt is the path and file name for your client user-options file, and the /home/davehil directory is where you want to store the dsmerror.log file, executable file, resource files, and dsm.sys file.

```
DSM_DIR=/home/davehil
DSM_CONFIG=/home/davehil/dsm.opt
DSM_LOG=/home/davehil
export DSM_DIR DSM_CONFIG DSM_LOG
```

Set C shell variables

For the C shell, add the DSM_CONFIG, DSM_LOG and DSM_DIR variables to the `.cshrc` file in your $HOME directory.

The following is an example, where /home/davehil/dsm.opt is the path and file name for your client user-options file, and the /home/davehil directory is where you want to store the dsmerror.log file, executable file, resource files, and dsm.sys file.

```
setenv DSM_DIR /home/davehil
setenv DSM_CONFIG /home/davehil/dsm.opt
setenv DSM_LOG /home/davehil
```

Set API environment variables

If you installed the Tivoli Storage Manager client API, set the following environment variables.

DSMI_DIR

Points to your installation directory. The files dsmtca, dsm.sys, and the language files must reside in the directory pointed to by DSMI_DIR. This environment variable must be present.

DSMI_CONFIG

Full path name of your own client user-options file (dsm.opt).

DSMI_LOG

Path for dsierror.log (cannot be a symbolic link).

Note: End users of applications developed with the API should consult the installation directions for that application for special path names or guidelines for options. Ensure that directories in the environment variables are specified in the path statement. The location of the API library is especially important.

For more information about the Tivoli Storage Manager client API, see *IBM Tivoli Storage Manager Using the Application Programming Interface*, SC32-0147.

Configuring the Web client

This task guides you through the steps to configure the Web client.

To configure the Web client from the command line, perform the following steps:

1. If *MANAGEDServices scheduler* is set, specify *MANAGEDServices webclient scheduler* in the dsm.sys file.
2. Ensure that you specify passwordaccess *generate* in the dsm.sys file.
3. To generate the Tivoli Storage Manager password, start the backup-archive client by entering the following command:

   ```
   dsmc query session
   ```

   When prompted, enter your user ID and password.

   Tip: On Mac OS X, you can also generate a password by selecting the TSM Tools for Administrators application and starting Tivoli Storage Manager.
4. Start the client acceptor daemon (CAD) by entering the following command:

   ```
   dsmcad
   ```

   Tip: On Mac OS X you can start the client acceptor with the TSM Tools for Administrators application, by selecting **Start the TSM Connect Agent**.
5. To access the Web client, enter the following URL from any supported browser, where *your_machine_name* is the host name of the workstation running the Web client.:

   ```
   http://your_machine_name:1581
   ```

Port 1581 is the default port number. You can set a different port number using the httpport option.

The Tivoli Storage Manager Remote Client Agent daemon must not be started manually. It is automatically started by the Tivoli Storage Manager Client Acceptor daemon when needed.

The options applicable only to the **dsmcad** program are httpport, managedservices, and webports. You can also use options such as optfile or errorlogname. You can use the managedservices option to specify whether the Tivoli Storage Manager client acceptor daemon also manages the Tivoli Storage Manager scheduler.

All Web client messages are written to the Web client log file, dsmwebcl.log. Error messages are written to the error log file dsmerror.log, or the file you specify with the errorlogname option. The dsmwebcl.log and dsmerror.log files reside in the directory you specify with the DSM_LOG environment variable or in the current working directory. Set this environment variable to the directory where the log files should reside. The root directory is *not* a valid value for DSM_LOG. Specify a directory other than the root directory.

After installing and configuring the Web client on your workstation you can use the Web client to perform backup, archive, restore, and retrieve operations.

Related concepts

"Scheduling options" on page 254

Related tasks

"Starting a Web client session" on page 111

Related reference

"Httpport" on page 343

"Passwordaccess" on page 380

Configuring the client scheduler

This task guides you through the steps to configure the client scheduler.

Your Tivoli Storage Manager administrator can schedule Tivoli Storage Manager to perform tasks automatically.

For example, you can automatically back up files at the end of each day or archive some of your files every Friday. This procedure, known as *central scheduling*, is a cooperative effort between the server and your client node. Your administrator associates clients with one or more schedules that are part of the policy domain maintained in the server database. The Tivoli Storage Manager administrator defines central scheduling on the server and you start the client scheduler on your workstation. Once you start the client scheduler, further intervention is not necessary.

With client scheduling, you can perform the following tasks:
- Display information about available schedules.
- Display information about work that the schedule has completed.
- Modify scheduling options in the dsm.sys file.

The following information is a comparison of the client acceptor daemon-managed services and the default traditional scheduler services methods.

client acceptor daemon-managed services
- Defined using the managedservices schedule option and started with client acceptor daemon services (dsmcad).
- The client acceptor daemon starts and stops the scheduler process as needed for each scheduled action.
- Requires fewer system resources when idle.
- Tivoli Storage Manager client options and Tivoli Storage Manager server override options are refreshed each time the client acceptor daemon services start a scheduled backup.
- Cannot be used with SESSIONINITiation=SERVEROnly backups.

Tivoli Storage Manager traditional scheduler services
- Started with command `dsmc sched` command.
- Remains active, even after scheduled backup is complete.
- Requires higher use of system resources when idle.
- Tivoli Storage Manager client options and Tivoli Storage Manager server override options are only processed once when `dsmc sched` is started.
- You must restart the scheduler process for updated Tivoli Storage Manager options to take effect.

Tip: Restart the traditional scheduler periodically to free system resources previously used by system calls.

The Tivoli Storage Manager client acceptor daemon can manage the scheduler. In this case, the CAD serves as an external timer for the scheduler. When the scheduler is started, it queries the server for the next scheduled event. The event is either run immediately or the scheduler exits. The CAD restarts the scheduler when it is time to run the scheduled event. This reduces the number of background processes on your workstation and resolves memory retention problems that can occur when running the scheduler service without CAD management.

You cannot use the dsmcad for scheduling when you set the sessioninitiation option to *serveronly*.

Use the managedservices option in your client system-options file (dsm.sys) to specify whether the CAD manages the scheduler.

Important: If you run the client scheduler on the command line, the scheduler does not run as a background service.

Perform the following steps to configure the CAD to manage the client scheduler:
1. Install the Web client.
2. From the Tivoli Storage Manager GUI, select **Edit** → **Preferences**. Then select the **Web Client** category. Check the **Schedule** option in the **ManagedServices** options section. If you want to run the Web client also, check the **Both** option.
3. Start the Client Acceptor.

Note:
*   If you include files for encryption, ensure that the encryptkey *save* option is set in the options file. This option is set by selecting **Save Encryption Key Password Locally** on the Authorization panel in the preference editor. Setting this option enables unattended scheduled services. If the encryption key has not been previously saved, you must perform an attended backup of at least one file so that you get the encryption prompt and you can save the key.

Related concepts

Related tasks

Related reference

Start the client scheduler

This task guides you through the steps to schedule events using the GUI and the command-line client.

Scheduling events using the command-line client

This task guides you through the steps to schedule events using the command-line client.

You must be a system administrator to configure Tivoli Storage Manager to use the command-line client interface to handle scheduled events. The command-line tools must be installed to enable this function.

Important: If you run the client scheduler on the command line, the scheduler does not run as a background service.

Before starting the client scheduler using the client acceptor daemon, you must complete the following steps:

1. Ensure that the managedservices option includes *schedule* in the client systems options (dsm.sys) file.
2. Set the passwordaccess option to *generate* in the client systems options (dsm.sys) file.

If you include files for encryption processing, ensure that you select the **Save Encryption Key Password Locally** option in the Authorization Preferences window so that the client scheduler can perform unattended scheduled services without prompting the user for the encryption key. If the encryption key has not been previously saved, you must perform an attended backup of at least one file so that the encryption prompt is given and the key is saved.

To start the client scheduler on your client node and connect to the server schedule:

1. Change to the Tivoli Storage Manager installation directory and enter the following command:

       dsmc schedule

   When you start the client scheduler, it runs continuously until you close the window, end the process, or log off your system.
2. If the Tivoli Storage Manager client executable directory is not in your PATH environment variable, change to the installation directory and enter the following command:

       ./dsmc schedule
3. To run the **schedule** command in the background and to keep the client scheduler running, even if you log off your system, enter the following:

       nohup dsmc schedule 2> /dev/null &

If a Tivoli Storage Manager password is required for your workstation and you want to run the **schedule** command in the background, enter the password with the command.

Root User: To start the client scheduler automatically, ensure that the passwordaccess option is set to `generate` in dsm.sys, then follow the procedure below for your operating system:

To start each client scheduler automatically, add an entry to the /etc/inittab file. Typically, the run level to use is 2, 3, 4, 5, or 6, depending on the operating system and its configuration. Consult documentation for your operating system for details on run levels.

Verify the correct syntax for the entry by consulting documentation for your operating system.

Here are some examples:

For AIX, add the following entry to the /etc/inittab file:

```
itsm:2:once:/usr/bin/dsmc sched > /dev/null 2>&1 # TSM scheduler
```

In this example, the run level is set to 2.

For HP-UX, add the following entry to the /etc/inittab file:

```
itsm:3456:once:/usr/bin/dsmc sched > /dev/null 2>&1 # TSM scheduler
```

In this example, the run level is set to 3, 4, 5, and 6.

For Solaris, add the following entry to the /etc/inittab file:

```
itsm:23:once:/usr/bin/dsmc sched > /dev/null 2>&1 # TSM scheduler
```

In this example, the run level is set to 2 and 3.

Note: You must include the redirection to /dev/null in the command.

For Mac OS X:

The system administrator must generate a password so that Tivoli Storage Manager can store the password in the TSM.PWD file. This can be done either with TSM Tools for Administrators or with the command line.

A system administrator must use either of the following methods to enable the Tivoli Storage Manager client acceptor daemon to launch the command-line client in schedule mode to handle scheduled events when you start the system.

Method 1 (preferred)
- Use TSM Tools for Administrators and Start the TSM Connect Agent applications. This installs the client acceptor daemon as a system startup item so the client acceptor daemon starts after system restarts. The client acceptor daemon is also started immediately, so you do not need to restart the system to handle scheduled events.

Method 2
- Use the shell script in "/Library/Application Support/tivoli/tsm/ client/ba/bin" to install the client acceptor daemon as a startup item. The script name is StartCad.sh.

Complete the following steps to start the client acceptor daemon manually, and to check that it is running.

1. To check whether the client acceptor daemon is running, enter the following command:

   ```
   sudo ps -x | grep dsmcad
   ```

   If the client acceptor daemon is running, one of the processes listed has the path /usr/bin/dsmcad.

2. To start the client acceptor daemon manually, enter the following command in a terminal window:

   ```
   sudo /sbin/SystemStarter start dsmcad
   ```

The client scheduler can fail to initialize properly at IPL because TCP/IP is not fully initialized. You might need to customize the time for your system to compensate for this.

Tivoli Storage Manager does not recognize changes made to the dsm.opt or the dsm.sys file while the client scheduler is running. If you make changes to these files while the client scheduler is running, and you want to use the new values immediately, stop the client scheduler and restart it. For example, if you change the inclexcl option in your dsm.sys file to point to a different include-exclude options file, you must stop the client scheduler and restart it before Tivoli Storage Manager uses the new file.

To manually stop the client scheduler, enter the **kill** command if it is running in the background, or press **q** or **Ctrl+C** if it is running in the foreground. To restart the client scheduler, enter the **schedule** command again.

Tape prompting does not occur during a scheduled event regardless of the tapeprompt option setting in your options file.

Related tasks

"Configuring the client scheduler" on page 59

Related reference

"Managedservices" on page 365

"Passwordaccess" on page 380

Configuring Tivoli Storage Manager client/server communication across a firewall

In most cases, the Tivoli Storage Manager server and clients can work across a firewall.

Every firewall is different, so the firewall administrator might need to consult the instructions for the firewall software or hardware in use.

There are two methods for enabling client and server operations through a firewall:

Method 1:
> To allow clients to communicate with a server across a firewall, the following ports must be opened in the firewall by the firewall administrator:

> **TCP/IP port**
>> To enable the backup-archive client, command-line admin client, and the scheduler to run outside a firewall, the port specified by the server option *tcpport* (default 1500) must be opened by the firewall administrator. This port is set on the client and the server using the *tcpport* option. The setting must be the same on the client and server. This allows Tivoli Storage Manager scheduler communications in both *polling* and *prompted* mode, CAD-managed schedulers, and regular backup-archive client operations.

>> **Note:** The client cannot use the port specified by the *tcpadminport* option (on the server) for a client session. That port can be used for administrative sessions only.

> **HTTP port**
>> To allow the Web client to communicate with remote workstations

across a firewall, the HTTP port for the remote workstation must be opened. Use the *httpport* option in the remote workstation client options file to specify this port. The default HTTP port is 1581.

TCP/IP ports for the remote workstation
The two TCP/IP ports for the remote workstation client must be opened. Use the *webports* option in the remote workstation client options file to specify these ports. If you do not specify the values for the *webports* option, the default zero (0) causes TCP/IP to randomly assign two free port numbers.

TCP/IP port for administrative sessions
Specifies a separate TCP/IP port number on which the server is waiting for requests for administrative client sessions, allowing secure administrative sessions within a private network.

Method 2:
For the client scheduler in prompted mode, it is unnecessary to open *any* ports on the firewall. If you set the *sessioninitiation* option to *serveronly*, the client will not attempt to contact the server. *All sessions are initiated by server prompted scheduling* on the port defined on the client with the *tcpclientport* option. The *sessioninitiation* option only affects the behavior of the client scheduler running in the prompted mode.

The Tivoli Storage Manager server must set the SESSIONINITiation parameter on the **register node** and **update node** commands for each node. If the server specifies SESSIONINITiation=*clientorserver*, the default, the client can decide which method to use. If the server specifies SESSIONINITiation=*serveronly*, all sessions are initiated by the server.

Note:
1. If *sessioninitiation* is set to *serveronly*, the value for the *tcpclientaddress* client option must be the same as the value for the *HLAddress* option of the **update node** or **register node** server command. The value for the *tcpclientport* client option must be the same as the value for the *LLAddress* option of the **update node** or **register node** server command.

2. If you set the *sessioninitiation* option to *serveronly*, with the exception of CAD-managed schedulers, the command-line client, backup-archive client Java GUI, and Web client GUI still attempts to initiate sessions, but are blocked by the Tivoli Storage Manager server for nodes that have the *sessioninitiation* option set to *serveronly*.

3. When configuring the Tivoli Storage Manager scheduler on a client workstation for the first time, the scheduler service might be unable to authenticate to the server when the server contacts the client scheduler to run a schedule. This can happen when the *passwordaccess* is set to generate and the Tivoli Storage Manager server is behind a firewall and the encrypted password cannot be locally stored before the scheduler is started. To correct this problem, you need to run the scheduler from the command line (dsmc schedule), wait until a scheduled operation starts, and enter the password for your node when prompted.

4. The Tivoli Storage Manager client cannot prompt for the encryption key password in scheduler mode. If you are using Tivoli Storage Manager data encryption, you must run an initial interactive backup once to set up the encryption key by opening the TCP/IP connection from the client workstation to the server workstation. See **Method 1** for more

information about setting up this communication. After the encryption key is set, you can use server-initiated sessions to back up the files using Tivoli Storage Manager encryption.

If you set the *sessioninitiation* option to *client*, the client initiates sessions with the server (**Method 1**) by communicating on the TCP/IP port defined with the *server* option **tcpport**. This is the default. Server prompted scheduling can be used to prompt the client to connect to the server.

When using Tivoli Storage Manager across a firewall, consider the following:

- In *prompted* mode the Tivoli Storage Manager server needs to contact the client. In order to do this, some software might need to be installed on the Tivoli Storage Manager server to route the request through the firewall. This software routes the server request through a socks port on the firewall. This is typically called *socksifying* a system. Proxies are not supported, because they only route a few types of communication protocols (HTTP, FTP, GOPHER). Tivoli Storage Manager communications are not routed by proxies. It is important to note that the client creates a new connection to the Tivoli Storage Manager server when prompted. This means that the firewall configuration discussed above must be in place.

Related tasks

"Configuring the client scheduler" on page 59

Related reference

"Sessioninitiation" on page 414

"Tcpadminport" on page 434

"Tcpport" on page 439

"Webports" on page 455

Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer

Secure Sockets Layer (SSL) allows industry standard SSL-based secure communications between the Tivoli Storage Manager client and server.

The following client components support SSL:
- Command-line client
- Administrative command-line client
- Java GUI
- Client API

Only outgoing client-server connections support SSL. Incoming connections (for example, CAD, server-initiated schedule connections) do not support SSL. Client-to-client communications and Web GUI do not support SSL.

Each Tivoli Storage Manager server that is enabled for SSL must have a unique certificate. The certificate can be one of the following types:
- A certificate that is self-signed by Tivoli Storage Manager
- A certificate that is issued by a Certificate Authority (CA). The CA can be from a company such as VeriSign or Thawte, or an internal CA, maintained within your company.

Follow these steps to enable SSL communication with a self-signed certificate:

1. Obtain the Tivoli Storage Manager server self-signed certificate.

2. Configure the clients
   - To use SSL, each client must import the self-signed server certificate. You can do this using the GSKit command-line utility, gsk8capicmd.

   Important: On a 64-bit platform, use `gsk8capicmd_64` instead of `gsk8capicmd`.
3. For a disaster recovery of the Tivoli Storage Manager server, if the certificate has been lost, a new one is automatically generated by the server. Each client must obtain and import the new certificate.

Follow these steps to enable SSL communication with a CA-signed certificate:
1. Obtain the CA root certificate.
2. Configure the clients.
   - To use SSL, each client must import the root certificate of the CA. You can do this using the GSKit command-line utility, gsk8capicmd.

   Important: On a 64-bit platform, use `gsk8capicmd_64` instead of `gsk8capicmd`.

   Tip: After you have completed this step, if the server gets a new certificate signed by the same CA, the client does not need to import the root certificate again.
3. If you are recovering the Tivoli Storage Manager as part of disaster recovery, you must install the SSL certificate on the server again. If the certificate was lost, you must get a new one. You do not need to reconfigure the client if the new certificate has been signed by a CA.

If you are configuring SSL on the Tivoli Storage Manager client for the first time, you need to create the client local key database, `dsmcert.kdb`. To create the client local key database, run the following command from the DSM_DIR directory:

```
gsk8capicmd -keydb -create -populate
  -db dsmcert.kdb -pw password -stash
```

After you create the local key database, you must import the server certificate, or the CA root certificate.

If you use a self-signed certificate

Each Tivoli Storage Manager server generates its own certificate. The certificate has a fixed file name of cert.arm, and it is stored on the server workstation in the server instance directory, for example: `/opt/tivoli/tsm/server/bin/cert.arm`. If this file does not exist when you specify the SSLTCPPORT or SSLTCPADMINPORT server option in dsmserv.sys, the file is created when you restart the server with these options. Also, a self-signed SSL certificate is generated and stored in the server instance directory.

Follow these steps to set up the SSL connection to a server:
1. Obtain the certificate, the cert.arm file, from the server administrator.
2. Import the certificate into the local-client key database, and enable trust to the certificate by placing the certificate file, cert.arm, on your client workstation. After you place the certificate file on your workstation, run the following command:
   ```
   gsk8capicmd -cert -add -db dsmcert.kdb -pw <password>
     -label "TSM server <servername> self-signed key"
     -file <path to cert.arm> -format ascii -trust enable
   ```

If you use a certificate issued by a CA

If the certificate was issued by a Certificate Authority (CA) such as VeriSign or Thawte, the client is ready for SSL and you can skip the following steps. See "Certificate Authorities root certificates" on page 68 for the preinstalled list of the root certificates of the external Certificate Authorities.

If the certificate was not issued by one of the well-known Certificate Authorities, follow these steps:

1. Obtain the root certificate of the signing CA.
2. Import the root certificate into the local-client key database file, and enable trust to the certificate by placing the certificate file on your client workstation. After you place the certificate file on your workstation, run the following command:

   ```
   gsk8capicmd -cert -add -db dsmcert.kdb -pw <password>
     -label "XYZ Certificate Authority" -file
     <path to CA root certificate> -format ascii -trust enable
   ```

Important:

1. An arbitrary password provided by you is used to encrypt the key database. The password is automatically stored encrypted in the stash file (`dsmcert.sth`). The stash file is used by the Tivoli Storage Manager client to retrieve the key database password.
2. More than one server certificate can be added to the client key database file so that the client can connect to different servers. Different certificates must have different labels. The label names are not important, but use meaningful names. Also, more than one CA root certificate can be added to the client key database.
3. If you do not run the preceding commands from the DSM_DIR directory, you must copy `dsmcert.kdb` and `dsmcert.sth` into that directory.
4. By default, local key database files have root ownership and permissions and cannot be read by other users. If you plan to run the Tivoli Storage Manager client as a non-root user, you must update the permissions. For example, to grant read access to all users and groups, run the following command:

   ```
   # chmod go+r dsmcert.*
   ```
5. Due to system performance concerns, use SSL only for sessions where it is needed. Consider adding additional processor resources on the Tivoli Storage Manager server system to manage the increased requirements.

After the server certificate has been added to the client key database, add the ssl *yes* option to the client options file, and update the value of the tcpport option. It is important to understand that the server is normally set up for SSL connections on a different port. In other words, two ports are opened on the server:

1. One port accepts regular non-SSL client connections
2. Another port accepts SSL connections only

You cannot connect to a non-SSL port with an SSL-enabled client, and vice versa.

If the value of tcpport is incorrect, the client cannot connect to the server. This is normal behavior that you can correct by specifying the correct port number.

Certificate Authorities root certificates

When the Tivoli Storage Manager Backup-Archive client is delivered, it includes a list of root certificates for a number of common Certificate Authorities.

Here is the list of root certificates for a number of common Certificate Authorities that are delivered with the client:

- Entrust.net Global Secure Server Certification Authority
- Entrust.net Global Client Certification Authority
- Entrust.net Client Certification Authority
- Entrust.net Certification Authority (2048)
- Entrust.net Secure Server Certification Authority
- VeriSign Class 3 Public Primary Certification Authority
- VeriSign Class 2 Public Primary Certification Authority
- VeriSign Class 1 Public Primary Certification Authority
- VeriSign Class 4 Public Primary Certification Authority - G2
- VeriSign Class 3 Public Primary Certification Authority - G2
- VeriSign Class 2 Public Primary Certification Authority - G2
- VeriSign Class 1 Public Primary Certification Authority - G2
- VeriSign Class 4 Public Primary Certification Authority - G3
- VeriSign Class 3 Public Primary Certification Authority - G3
- VeriSign Class 2 Public Primary Certification Authority - G3
- VeriSign Class 1 Public Primary Certification Authority - G3
- Thawte Personal Premium CA
- Thawte Personal Freemail CA
- Thawte Personal Basic CA
- Thawte Premium Server CA
- Thawte Server CA
- RSA Secure Server Certification Authority

To use certificates issued by any other Certificate Authority you must install the root certificate of the Certificate Authority on all clients as part of the client configuration.

Configure your system for journal-based backup

This topic guides you through the steps to install and configure journal-based backup on your system.

Related concepts
"Journal daemon configuration"

Journal daemon configuration

Journal-Based Backup is enabled by installing and configuring the Tivoli Storage Manager journal daemon.

Configure the Tivoli Storage Manager journal daemon by editing the journal daemon configuration sample file, tsmjbbd.ini.smp, and saving it as tsmjbbd.ini. Both files should be in the default install directory.

Start the Tivoli Storage Manager journal daemon by invoking the tsmjbbd executable file, after you have configured the tsmjbbd.ini file. Also, run the jbbinittab script file to add an entry to the /etc/inittab file, to start the journal daemon after restarting your system. The tsmjbbd executable file and the jbbinittab script file should be in the default install directory.

Note:

1. The journal database processing has changed and requires resetting existing journals preserved through the journal daemon PreserveDbOnExit setting. Environments that do not use the PreserveDbOnExit setting are not affected.

   The journal daemon is reset the first time the upgraded journal daemon is started. After the journal reset, the next backup is a full non-journal incremental backup. When this backup completes successfully, subsequent backups use the journal.

2. Network and removable file systems are not supported.

3. Journal-based backup uses the filepath kernel extension to monitor file system changes. Filepath does not monitor the following directories:
   - /bin
   - /dev
   - /etc
   - /lib
   - /usr/bin
   - /usr/lib
   - /usr/share

   Therefore, journal-based backup will not process any changes made to these directories. Changes to these directories are processed if you perform periodic full incremental backups by using the -nojournal option. Periodic (weekly) full incremental backups can be a complement to daily journal-based backup.

   You can also perform individual incremental backups of these directories by using the -nojournal option. For example:

   ```
   dsmc incremental /bin/ /dev/ /etc/ /lib/ /usr/bin/ /usr/lib/ /usr/share/
     -subdir=yes -nojournal
   ```

The configuration file is periodically checked for updates to the list of journaled file systems.

File systems can be added to or removed from the list of monitored file systems without stopping and restarting the journal daemon.

Attention: If you bring a file system offline without setting the PreserveDbOnExit value of 1, the journaled file system journal database is deleted. PreserveDbOnExit=1 specifies that the journaled file system journal database is not deleted when the journal file system goes offline. The database is also valid when the journal file system comes back online.

The following is the syntax for stanza and stanza settings:

Syntax for stanzas:
   [*StanzaName*]

Syntax for stanza settings:
> *stanzaSetting=value*

Note:

1. You can specify comments in the file by beginning the line with a semicolon.
2. Stanza and value names are not case sensitive.
3. Numeric values can be specified in hexadecimal by preceding the value with `0x` otherwise they are interpreted as decimal.
4. There is no correlation between these settings and any settings in the backup-archive client options file. The journal daemon is a completely independent process and does not process backup-archive client options.

JournalSettings stanza

Settings under this stanza are global and apply to the entire journal daemon.

The following is the syntax for the *JournalSettings* stanza:

Syntax for *JournalSettings* stanza:
> **[*JournalSettings*]**

Syntax for stanza settings:
> *JournalSettings=value*

You can specify the following *JournalSettings* values:

ErrorLog
> Specifies the log file where detailed error messages generated by the journal daemon are written. The default value is `jbberror.log` in the directory of the daemon executable. For example:
>
>  ```
>  ErrorLog=/logs/jbberror.log
>  ```

JournalDir
> Directory where journal database files are stored and written.
>
> If the path given is an absolute (for example, it begins with a `dir` delimiter) `pathname`, this is the directory used. If the path given is a relative directory name, then this path is appended to each file system name and the resulting path name is used.
>
> The default is a directory named `.tSm_JoUrNaL` (used within each file system being journaled).
>
> The advantage of having the journal database on the file system being monitored is that the database stays with the file system. The disadvantage is that the updates to the database must be processed and discarded.
>
> **Important:** Directing the database to a non-journaled file system, unless this file system is shared in a cluster environment.
>
> This setting applies to all journaled file systems but can be overridden with an override stanza for each journal file system.

JournalExcludeList stanza

This list of exclude statements filters changes from being recorded in the journal database.

Changes to objects which match statements in this stanza are ignored and are not recorded in the journal database.

Note:

1. Excluding files from the journal has no bearing on those files being excluded by the backup client, other than preventing the file names from being sent to the backup client to be processed during journal-based backup. A file that is not excluded from the journal should still be excluded by the backup-archive client, if there is a matching exclude statement in the client options file.

2. The journal daemon only provides a subset of the INCLUDE/EXCLUDE function provided by the backup-archive client. The journal daemon does not support INCLUDE statements and it does not support the *exclude.dir* option.

There is no correlation between the journal exclude list and the backup-archive client exclude list.

The following pattern matching meta characters are supported:

%   Matches exactly one character.

*   Matches zero or more characters.

%EnvVar%
        Expands environment variable.

The following is an exclude statement syntax example:

```
[JournalExcludeList]
*.jbb.jbbdb
*.jbbInc.jbbdb
```

JournaledFileSystemSettings stanza

Settings under this stanza apply to each specified journaled file system unless they are overridden for individual file systems in an override stanza.

The following is the syntax for the JournaledFileSystemSettings stanza:

Syntax for *JournaledFileSystemSettings* stanza:
        [*JournaledFileSystemSettings*]

Syntax for stanza settings:
        JournaledFileSystemSetting=value

You can specify the following *JournaledFileSystemSettings* values:

JournaledFileSystems
        Specifies a space delimited list of file systems to journal. Full file system specifications and Windows virtual mount points are supported. There is no default value. You must specify at least one journaled file system for the journal daemon to run. Journaled file systems can be added or removed online without having to restart the daemon. For example:

           JournaledFileSystems=/home /other

        Important: The journal selects object names based strictly on a string match. The implication for the user is that care must be taken when selecting file systems to journal. For example, suppose you have a file system /jbb and another file system called /jbb/mnt1. If you ask the journal to monitor just /jbb, then all the changes for /jbb/mnt1 also match this string and are entered in the database. When, however, you do a back up on the client, it parses the name based on file systems, realizes the journal is not monitoring this file system and then tells the journal to remove the /jbb/mnt1 files from the database. The solution is to either monitor both or use the JournalExludeList. The same is true for the virtual

mount point options. You must be consistent with this list. For example, if you specify /home/student1 as a virtual mount point in your dsm.sys option file and you want to journal /home, then you must specify JournaledFileSystems=/home /home/student1. In this case, two separate databases are created.

JournalDbSize
> Specifies the maximum size the journal database can grow. The journal database size is expressed in bytes. A value of zero (0) indicates that the database size is limited only by the capacity of the file system containing the journal database. The default is 0 (unlimited). For example:
>
> ```
> JournalDBSize=0x10000000
> ```

NotifyBufferSize, *DirNotifyBufferSize*
> Specify change notification buffer sizes for a journaled file system. A large amount of change activity on a journaled file system might require this to be increased. The default is 0x00020000 (128 k) for files and 0x00010000 (64 k) for directories.
>
> ```
> NotifyBufferSize=0x00200000
> ```

PreserveDbOnExit setting
> This setting allows a journal to remain valid when a journaled file system goes offline and comes back online. This is useful for preserving the journal during system reboots, and resource movement.
>
> This setting allows a journal-based backup to continue processing when the daemon is restarted (or the file system comes back online) without performing a full incremental backup.
>
> **Note:** Any change activity which occurs while the journal daemon is not running (or the file system is offline) is not recorded in the journal.
>
> A value of 1 specifies that the journaled file system journal database is not deleted when the journal file system goes offline. The database is also valid when the journal file system comes back online. This value should be used with caution because any file system change activity which occurs while the journaled file system is offline is not reflected in the journal database. The default setting of 0 deletes the journaled file system journal database.
>
> **Note:** The journal is only preserved when a journaled file system comes offline normally or is brought offline when the resource is no longer available and you specify the deferFsMonStart setting. If a file system comes offline due to an error such as a notification buffer overrun, the journal is not preserved.
>
> **Note:** Set PreserveDBonExit only when you can ensure there is a controlled shutdown of the journal service. The scope of "controlled shutdown" includes stopping the journal service in order to reboot the system, failing over a cluster resource, or moving a cluster resource. The journal database can become corrupted if the shutdown is not controlled. Therefore, you should perform the following steps if the journal service was not shut down in a controlled manner or if the journal database was otherwise taken offline in an uncontrolled manner.
> 1. Stop the journal service (if it is running)
> 2. Delete the corrupted journal databases
> 3. Restart the journal service
> 4. Perform an incremental backup

An example for not deleting the journal database upon exit is:

```
preserveDBOnExit=1
```

deferFSMonStart setting
> This setting defers an attempt to begin monitoring a file system in the following cases:
> - When the specified journaled file system is not valid or available
> - The journal directory for the specified journaled file system cannot be accessed or created
>
> Resources are checked at the interval you specify using the *deferRetryInterval* setting.
>
> A value of 1 indicates that the setting is on. A value of 0 indicates that the setting is off. The default value is off (set to 0) .

deferRetryInterval setting
> This setting specifies the value in seconds that a deferred file systems with the *deferRetryInterval* setting enabled are checked for availability and brought online. The default value is 5 seconds.

logFSErrors setting
> A value of 1 indicates that all errors encountered accessing a journaled file system or journal directory should be logged. A value of zero indicates that logging of errors encountered while checking deferred file systems and journal directories is suppressed. This is usually used in conjunction with the *deferFSMonStart* setting to eliminate excessive File System unavailable messages from being written to the logs when bringing a journaled file system on-line is deferred. The default value is 1 (log all errors).

Related concepts

"Overriding stanzas"

Overriding stanzas

Any setting in the *JournaledFileSystemSettings* stanza, except for the buffer sizes, can be overridden for a particular journaled file system by creating an override stanza.

HookFileName
> In order for the journal to begin monitoring a file system, it must know the name of an existing file in that file system. This setting specifies an existing file. Access to this file is then used as a test of whether or not this file system is online. (The system definition of mounted cannot be used because we allow the use of virtual mount points in the backup-archive client. This means that the Tivoli Storage Manager system can treat a directory as a (virtual) file system).
>
> Therefore, if this file system can be mounted and unmounted, a *HookFileName* needs to be provided.
>
> If a *HookFileName* is not entered, the journal daemon attempts to create a temporary file in the highest directory, use it to begin monitoring, and then delete it.

The following is the syntax for the *JournaledFileSystemSettings* stanza:

Syntax for JournaledFileSystemSettings stanza:
> [*JournaledFileSystemSettings.fs*]

Syntax for stanza settings:
> *JournaledFileSystemSetting=override value*

For example, the override stanza name for `/home` would be:

```
JournaledFileSystemSettings./home
HookFileName=/home/doNotDeleteThisFile
```

Client-side data deduplication

Data deduplication is a method of reducing storage needs by eliminating redundant data.

Overview

Two types of data deduplication are available on Tivoli Storage Manager: *client-side data deduplication* and *server-side data deduplication*.

Client-side data deduplication is a data deduplication technique that is used on the backup-archive client to remove redundant data during backup and archive processing before the data is transferred to the Tivoli Storage Manager sever. Using client-side data deduplication can reduce the amount of data that is sent over a local area network.

Server-side data deduplication is a data deduplication technique that is done by the server. The Tivoli Storage Manager administrator can specify the data deduplication location (client or server) to use with the **DEDUP** parameter on the **REGISTER NODE** or **UPDATE NODE** server command.

Enhancements

With client-side data deduplication, you can:
- Exclude specific files on a client from data deduplication.
- Enable a data deduplication cache that reduces network traffic between the client and the server. The cache contains extents that were sent to the server in previous incremental backup operations. Instead of querying the server for the existence of an extent, the client queries its cache.

  Specify a size and location for a client cache. If an inconsistency between the server and the local cache is detected, the local cache is removed and repopulated.
- Enable both client-side data deduplication and compression to reduce the amount of data that is stored by the server. Each extent is compressed before being sent to the server. The trade-off is between storage savings and the processing power that is required to compress client data. In general, if you compress and deduplicate data on the client system, you are using approximately twice as much processing power as data deduplication alone.

  The server can work with deduplicated, compressed data. In addition, backup-archive clients earlier than V6.2 can restore deduplicated, compressed data.

Client-side data deduplication uses the following process:
- The client creates extents. *Extents* are parts of files that are compared with other file extents to identify duplicates.
- The client and server work together to identify duplicate extents. The client sends non-duplicate extents to the server.

- Subsequent client data-deduplication operations create new extents. Some or all of those extents might match the extents that were created in previous data-deduplication operations and sent to the server. Matching extents are not sent to the server again.

Benefits

Client-side data deduplication provides several advantages:
- It can reduce the amount of data that is sent over the local area network (LAN).
- The processing power that is required to identify duplicate data is offloaded from the server to client nodes. Server-side data deduplication is always enabled for deduplication-enabled storage pools. However, files that are in the deduplication-enabled storage pools and that were deduplicated by the client, do not require additional processing.
- The processing power that is required to remove duplicate data on the server is eliminated, allowing space savings on the server to occur immediately.

Client-side data deduplication has a possible disadvantage. The server does not have whole copies of client files *until* you back up the primary storage pools that contain client extents to a non-deduplicated copy storage pool. (*Extents* are parts of a file that are created during the data-deduplication process.) During storage pool backup to non-deduplicated storage pool, client extents are reassembled into contiguous files.

Server-side data deduplication offers more protection against data loss. By default, primary sequential-access storage pools that are set up for data deduplication must be backed up to non-deduplicated copy storage pools before they can be reclaimed and before duplicate data can be removed. The default ensures that the server has copies of whole files at all times, in either a primary storage pool or a copy storage pool.

Important: For further data reduction, you can enable client-side data deduplication and compression together. Each extent is compressed before it is sent to the server. Compression saves space, but it increases the processing time on the client workstation.

In a data deduplication-enabled storage pool (file pool) only one instance of a data extent is retained. Other instances of the same data extent are replaced with a pointer to the retained instance.

When client-side data deduplication is enabled, and the server has run out of storage in the destination pool, but there is a next pool defined, the server will stop the transaction. The Tivoli Storage Manager client retries the transaction without client-side data deduplication. To recover, the Tivoli Storage Manager administrator must add more scratch volumes to the original file pool, or retry the operation with deduplication disabled.

For client-side data deduplication, the Tivoli Storage Manager server must be Version 6.2 or higher.

Prerequisites

The following prerequisites apply to client-side data deduplication:
- When a client backs up or archives a file, the data is written to the primary storage pool that is specified by the copy group of the management class that is

bound to the data. To deduplicate the client data, the primary storage pool must be a sequential-access disk (FILE) storage pool that is enabled for data deduplication.

- The value of the DEDUPLICATION option on the client must be set to YES. You can set the DEDUPLICATION option in the client options file, in the preference editor of the IBM Tivoli Storage Manager client GUI, or in the client option set on the Tivoli Storage Manager server. Use the **DEFINE CLIENTOPT** command to set the DEDUPLICATION option in a client option set. To prevent the client from overriding the value in the client option set, specify **FORCE=YES**.

- Client-side data deduplication must be enabled on the server. To enable client-side data deduplication, use the **DEDUPLICATION** parameter on the **REGISTER NODE** or **UPDATE NODE** server command. Set the value of the parameter to CLIENTORSERVER.

- Files on the client must not be excluded from client-side data deduplication. By default, all files are included. You can optionally exclude specific files from data deduplication.

- Files on the client must not be encrypted. Encrypted files and files from encrypted file systems cannot be deduplicated.

- Files must be more than 2 KB. Files that are 2 KB or less are not deduplicated.

The server can limit the maximum transaction size for data deduplication by setting the CLIENTDEDUPTXNLIMIT option on the server. See the *Administrator's Guide* for details.

The following operations take precedence over client-side data deduplication:
- LAN-free data movement
- Simultaneous-write operations
- Data encryption

Important: Do not schedule or enable any of those operations during client-side data deduplication. If any of those operations occur during client-side data deduplication, client-side data deduplication is turned off, and a message is written to the error log.

The setting on the server ultimately determines whether client-side data deduplication is enabled. See Table 42.

Table 42. Data deduplication settings: Client and server

Value of the client DEDUPLICATION option	Setting on the server	Data deduplication location
Yes	On either the server or the client	Client
Yes	On the server only	Server
No	On either the server or the client	Server
No	On the server only	Server

Encrypted files

The Tivoli Storage Manager server and the backup-archive client cannot deduplicate encrypted files. If an encrypted file is encountered during data deduplication processing, the file is not deduplicated, and a message is logged.

Tip: You do not have to process encrypted files separately from files that are eligible for client-side data deduplication. Both types of files can be processed in the same operation. However, they are sent to the server in different transactions.

As a security precaution, you can take one or more of the following steps:
- Enable storage-device encryption together with client-side data deduplication.
- Use client-side data deduplication only for nodes that are secure.
- If you are uncertain about network security, enable Secure Sockets Layer (SSL).
- If you do not want certain objects (for example, image objects) to be processed by client-side data deduplication, you can exclude them on the client. If an object is excluded from client-side data deduplication and it is sent to a storage pool that is set up for data deduplication, the object is deduplicated on server.
- Use the **SET DEDUPVERIFICATIONLEVEL** command to detect possible security attacks on the server during client-side data deduplication. Using this command, you can specify a percentage of client extents for the server to verify. If the server detects a possible security attack, a message is displayed.

Related tasks
"Configuring the client for data deduplication"

Related reference
"Deduplication" on page 293

"Exclude options" on page 321

"Dedupcachepath" on page 291

"Dedupcachesize" on page 292

"Enableddedupcache" on page 312

"Ieobjtype" on page 345

Configuring the client for data deduplication

Configure the client so that you can use data deduplication to back up or archive your files.

Before you configure your client to use data deduplication, ensure that the requirements listed in "Client-side data deduplication" on page 74 are met:
- The server must enable the client for client-side data deduplication with the **DEDUP=CLIENTORSERVER** parameter on either the **REGISTER NODE** or **UPDATE NODE** command.
- The storage pool destination for the data must be a data deduplication-enabled storage pool.
- Ensure that your files are bound to the correct management class.
- Files must be larger than 2 KB.

A file can be excluded from client-side data deduplication processing. By default, all files are included. Refer to the exclude.dedup option for details.

The server can limit the maximum transaction size for data deduplication by setting the CLIENTDEDUPTXNLIMIT option on the server. See *Tivoli Storage Manager Administrator's Guide* for details.

Follow these steps to enable data deduplication on the client:
1. Add the `deduplication yes` option to the `dsm.sys` file.
2. You can also set this option in the GUI:

a. Select **Edit > Client Preferences**
   b. Click the **Deduplication** tab
   c. Select the **Enable Deduplication** check box
   d. Click **OK** to save your selections and close the Preferences Editor

After you have configured the client for data deduplication, start a backup or archive operation. When the operation completes, the backup or archive report shows the amount of data that was deduplicated in this operation, and how many files were processed by client-side data deduplication.

If you do not have enough disk space for the backup or archive operation, you can enable client-side data deduplication without local data deduplication cache on the client using these steps:

1. Add the `deduplication yes` option to the `dsm.sys` file. You can also set this option in the GUI.
2. Turn off the local data deduplication cache by adding the `ENABLEDEDUPCACHE NO` option to the `dsm.sys` file.

You can also set this option in the Tivoli Storage Manager backup-archive client preference editor by clearing the **Enable Deduplication** check box.

The following example uses the query session command to show the type of data that was processed for data deduplication:

```
tsm> q sess
TSM Server Connection Information

Server Name.............: SERVER1
Server Type.............: Windows
Archive Retain Protect..: "No"
Server Version..........: Ver. 6, Rel. 2, Lev. 0.0
Last Access Date........: 08/25/2009 13:38:18
Delete Backup Files.....: "No"
Delete Archive Files....: "Yes"
Deduplication...........: "Client Or Server"

Node Name...............: AVI
User Name...............:
```

The following example uses the query management class command to show the type of data that was processed for data deduplication:

```
tsm> q mgmt -det
Domain Name : DEDUP
Activated Policy Set Name : DEDUP
Activation date/time : 08/24/2009 07:26:09
Default Mgmt Class Name : DEDUP
Grace Period Backup Retn. : 30 day(s)
Grace Period Archive Retn.: 365 day(s)


MgmtClass Name : DEDUP
Description : dedup - values like standard
Space Management Technique : None
Auto Migrate on Non-Usage : 0
Backup Required Before Migration: YES
Destination for Migrated Files : SPACEMGPOOL
Copy Group
Copy Group Name........: STANDARD
Copy Type..............: Backup
Copy Frequency.........: 0 day(s)
Versions Data Exists...: 2 version(s)
```

```
          Versions Data Deleted..: 1 version(s)
          Retain Extra Versions..: 30 day(s)
          Retain Only Version....: 60 day(s)
          Copy Serialization.....: Shared Static
          Copy Mode..............: Modified
          Copy Destination.......: AVIFILEPOOL
          Lan Free Destination...: NO
          Deduplicate Data.......: YES

          Copy Group Name........: STANDARD
          Copy Type..............: Archive
          Copy Frequency.........: Cmd
          Retain Version.........: 365 day(s)
          Copy Serialization.....: Shared Static
          Copy Mode..............: Absolute
          Retain Initiation......: Create
          Retain Minimum.........: 65534 day(s)
          Copy Destination.......: FILEPOOL
          Lan Free Destination...: NO
          Deduplicate Data.......: YES

          ANS1900I Return code is 0.
```

Related concepts

"Client-side data deduplication" on page 74

Related reference

"Deduplication" on page 293

"Enablededupcache" on page 312

"Exclude options" on page 321

Excluding files from data deduplication

You can exclude a file from data deduplication during backup or archive processing.

You can exclude only files for archive data deduplication. You can exclude files and images (where applicable) for backup data deduplication.

If you do not want certain files to be processed by client-side data deduplication, you can exclude files from data deduplication processing using the GUI:

1. Click **Edit > Client Preferences**.
2. Click the **Include-Exclude** tab.
3. Click **Add** to open the **Define Include-Exclude Options** window.
4. Select a category for processing.
   - To exclude a file from data deduplication during archive processing, select **Archive** in the **Category** list.
   - To exclude a file from data deduplication during backup processing, select **Backup** in the **Category** list.
5. Select **Exclude.Dedup** in the **Type** list.
6. Select an item from the **Object Type** list.
   - For archive processing, only the **File** object type is available.
   - For backup processing, select one of the following object types:
     - **File**
     - **Image**
7. Specify a file or pattern in the **File or Pattern** field. You can use wildcard characters. If you do not want to type a file or pattern, click **Browse** to open a

selection window and select a file. For mounted file spaces, you can choose the directory mount point from the selection window.

8. Click **OK** to close the Define Include-Exclude Options window. The exclude options that you defined are in an exclude statement at the bottom of the Statements list box in the **Include-Exclude Preferences** tab.

9. Click **OK** to save your selections and close the Preferences Editor.

You can also exclude files from data deduplication processing in the `dsm.sys` file:

1. Add the `deduplication yes` option.

2. Exclude the files in the /Users/Administrator/Documents/Taxes/ directory by adding `EXCLUDE.dedup /Users/Administrator/Documents/Taxes/.../*`

3. Exclude client-side data deduplication for image backup of file system /home by adding `EXCLUDE.DEDUP /home/*/* IEOBJTYPE=Image`.

Important: If an object is sent to a data deduplication pool, data deduplication occurs on the server, even if the object is excluded from client-side data deduplication.

Related concepts

"Client-side data deduplication" on page 74

Related reference

"Deduplication" on page 293

"Enablededupcache" on page 312

"Exclude options" on page 321

Configuring the client to back up and archive Tivoli Storage Manager FastBack data

Before you can back up or archive Tivoli Storage Manager FastBack client data, there are some configuration tasks that you must complete.

First ensure that you have configured the backup-archive client and that you installed the Tivoli Storage Manager FastBack client.

Install the Tivoli Storage Manager FastBack client using the information in *Tivoli Storage Manager FastBack Installation and User's Guide*, here: http://publib.boulder.ibm.com/infocenter/tsmfbinf/v6.

After you install the Tivoli Storage Manager FastBack client, complete the following tasks:

1. Register a Tivoli Storage Manager node for each Tivoli Storage Manager FastBack client where Tivoli Storage Manager backs up or archives data. The node name must be the short host name of the FastBack client.

   This is a one-time configuration performed once for each Tivoli Storage Manager FastBack client whose volumes need to be backed up or archived.

   This registration step must be performed manually only when the Tivoli Storage Manager client is used as a stand-alone application.

   The Administration Center does this node registration automatically when the user creates schedules for archiving or backing up FastBack data using the Administration Center.

2. Use the server **GRANT PROXY** command to grant proxy authority to your current Tivoli Storage Manager client node on each node representing a

FastBack client created in step 1. The FastBack node should be the target, and the current Tivoli Storage Manager client node should be the proxy.

This is a one-time configuration, and is performed by the Administration Center if the backup or archive is initiated by the Administration Center.

3. Run the **set password** command to store the credentials of the FastBack repositories where the Tivoli Storage Manager client connects. Run the `set password -type=fastback` command once for each repository where the Tivoli Storage Manager client is expected to connect.

The credentials that are stored depends on these configurations:

- Tivoli Storage Manager client on a FastBack server
- Tivoli Storage Manager client on a FastBack DR Hub
- Tivoli Storage Manager client on a dedicated proxy workstation

See *IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack* for information about integrating Tivoli Storage Manager and Tivoli Storage Manager FastBack.

Related concepts

"Installation requirements for backing up and archiving Tivoli Storage Manager FastBack client data" on page 15

Related reference

"Set Password" on page 555

Cluster environment configuration and use

The term *cluster* has different meanings in different environments. It can mean highly available, high performance, load balancing, grid computing, or some combination of all of these terms.

There are currently several clustering products available for UNIX and Linux, and this section defines those aspects of a clustering environment that need to exist in order for this backup methodology to work correctly. A basic understanding of how your cluster software functions is needed. Cluster software related activities such as the development of application start and stop scripts are not described in this section.

A cluster environment refers to a UNIX or a Linux environment which exhibits the following characteristics:

- Disks are shared between physical workstations, either in an exclusive fashion (only one host has access to the logical disk at any one time) or in a concurrent fashion.
- Disks appear as local disks to the host and not as network resources.

  Important: Mount the file systems locally to the system, not through a LAN-based file share protocol such as network file system (NFS).

- Mount points of local disks are identical on each physical host in the environment (if file system /group1_disk1 fails from NodeA to NodeB, it is mounted on NodeB as /group1_disk1).

Overview of cluster environments

Cluster environments can be set up in many different configurations. This section describes the most popular cluster configurations.

Active/Active: Pool cluster resources

In an active/active configuration, each node is actively managing at least one resource and is configured as a backup for one or more resources in the cluster. Active/active is the most common form of a cluster environment.

Active/Passive: Fault tolerant

In an active/passive configuration, one node actively manages the resource.

The other node is only used if the primary node experiences a fault and the resource needs to failover. An active/passive cluster is a subtype of an active/active cluster.

Concurrent access

In a concurrent configuration, more than one node manages a resource. When a fault occurs, the resource continues to be managed by the other nodes.

Configuring the backup-archive client in a cluster environment

The Tivoli Storage Manager Backup-Archive Client is designed to manage the backup of cluster drives by placing the backup-archive client within the context of the cluster's resource groups.

This gives the advantage of backing up data from local resources (as opposed to accessing the data across the network) to maximize the performance of the backup operation and to manage the backup data relative to the resource group. Therefore, the backup-archive client can always back up data on cluster resources as if the data were local data and maximize backup performance. This ensures that critical data is getting backed up across system failures.

For example, an active/active cluster environment has three physical hosts in the cluster named NodeA, NodeB, and NodeC.

The nodes have the following qualities:
* NodeA owns the cluster resource with file systems /A1 and /A2
* NodeB owns the cluster resources with file systems /B1 and /B2
* NodeC owns the cluster resources with file systems /C1 and /C2

Note: NodeA might also have two non-clustered volumes, /fs1 and /fs2, that must be backed up.

For best backup performance, you might want all nodes in the cluster to perform the backups of the shared file systems that they own. When a node failover occurs, the backup tasks of the failed node shift to the node to which the failover occurred. For example, when NodeA fails over to NodeB, the backup of /A1 and /A2 moves to NodeB.

The following are prerequisites before configuring the backup-archive client to back up cluster and non-cluster volumes:
* A separate backup-archive client scheduler process must be run for each resource group being protected. In normal conditions, each node would have two scheduler processes: one for the cluster resources, and one for the local file systems. After a failure, additional scheduler processes are started on a node in order to protect the resources that have moved over from another node.

- The backup-archive client password files must be stored on cluster disks so that after a failure, the generated backup-archive client password is available to the takeover node.
- The file systems to be protected as part of a resource group are defined using the backup-archive client domain option. The domain option is specified in the `dsm.sys` file, which should also be stored on a cluster disk so that it can be accessed by the takeover node.

Follow the steps below to configure the Tivoli Storage Manager backup-archive client in a cluster environment.

1. Register backup-archive client node definitions on the Tivoli Storage Manager Server. All nodes in the cluster must be defined on the Tivoli Storage Manager Server. If you are defining multiple cluster resources in a cluster environment to failover independently, then unique node names must be defined per resource group. For the above sample three-way active/active cluster configuration, define three nodes (one per resource), as follows: (1) `tsm: IBM>register node nodeA nodeApw domain=standard`, (2) `tsm: IBM>register node nodeB nodeBpw domain=standard`, (3) `tsm: IBM>register node nodeC nodeCpw domain=standard`.

2. Configure the backup-archive client system-options file. Each node in the cluster must have separate server stanzas for each cluster resource group in order to be backed up in each respective `dsm.sys` file. You must ensure that the server stanzas are identical in the system option files on each node. Alternatively, you can place the `dsm.sys` file on a shared cluster location. The server stanzas defined to back up clustered volumes must have the following special characteristics:

   - The nodename option must refer to the client node name registered on the Tivoli Storage Manager Server. If the client node name is not defined, the node name defaults to the host name of the node, which might conflict with other node names used for the same client system.

     Important: Use the nodename option to explicitly define the client node.
   - The tcpclientaddress option must refer to the service IP address of the cluster node.
   - The passworddir option must refer to a directory on the shared volumes that are part of the cluster resource group.
   - The errorlogname and schedlogname options must refer to files on the shared volumes that are part of the cluster resource group to maintain a single continuous log file.
   - All include exclude statements must refer to files on the shared volumes that are part of the cluster resource group.
   - If you use the inclexcl option, it must refer to a file path on the shared volumes that are part of the cluster group.
   - The stanza names identified with the servername option must be identical on all systems.

3. Other backup-archive client options can be set as needed. In the following example, all three nodes, NodeA, NodeB, and NodeC, must have the following three server stanzas in their `dsm.sys` file:

```
Servername        clinton_nodeA
nodename          NodeA
commmethod        tcpip
tcpport           1500
tcpserveraddress  clinton.ibm.com
tcpclientaddres   nodeA.ibm.com
```

```
passwordaccess    generate
passworddir       /A1/tsm/pwd
managedservices   schedule
schedlogname      /A1/tsm/dsmsched.log
errorlogname      /A1/tsm/errorlog.log

Servername        clinton_nodeB
nodename          NodeB
commmethod        tcpip
tcpport           1500
tcpserveraddress  clinton.ibm.com
tcpclientaddres   nodeB.ibm.com
passwordaccess    generate
passworddir       /B1/tsm/pwd
managedservices   schedule
schedlogname      /B1/tsm/dsmsched.log
errorlogname      /B1/tsm/errorlog.log

Servername        clinton_nodeC
nodename          NodeC
commmethod        tcpip
tcpport           1500
tcpserveraddress  clinton.ibm.com
tcpclientaddres   nodeC.ibm.com
passwordaccess    generate
passworddir       /C1/tsm/pwd
managedservices   schedule
schedlogname      /C1/tsm/dsmsched.log
errorlogname      /C1/tsm/errorlog.log
```

4. Configure the backup-archive client user-options file. The options file
   (dsm.opt) must reside on the shared volumes in the cluster resource group.
   Define the DSM_CONFIG environment variable to refer to this file. Ensure that
   the dsm.opt file contains the following settings:

   - The value of the servername option must be the server stanza in the
     dsm.sys file which defines parameters for backing up clustered volumes.

   - Define the clustered file systems to be backed up with the domain option.

     Note: Ensure that you define the domain option in the dsm.opt file or
     specify the option in the schedule or on the Tivoli Storage Manager
     command-line client. This is to restrict clustered operations to cluster
     resources and non-clustered operations to non-clustered resources.

   In the example, nodes NodeA, NodeB, and NodeC set up their corresponding
   dsm.opt file and DSM_CONFIG environment variable as follows:

   NodeA:

   1) Set up the /A1/tsm/dsm.opt file:

   ```
   servername clinton_nodeA
   domain     /A1 /A2
   ```

   2) Issue the following command or include it in your user profile:

   ```
   export DSM_CONFIG=/A1/tsm/dsm.opt
   ```

   NodeB:

   1) Set up the /B1/tsm/dsm.opt file:

   ```
   servername clinton_nodeB
   domain     /B1 /B2
   ```

   2) Issue the following command or include it in your user profile:

```
export DSM_CONFIG=/B1/tsm/dsm.opt
```

NodeC:

1) Set up the /C1/tsm/dsm.opt file:

```
servername clinton_nodeC
domain     /C1 /C2
```

2) Issue the following command or include it in your user profile:

```
export DSM_CONFIG=/C1/tsm/dsm.opt
```

5. Set up the schedule definitions for each cluster resource group. After the basic setup is completed, define the automated schedules to back up cluster resources to meet the backup requirements. The procedure illustrates the schedule setup by using the built-in Tivoli Storage Manager scheduler. If you are using a vendor-acquired scheduler, refer to the documentation provided by the scheduler vendor.

   - Define a schedule in the policy domain where cluster nodes are defined. Ensure that the schedule's startup window is large enough to restart the schedule on the failover node in case of a failure and fallback event. This means that the schedule's duration must be set to longer than the time it takes to complete the backup of the cluster data for that node, under normal conditions.

     If the reconnection occurs within the start window for that event, the scheduled command is restarted. This scheduled incremental backup reexamines files sent to the server before the failover. The backup then "catches up" to where it stopped before the failover situation.

     In the following example, the clus_backup schedule is defined in the standard domain to start the backup at 12:30 A.M. every day with the duration set to two hours (which is the normal backup time for each node's data).

     ```
     tsm: IBM>define schedule standard clus_backup action=incr
     starttime=00:30 startdate=TODAY  Duration=2
     ```

   - Associate the schedule with the all of the backup-archive client nodes defined to backup cluster resources, as follows: (1) tsm: IBM>define association standard clus_backup nodeA, (2) tsm: IBM>define association standard clus_backup nodeB, (3) tsm: IBM>define association standard clus_backup nodeC.

6. Set up the scheduler service for backup. On each client node, a scheduler service must be configured for each resource that the node is responsible for backing up, under normal conditions. The DSM_CONFIG environment variable for each resource scheduler service must be set to refer to the corresponding dsm.opt file for that resource. For the sample configuration, the following shell scripts must be created to allow dsmcad processes to be started, as needed, from any node in the cluster.

   ```
   NodeA: /A1/tsm/startsched
   #!/bin/ksh
   export DSM_CONFIG=/A1/tsm/dsm.opt
   dsmcad
   NodeB: /B1/tsm/startsched
   #!/bin/ksh
   export DSM_CONFIG=/B1/tsm/dsm.opt
   dsmcad
   NodeC: /C1/tsm/startsched
   #!/bin/ksh
   export DSM_CONFIG=/C1/tsm/dsm.opt
   dsmcad
   ```

7. Define the Tivoli Storage Manager Backup-Archive Client to the cluster application. To continue the backup of the failed resource after a failover condition, the Tivoli Storage Manager scheduler service (for each cluster client node) must be defined as a resource to the cluster application in order to participate in the failover processing. This is required in order to continue the backup of the failed resources from the node that takes over the resource. Failure to do so would result in the incomplete backup of the failed resource. The sample scripts in step 5 can be associated with the cluster resources to ensure that they are started on nodes in the cluster while the disk resources being protected move from one node to another. The actual steps required to set up the scheduler service as a cluster resource are specific to the cluster software. Refer to your cluster application documentation for additional information.

8. Ensure each node's password is generated and cached correctly in the location specified using the passworddir option. This can be validated by performing the following steps:

   a. Validate that each node can connect to the Tivoli Storage Manager Server without the password prompt. You can do this by running the backup-archive client command line interface and issuing the following command on each node:

      ```
      #dsmc query session
      ```

      If you are prompted to submit your password, enter the password to run the command successfully and rerun the command. The second time, the command should run without the prompt for the password. If you get prompted for the password, check your configuration.

   b. Validate that the other nodes in the cluster can start sessions to the Tivoli Storage Manager Server for the failed-over node. This can be done by running the same commands, as described in the step above, on the backup nodes. For example, to validate if NodeB and NodeC can start a session as NodeA in the failover event without prompting for the password, perform the following commands on NodeB and NodeC

      ```
      #export DSM_CONFIG=/A1/tsm/dsm.opt
      #dsmc query session
      ```

      The prompt for the password might appear at this time, but this is unlikely. If you are prompted, the password was not stored in the shared location correctly. Check the passworddir option setting used for NodeA and follow the configuration steps again.

   c. Ensure that the schedules are run correctly by each node. You can trigger a schedule by setting the schedule's start time to now. Remember to reset the start time after testing is complete.

      ```
      tsm: IBM>update sched standard clus_backup starttime=now
      ```

   d. Failover and fallback between nodeA and nodeB, while nodeA is in the middle of the backup and the schedule's start window, is still valid. Verify that the incremental backup continues to run and finish successfully after failover and fallback.

   e. Issue the command below to cause a node's (nodeA) password to expire. Ensure that backup continues normally under normal cluster operations, as well as failover and fallback:

      ```
      tsm: IBM>update node nodeA forcep=yes
      ```

9. Configure the backup-archive client to back up local resources.

   a. Define client nodes on the Tivoli Storage Manager Server. Local resources should never be backed up or archived using node names defined to back up cluster data. If local volumes that are not defined as cluster resources

are backed up, separate node names (and separate client instances) must be used for both non-clustered and clustered volumes.

In the following example, assume that only NodeA has local file systems /fs1 and /fs2 to be backed up. In order to manage the local resources, register a node NodeA_local on the Tivoli Storage Manager server: tsm: IBM>register node nodeA_local nodeA_localpw domain=standard.

b. Add a separate stanza in each node's system options file dsm.sys that must back up local resources with the following special characteristics:

- The value of the tcpclientaddress option must be the local host name or IP address. This is the IP address used for primary traffic to and from the node.

- If the client backs up and restores non-clustered volumes without being connected to the cluster, the value of the tcpclientaddress option must be the boot IP address. This is the IP address used to start the system (node) before it rejoins the cluster:

  Example stanza for NodeA_local:

  ```
  Servername       clinton_nodeA_local
  nodename         nodeA_local
  commmethod       tcpip
  tcpport          1500
  tcpserveraddress clinton.ibm.com
  tcpclientaddres  nodeA_host.ibm.com
  passwordaccess   generate
  managedservices  schedule
  ```

c. Define the user options file dsm.opt in a path that is on a non-clustered resource.

- The value of the servername option must be the server stanza in the dsm.sys file which defines parameters for backing up non-clustered volumes.

- Use the domain option to define the non-clustered file systems to be backed up.

Note: Ensure that you define the domain option in the dsm.opt file or specify the option in the schedule or on the Tivoli Storage Manager client command line, in order to restrict the backup-archive operations to non-clustered volumes.

In the following example, nodeA uses the following /home/admin/dsm.opt file and sets up the DSM_CONFIG environment to refer to /home/admin/A1.dsm.opt.

Contents of /home/admin/A1.dsm.opt

```
servername ibm_nodeA_local
domain     /fs1 /fs2


export DSM_CONFIG=/home/admin/A1.dsm.opt
```

d. Define and set up a schedule to perform the incremental backup for non-clustered file systems.

```
tsm: IBM>define schedule standard local_backup action=incr
starttime=00:30 startdate=TODAY  Duration=2
```

Associate the schedule with all of the backup-archive client nodes that are defined to backup non-clustered resources.

```
tsm: IBM>define association standard nodeA_local
```

10. Restore cluster file system data. All volumes in a cluster resource are backed up under the target node defined for that cluster resource. If you need to

restore the data that resides on a cluster volume, it can be restored from the client node that owns the cluster resource at the time of the restore. The backup-archive client must use the same user options file (dsm.opt) that was used during the backup to restore the data. There are no additional setup requirements necessary to restore data on cluster volumes.

11. Restore local file system data. The non-clustered volumes are backed up under the separate node name setup for non-clustered operations. In order to restore this data, Tivoli Storage Manager Backup-Archive Client must use the same user options file dsm.opt that was used during the backup. In the example, set environment variable DSM_CONFIG to refer to /home/admin/A1.dsm.opt prior to performing a Tivoli Storage Manager client restore for the local node nodeA_local.

Related concepts

Chapter 5, "Restoring your data," on page 177

Enabling Tivoli Storage Manager Web client access in a Cluster Environment

If the Tivoli Storage Manager Web client access is needed during a failover condition, you must configure the Tivoli Storage Manager Web client acceptor daemon associated with the cluster to failover along with the cluster resource.

After you have completed the configuration steps described in the *Configuring the backup-archive client in a cluster environment* section, perform the additional steps described below to complete the Web client access setup:

1. Set up the CAD to manage the Web client and scheduler. Tivoli Storage Manager CAD should be set up to manage schedulers as well as Web client access. This reduces the number of daemons that need to be configured as cluster applications and thus simplifies the configuration and administration. When a failover occurs, the Tivoli Storage Manager CAD starts on the node that is managing the takeover.

2. Update the managedservices option in the system-options file dsm.sys on each node for each server stanza, as shown below for NodeA

```
Servername          clinton_NodeA
nodename            NodeA
commmethod          tcpip
tcpp                1500
tcps                clinton.sanjose.ibm.com
tcpclientaddres     nodeA.sanjose.ibm.com
passwordaccess      generate
passworddir         /A1/tsm/pwd
schedlogn           /A1/tsm/dsmsched.log
errorlogname        /A1/tsm/errorlog.log
managedservices     webclient schedule
```

3. Set up the CAD to use a known http port. By default, the CAD uses http port 1581, when available, for the Web client access. If this port is not available, the CAD finds the first available port, starting with 1581. In a failover condition of an active-active cluster configuration, a failover cluster host system is probably running multiple instances of the CAD. If default settings are used for the http port, the failover node uses any available port for the CAD being failed over, since the default port is probably in use by the failover host's current CAD processes. This causes problems for the Web client associated with the CAD that failed over, as the new http port is not known to the Web client users. You might use the httpport option to specify the specific ports for the Web client access for each resource. This allows you to always use the same port when connecting from a web browser, independent of the node serving the cluster

resource. Add the httpport option in the system-options file (dsm.sys) on each node for each server stanza as follows, making sure that each stanza uses a unique value:

```
Servername          clinton_NodeA
nodename            NodeA
commmethod          tcpip
tcpp                1500
tcps                clinton.sanjose.ibm.com
tcpclientaddres     nodeA.sanjose.ibm.com
passwordaccess      generate
passworddir         /A1/tsm/pwd
managedservices     webclient schedule
schedlogn           /A1/tsm/dsmsched.log
errorlogname        /A1/tsm/errorlog.log
httpport            1510


Servername          clinton_NodeB
nodename            NodeB
commmethod          tcpip
tcpp                1500
tcps                clinton.sanjose.ibm.com
tcpclientaddres     nodeB.sanjose.ibm.com
passwordaccess      generate
passworddir         /B1/tsm/pwd
managedservices     webclient schedule
schedlogn           /B1/tsm/dsmsched.log
errorlogname        /B1/tsm/errorlog.log
httpport            1511


Servername          clinton_NodeC
nodename            NodeC
commmethod          tcpip
tcpp                1500
tcps                clinton.sanjose.ibm.com
tcpclientaddres     nodeC.sanjose.ibm.com
passwordaccess      generate
passworddir         /C1/tsm/pwd
managedservices     webclient schedule
schedlogn           /C1/tsm/dsmsched.log
errorlogname        /C1/tsm/errorlog.log
httpport            1512
```

Migrating legacy AIX HACMP setups

If you are currently using the Tivoli Storage Manager Backup-Archive Client in an HACMP environment using the clusternode option, you must update your current configurations. The clusternode option is no longer supported.

Perform the following steps to update your current configurations:

1. Update the backup-archive client system-options file. As with the clusternode option, each node in the cluster must continue to have separate server stanzas for each cluster resource group to be backed up in each respective dsm.sys file. The existing dsm.sys file for NodeA might appear as follows:

```
Servername        clinton_nodeA
commmethod        tcpip
tcpp              1500
tcps              clinton.sanjose.ibm.com
tcpclientaddres   nodeA.sanjose.ibm.com
passwordaccess    generate
passworddir       /A1
clusternode       yes
managedservices   schedule
schedlogn         /A1/dsmsched.log
errorlogname      /A1/errorlog.log
```

2. Notice that no nodename option is used in this sample. Make the following changes to the existing `dsm.sys` file for `NodeA`.
   - Remove the clusternode option.
   - Specify a nodename option if you do not have one already specified.

3. The new `dsm.sys` file for `NodeA` should appear as follows:

```
Servername        clinton_nodeA
commmethod        tcpip
nodename          myclus (myclus is the existing cluster name )
tcpp              1500
tcps              clinton.sanjose.ibm.com
tcpclientaddres   nodeA.sanjose.ibm.com
passwordaccess    generate
passworddir       /A1
managedservices   schedule
schedlogn         /A1/dsmsched.log
errorlogname      /A1/errorlog.log
```

4. Register backup-archive client nodes on the Tivoli Storage Manager server. If new backup-archive client nodes are added in the first step to replace the current default value of the cluster node name, register those nodes on the Tivoli Storage Manager Server.

5. Update schedule definitions. If new backup-archive client nodes are added in the previous step, ensure that the backup schedule definitions used earlier to back up this node's data are now associated with the new client node names.

6. Validate the setup. Follow step 6 described in the "Configuring the backup-archive client in a cluster environment" section.

AIX configuration considerations prior to performing snapshot-based file backups and archives

If you are configuring your Tivoli Storage Manager AIX client to perform snapshot-based file backups and archives, there are some items that you need to consider.

- Ensure that the volume group containing the file system to be snapshot has sufficient disk space to allow JFS2 external snapshots to be created for the file system.
- Tivoli Storage Manager is using a default size of 100 percent of the file system size for the snapshot size. This value was found to be most appropriate for file systems with even moderate file system activity. If you need to lower this value based on your experience with your own file system activity, you can use the snapshotcachesize option to fine-tune this value.
- Do not enable internal snapshots when creating new JFS2 file systems on AIX 6.1 or later for all file systems managed by Tivoli Storage Manager. Tivoli Storage Manager uses external snapshots and JFS2 does not allow the creation of external and internal snapshots concurrently for the same file system.

Related reference
"Snapshotcachesize" on page 422

Configuring the file server connection for Tivoli Storage Manager

You must configure the NetApp file server connection information to run the snapshot difference incremental backup command on Tivoli Storage Manager client.

1. Define a new user on your NetApp file server using the following steps:
   a. Add the user ID to a group that permits users to log in to the file server with http and running API commands.
   b. From the file server, enter the following command to list the user ID to verify the settings and verify that the output is similar:

   ```
   useradmin user list snapdiff_user
   ```

   ```
   Name: snapdiff_user
   Info:
   Rid: 131077
   Groups: snapdiff_group
   Full Name:
   Allowed Capabilities: login-http-admin,api-*
   ```

   c. If the security.passwd.firstlogin.enable option for the user ID on the NetApp server is set to "on", ensure that all groups have the login-telnet and cli-passwd* capabilities.

   Tip: When security.passwd.firstlogin.enable option is enabled, the user ID is set to "expired" when created. The user cannot run any commands, including snapshot difference incremental, until their password is changed. Users in groups that do not have these capabilities cannot log in to the storage system. Refer to NetApp documentation for details on defining a user id and a password on the NetApp file server.

2. Export the NetApp volumes and consider the following recommended settings:

   Tip: See the NetApp documentation for details on exporting the NetApp volumes for use with AIX, or Linux hosts.
   - Map the NetApp volumes using NFS mount.
   - Ensure the NetApp volumes have the UNIX security setting

3. Set the user ID, and password on Tivoli Storage Manager client for the user ID that you created in step 1 using the following steps:
   a. Log in as the root user ID.
   b. From the Tivoli Storage Manager client command line, enter the following command:

   ```
   dsmc set password –type=filer my_file_server snapdiff_user newPassword
   ```

   Substitute the following values:

   my_file_server
   > This value is the fully qualified hostname of your NetApp file server.

   snapdiff_user
   > This value is the user ID that you created in step 1.

| **Related reference**

Register your workstation with a server

Before you can use Tivoli Storage Manager, you must set up a node name and password and your node must be registered with the server.

The process of setting up a node name and password is called *registration*. There are two types of registration: *open* and *closed*.

Your Tivoli Storage Manager administrator chooses the type of registration for your site.

You must be a root user or authorized user to perform this required task.

If you plan to use a Web client, you must have an administrative user ID with system privilege, policy privilege, client access authority, or client owner authority. When a new node is registered, an administrative user ID is automatically created for the node. By default, this node has client owner authority.

Closed registration

With closed registration, a Tivoli Storage Manager administrator must register your workstation as a client node with the server. If your enterprise uses closed registration, you must provide some information to your Tivoli Storage Manager administrator.

You must provide the following items to your Tivoli Storage Manager administrator:
- Your node name (the value returned by the **hostname** command, the name of your workstation, or the node name you specified with the *nodename* option). If you do not specify a node name with the *nodename* option, the default login ID is the name that the **hostname** command returns.
- The initial password you want to use, if required.
- Contact information, such as your name, user ID, and phone number.

Your Tivoli Storage Manager administrator defines the following for you:
- The policy domain to which your client node belongs. A policy domain contains policy sets and management classes that control how Tivoli Storage Manager manages the files you back up and archive.
- Whether you can compress files before sending them to the server.
- Whether you can delete backup and archive data from server storage.

Open registration

With open registration, a system administrator can register your workstation as a client node with the server.

The first time you start a session, Tivoli Storage Manager prompts you for information necessary to register your workstation with the server identified in your client options file. You need to supply your node name, a password, and contact information.

When you use open registration:
- Your client node is assigned to a policy domain named **standard**.
- You can delete archived copies of files from server storage, but not backup versions of files.

If necessary, your Tivoli Storage Manager administrator can change these defaults later.

Creating an include-exclude list

If you do not create an include-exclude list, Tivoli Storage Manager considers all files for backup services and uses the default management class for backup and archive services.

You must be an authorized user to perform this task.

This is an optional task but an important one.

You can create an include-exclude list to exclude a specific file or groups of files from backup services, and to assign specific management classes to files. Tivoli Storage Manager backs up any file that is not explicitly excluded. You should exclude Tivoli Storage Manager client directories from backup services. You can use the **query inclexcl** command to display a list of include and exclude statements in the order they are examined when determining whether an object is to be included.

Specify the include-exclude list in your dsm.sys file. If you define more than one server in your dsm.sys file, each server must have its own include-exclude list. This list can also contain include-exclude statements obtained from the include-exclude files you specify with the inclexcl option.

When the client processes include-exclude statements, the include-exclude statements within the include-exclude file are placed at the position occupied by the inclexcl option in dsm.sys, in the same order, and processed accordingly.

You can use the following methods to create an include-exclude list or specify an include-exclude file:

1. You can add include-exclude statements in the Tivoli Storage Manager or Web client directory tree. The online help provides detailed instructions.
2. Open the **Edit** menu and select **Client Preferences**. In the Preferences dialog, click the **Include/Exclude** section. You can specify an INCLEXCL file using the Preferences editor. However, you cannot create the INCLEXCL file using the Preferences editor.
3. Create the include-exclude list manually, following the steps listed.

You can create an include-exclude list manually by performing the following steps:

1. Determine your include and exclude requirements.
2. Locate the server stanza in your dsm.sys file. Each server stanza must have its own include-exclude list.

3. Enter your include and exclude statements. Tivoli Storage Manager evaluates all exclude.fs and exclude.dir statements *first* (regardless of their position within the include-exclude list), and removes the excluded file spaces, directories, and files from the list of objects available for processing. All other include-exclude statements are processed from the bottom of the list up. Therefore, it is important to enter all your include-exclude statements in the proper order. For example, in the following include-exclude list the `includefile.cpp` file *is not* backed up:

```
include /Users/user01/Documents/includefile.cpp
exclude /Users/user01/Documents/.../*
```

However, in the following include-exclude list the `includefile.cpp` file *is* backed up:

```
exclude /Users/user01/Documents/.../*
include /Users/user01/Documents/includefile.cpp
```

4. Save the file and close it. For Mac OS X, ensure that you save the file as plain text. Do not add the `.txt` extension.
5. Restart your Tivoli Storage Manager client to enable your include-exclude list.

Related concepts

"Considerations for Unicode-enabled clients" on page 350

"System files to exclude" on page 97

Chapter 9, "Storage management policies," on page 229

Related reference

"Inclexcl" on page 349

Include-exclude options

This topic provides brief descriptions of the include and exclude options that you can specify in your client options file, a minimum include-exclude list that excludes system files, a list of supported wildcard characters, and examples of how you might use wildcard characters with include and exclude patterns.

Exclude file spaces and directories

Use exclude.dir statements to exclude all files and subdirectories in the specified directory from processing.

Tivoli Storage Manager evaluates all exclude.dir statements *first* (regardless of their position within the include-exclude list), and removes the excluded directories and files from the list of objects available for processing. The exclude.dir statements override all include statements that match the pattern.

Table 43 lists the options you can use to exclude file spaces and directories from processing.

Table 43. Options for excluding file spaces and directories

Option	Description
exclude.fs "Exclude options" on page 321	Excludes file spaces matching the pattern. The client does not consider the specified file space for processing and the usual deleted-file expiration process cannot occur. If you exclude a file space that was previously included, existing backup versions remain on the server subject to retention rules specified in the associated management class definition.

Table 43. Options for excluding file spaces and directories (continued)

Option	Description

exclude.dir
"Exclude options"
on page 321
	Excludes a directory, its files, and all its subdirectories and their files from backup processing. For example, the statement `exclude.dir /test/dan/data1` excludes the `/test/dan/data1` directory, its files, and all its subdirectories and their files. Using the exclude.dir option is preferable over the standard exclude option to exclude large directories containing many files that you do not want to back up. You cannot use include options to override an exclude.dir statement. Only use exclude.dir when excluding an entire directory branch.
	
	• Use the following statements to exclude volumes /Volumes/disk2 altogether from backup processing. Note that the volume (/Volumes/disk2) is backed up, but all other directories on /Volumes/disk2 is excluded.
	
	  `exclude /Volumes/disk2/*`
	  `exclude.dir /Volumes/disk2/*`
	
	• An alternative method for excluding an entire volume from domain incremental backup is to use a domain statement to exclude the volume. For example:
	
	`domain "-/Volumes/disk2"`
	
	This alternative still permits selective backup processing of files on /Volumes/disk2.

Exclude files and directories from a journal-based backup

There are two methods of excluding files and directories from a journal-based backup.

Journal-based backup is only supported on AIX.

One method is to add exclude statements to the client options file to prevent the files or directories from being backed up during backup processing. The other method is to add exclude statements to the journal configuration file `tsmjbbd.ini`, to prevent journal entries from being added for the files or directories, which prevents them from being processed during a journal-based backup. If you are running AIX Version 6.1 or later, add an `exclude .snapshot` statement to the `tsmjbbd.ini` file for journal-based backup.

Note: There is no correlation between the two exclude statements. It is preferable to place exclude statements in `tsmjbbd.ini` to prevent them from entering the journal database and being processed during a journal-based backup.

Control processing with exclude statements

After Tivoli Storage Manager evaluates all exclude statements, the following options are evaluated against the remaining list of objects available for processing.

Table 44 lists the options that you can use to control processing with include and exclude statements.

Table 44. Options for controlling processing using include and exclude statements

Option	Description	Page
Back up processing		

Table 44. Options for controlling processing using include and exclude statements (continued)

Option	Description	Page
exclude exclude.backup exclude.file exclude.file.backup	*These options are equivalent*. Use these options to exclude a file or group of files from backup services and space management services (if the HSM client is installed). The exclude.backup option only excludes files from normal backup, but not from HSM.	"Exclude options" on page 321
include include.backup include.file	Use these options to include files or assign management classes for backup processing.	"Include options" on page 351
include.fs	Controls how Tivoli Storage Manager processes your file space for incremental backups.	"Include options" on page 351
Archive processing		
exclude.archive	Excludes a file or group of files from archive services.	"Exclude options" on page 321
include include.archive	*These options are equivalent*. Use these options to include files or assign management classes for archive processing.	"Include options" on page 351
Image processing		
exclude.fs.nas	Excludes file systems on the NAS file server from an image backup when used with the **backup nas** command. If you do not specify a NAS node name, the file system identified applies to all NAS file servers. The **backup nas** command ignores all other exclude statements including exclude.fs and exclude.dir statements. This option is for AIX and Solaris clients *only*.	"Exclude options" on page 321
exclude.image	Excludes mounted file systems and raw logical volumes that match the specified pattern from full image backup operations. Incremental image backup operations are unaffected by exclude.image. This option is valid for AIX, HP-UX, Solaris, and all Linux clients.	"Exclude options" on page 321
include.fs.nas	Use the include.fs.nas option to bind a management class to Network Attached Storage (NAS) file systems. To specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system image backup, use the toc option with the include.fs.nas option in your dsm.sys file. See "Toc" on page 444 for more information. This option is valid for AIX and Solaris clients *only*.	"Include options" on page 351

Option	Description	Page
include.image	Includes a file space or logical volume, assigns a management class, or allows you to assign one of several image backup processing options to a specific logical volume when used with the **backup image** command. The **backup image** command ignores all other include options. This option is valid for AIX, HP-UX, Solaris, Linux x86/x86_64, and Linux on POWER only.	"Include options" on page 351

System files to exclude

There are some system files that should be placed in the client options file so that they are excluded.

Attention: These system files are either locked by the operating system or they can cause problems during restore. These are system files that cannot be recovered without the possibility of corrupting the operating system, or temporary files with data that you can easily recreate.

Note: This section applies to Mac OS X only.

The implicitly generated statements can be seen in the lines of output of the **query inclexcl** command with the source "operating system".

The Tivoli Storage Manager client adds the following exclude statements to the include-exclude list from your dsm.sys file. Do not include any of these statements in the dsm.sys file, or duplicate entries occurs.

```
EXCLUDE.ARCHIVE "/.../Desktop DB"
EXCLUDE.BACKUP "/.../Desktop DB"
EXCLUDE.ARCHIVE "/.../Desktop DF"
EXCLUDE.BACKUP "/.../Desktop DF"
EXCLUDE.ARCHIVE /.vol
EXCLUDE.BACKUP /.vol
EXCLUDE.ARCHIVE /automount
EXCLUDE.BACKUP /automount
EXCLUDE.ARCHIVE /Network
EXCLUDE.BACKUP /Network
EXCLUDE.ARCHIVE /dev
EXCLUDE.BACKUP /dev
EXCLUDE.BACKUP /.vol/.../*
EXCLUDE.ARCHIVE /.vol/.../*
EXCLUDE.BACKUP /automount/.../*
EXCLUDE.ARCHIVE /automount/.../*
EXCLUDE.BACKUP /Network/.../*
EXCLUDE.ARCHIVE /Network/.../*
EXCLUDE.BACKUP /dev/.../*
EXCLUDE.ARCHIVE /dev/.../*
EXCLUDE.DIR /.vol
EXCLUDE.DIR /automount
EXCLUDE.DIR /Network
EXCLUDE.DIR /dev
```

Note:
1. Do not specify volumes with periods in the name (...). Tivoli Storage Manager uses the sequence of periods as part of include-exclude processing. Tivoli

Storage Manager reports an invalid include-exclude statement if a volume has a sequence of periods in the name. The volume *must* be renamed.

2. Objects that have a type of rhap and a creator of lcmt are excluded from processing. Generally, these are special file-system objects that can also be created with the **mknod** command or are UNIX mount points. The objects or mount points must be manually recreated as part of a full system restore.

You should have the following minimum include-exclude list in your include-exclude options file:

```
EXCLUDE    /.../dsmsched.log
EXCLUDE    /.../dsmprune.log
EXCLUDE    /.../dsmj.log
EXCLUDE    /.../dsmerror.log
EXCLUDE    /.../.hotfiles.bTree

EXCLUDE.DIR  /private/tmp
EXCLUDE.DIR  /private/var/vm
EXCLUDE.DIR  /private/var/tmp
EXCLUDE.DIR  /private/var/db/netinfo/local.nidb

EXCLUDE.DIR  /.../.Trashes
EXCLUDE.DIR  /.../.Spotlight-*
EXCLUDE.DIR  /.../Library/Caches
EXCLUDE.DIR  /.../.fseventsd
EXCLUDE.DIR  /Users/.../Library/Caches
```

Include and exclude files that contain wildcard characters

You must use special escape characters when including or excluding files and directories that contain wildcard characters.

Tivoli Storage Manager treats wildcard characters in different ways on different platforms.

The names of directories and files can contain different symbols. The types of symbols that are allowed depend on the operating system.

For example, on AIX, the names of directories or files can contain:

```
* ? : [ ]
```

To specify files and directories in include and exclude statements, you must use the escape character "\" to specify the wildcards. However, the escape character can only be used inside the character classes "[]".

The following examples illustrate how to specify files and directories that contain wildcard characters using the escape character and character classes in include-exclude statements.

To exclude the single directory /usr1/[dir2] from backup processing, enter the following in the dsm.sys file or the include-exclude file:

```
exclude.dir "/usr1/[\[]dir2[\]]"
```

To exclude the single file /usr1/fi*le1 from backup processing, enter the following statement in the dsm.sys file or the include-exclude file:

```
exclude "/usr1/fi[\*]le1"
```

Tip: If you use the Preferences Editor to include or exclude a single file or directory that contains wildcard characters, you must manually edit the include or exclude statement to escape the wildcard characters. The Preferences Editor does

not automatically escape the wildcard characters. Follow the previous examples to edit the include or exclude statements in the dsm.sys file or the include-exclude file.

Related concepts

"Wildcard characters" on page 464

Include and exclude groups of files with wildcard characters

You can use wildcard characters to include or exclude groups of files.

To specify groups of files that you want to include or exclude, use the wildcard characters listed in the following table. This table applies to include and exclude statements *only*.

A very large include-exclude list can decrease backup performance. Use wildcards and eliminate unnecessary include statements to keep the list as short as possible.

Table 45. Wildcard and other special characters

Character	Function
?	The match one character matches any single character *except* the directory separator; it does not match the end of the string. For example: • The **pattern** ab?, **matches** abc, but **does not match** ab, abab, or abzzz. • The **pattern** ab?rs, **matches** abfrs, but **does not match** abrs, or abllrs. • The **pattern** ab?ef?rs, **matches** abdefjrs, but **does not match** abefrs, abdefrs, or abefjrs. • The **pattern** ab??rs, **matches** abcdrs, abzzrs, but **does not match** abrs, abjrs, or abkkkrs.
*	The match-all character. For example: • The **pattern** ab*, **matches** ab, abb, abxxx, but **does not match** a, b, aa, bb. • The **pattern** ab*rs, **matches** abrs, abtrs, abrsrs, but **does not match** ars, or aabrs, abrss. • The **pattern** ab*ef*rs, **matches** abefrs, abefghrs, but **does not match** abefr, abers. • The **pattern** abcd.*, **matches** abcd.c, abcd.txt, but **does not match** abcd, abcdc, or abcdtxt.
/...	The match-*n* character matches zero or more directories.
[	The open character-class character begins the enumeration of a character class. For example: xxx[abc] matches xxxa, xxxb, or xxxc.
–	The character-class range includes characters from the first character to the last character specified. For example: xxx[a-z] matches xxxa, xxxb, xxxc, ... xxxz.
\	The literal escape character. When used within a character class, it treats the next character literally. When used outside of a character class, it is not treated in this way. For example, if you want to include the ']' in a character class, enter [...\]...]. The escape character removes the usual meaning of ']' as the close character-class character.
]	The close character-class character ends the enumeration of a character class.

Related concepts

"Wildcard characters" on page 464

Examples using wildcards with include and exclude patterns

Tivoli Storage Manager accepts the exclude.dir option which can be used to exclude directory entries. However, the include and exclude.dir options cannot be used together.

Note: In the dsm.sys file, the include and exclude options do not work with symbolic links to directories. For example, do not use /u in your include or exclude statements because /u is a symbolic link to the /home directory. Instead of entering:

```
include /u/tmp/save.fil
```

enter:

```
include /home/tmp/save.fil
```

However, the exclude option does work with symbolic links to directories when you enter a backup command with the absolute path that contains the symbolic link.

Table 46 shows how to use wildcard characters to include or exclude files.

Table 46. Using wildcard characters with include and exclude patterns

Task	Pattern
Exclude all files that end with `.doc`, except those found in the home directory of aleko, `Documents` directory.	`EXCLUDE /.../*.doc` `INCLUDE "/home/aleko/Documents/ *.doc"`
Exclude all files during backup with an extension of *bak*, except those found on the /usr file system in the `dev` directory.	`exclude /.../*.bak` `include /usr/dev/*.bak`
Exclude all files and directories under any `Documents` directory that might exist, *except* for the `Current` file of user `aleko`.	`EXCLUDE /.../Documents/.../*` `INCLUDE "/home/aleko/Documents/ Current"`
Exclude all files in any directory named "tmp" and its subdirectories, *except* for the file `/home/tmp/save.fil`.	`exclude /.../tmp/.../*` `include /home/tmp/save.fil`
Exclude any `.cpp` file in any directory on the `Vol1`, `Vol2`, `Vol3`, and `Vol4` volumes.	`EXCLUDE /Volumes/Vol[1-4]/.../*.cpp`
Exclude any `.cpp` file in any directory on the `Vol1`, `Vol2`, `Vol3`, and `Vol4` volumes.	`EXCLUDE /Volumes/Vol[1-4]/.../*.cpp`
Exclude any `.cpp` file in any directory on the `/fs1`, `/fs2`, `/fs3` and `/fs4` file systems.	`EXCLUDE /fs[1-4]/.../*.cpp`
Exclude the `.cpp` files found in the `/fs2/source` directory.	`EXCLUDE /fs2/source/*.cpp`
Exclude any `.o` file in any directory on the `/usr1`, `/usr2`, and `/usr3` file systems.	`exclude /usr[1-3]/.../*.o`
Exclude the `.o` files found in the root directory in the `usr2` file system *only*.	`exclude /usr2/*.o`
Exclude any file that resides under the `tmp` directory found in any file system.	`exclude /.../tmp/.../*`

Table 46. Using wildcard characters with include and exclude patterns (continued)

Task	Pattern
Exclude the entire directory structure `/var/spool` from all processing.	`exclude.dir /var/spool`
Exclude a single file system from backup processing.	`exclude.fs /fs1` `exclude.fs home:`
Exclude all file systems mounted anywhere in the `/test/myfs/fs01` and `/test/myfs/fs02` directory tree from backup processing.	`exclude.fs /test/myfs/fs01/.../*` `exclude.fs /test/myfs/fs02/*`
Exclude the `/home/mydir/test1` directory and any files and subdirectories under it.	`exclude.dir /home/mydir/test1`
Exclude all directories under the `/home/mydir` directory with names beginning with `test`.	`exclude.dir /home/mydir/test*`
Exclude all directories directly under the `/mydir` directory with names beginning with `test`, on any file system.	`exclude.dir /.../mydir/test*`
Exclude the raw logical volume from image backup.	`exclude.image /dev/hd0`

Exclude all symbolic links or aliases (aliases apply to Mac OS X) from backup processing, except for the `Docs` directory for `user1`.	`EXCLUDE.ATTRIBUTE.SYMLINK /.../*`
	`INCLUDE.ATTRIBUTE.SYMLINK /Users/`
	`  user1/Docs/*`

Related concepts

"Examples using wildcards with include and exclude patterns" on page 100

Related reference

"Exclude options" on page 321

Symbolic link and alias processing

Tivoli Storage Manager evaluates all exclude.fs and exclude.dir statements and removes the excluded file spaces and directories.

After this initial evaluation, Tivoli Storage Manager evaluates any include-exclude statements for controlling symbolic link and alias processing (exclude.attribute.symlink and include.attribute.symlink) against the remaining list of objects available for processing.

Alias processing applies to Mac OS X.

Table 47 defines options for controlling symbolic link and alias processing.

Table 47. Options for controlling symbolic link and alias processing

Option	Description	Page
exclude.attribute.symlink	Excludes a file or a group of files that are symbolic links or aliases from backup processing only.	"Exclude options" on page 321
include.attribute.symlink	Includes a file or a group of files that are symbolic links or aliases within broad group of excluded files for backup processing only.	"Include options" on page 351

Determine compression and encryption processing

Tivoli Storage Manager evaluates exclude.dir and any other include-exclude options controlling backup and archive processing, and then the decision is made about which files undergo compression and encryption processing.

The following options determine which files undergo compression and encryption processing.

Table 48. Options for controlling compression and encryption

Option	Description	Page
Compression processing		
exclude.compression	Excludes files from compression processing if compression=*yes* is specified. This option applies to backups and archives.	"Exclude options" on page 321
include.compression	Includes files for compression processing if compression=*yes* is specified. This option applies to backups and archives.	"Include options" on page 351
Encryption processing		
exclude.encrypt	Excludes files from encryption processing.	"Exclude options" on page 321
include.encrypt	Includes files for encryption processing. The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received. **Important:** The include.encrypt option is the only way to enable encryption on the Backup-Archive client. If no include.encrypt statements are used encryption will not occur.	"Include options" on page 351

Preview include-exclude list files

You can preview the list of objects to be backed up or archived according to the include-exclude list, prior to sending any data to the server.

The Tivoli Storage Manager client Java GUI directory tree shows detailed information of included and excluded objects. The directory tree windows in the Tivoli Storage Manager Client Java GUI allow you to select files and directories to include or exclude. You should use this **preview** command to make sure that you include and exclude the correct files. The following is a sample scenario for using the include-exclude preview function.

For example, follow these steps to back up the files on your /Users/home file space:
1. Bring up the Tivoli Storage Manager Client Java GUI and open the Backup tree. You can see all of the directories and files that have been excluded by your options file and other sources.
2. Scroll down the tree and notice that all of the *.o files in your /Volumes/home/mary/myobjdir are backed up.

3. You don't want to back up all of the `*.o` files, so you right click a `.o` file, and choose "View File Details" from the popup menu.
4. The dialog shows that these files are included, so click the "Advanced" button and create a rule to exclude all `.o` files from the `DATA:\home` file space.
5. A rule is created at the bottom of your options file. The current directory is refreshed in the Backup tree, and the `.o` files have the red `'X'`, meaning they are excluded.
6. When you look at other directories, they show the new excludes that you have added. Press "Backup" and back up the files on your `/home` file space.

Related reference

Include and exclude option processing

The Tivoli Storage Manager server can define include-exclude options using the inclexcl parameter in a client option set.

The include-exclude statements specified by the server are evaluated along with those in the client options file. The server include-exclude statements are always enforced and placed at the bottom of the include-exclude list and evaluated before the client include-exclude statements.

If the dsm.sys file include-exclude list contains one or more inclexcl options that specify include-exclude files, the include-exclude statements in these files are placed in the list position occupied by the inclexcl option and processed accordingly.

A very large include-exclude list can decrease backup performance. Use wildcards and eliminate unnecessary include statements to keep the list as short as possible.

When performing an incremental backup, Tivoli Storage Manager evaluates all exclude.fs and exclude.dir statements *first*, and removes the excluded file spaces, directories, and files from the list of objects available for processing.

After evaluating all exclude.fs and exclude.dir statements, Tivoli Storage Manager evaluates the include-exclude statements for controlling symbolic link or alias processing (exclude.attribute.symlink and include.attribute.symlink) from the bottom up and stops if it finds an include or exclude statement that matches the file it is processing. After the include-exclude statements for controlling symbolic link or alias processing are processed, Tivoli Storage Manager evaluates the remaining include-exclude list from the bottom up and stops when it finds an include or exclude statement that matches the file it is processing. The order in which the include and exclude options are entered therefore affects which files are included and excluded.

To display a list of all include-exclude statements in effect on your client workstation in the actual order they are processed, use the **query inclexcl** command.

The client program processes the list of include-exclude statements according to the following rules:
1. Files are checked; directories are *only* checked if the exclude.dir option is specified.

2. **File names are compared to the patterns in the include-exclude list from the bottom up.** When a match is found, the processing stops and checks whether the option is include or exclude. If the option is include, the file is backed up. If the option is exclude, the file *is not* backed up.

   Note:
3. If a match *is not* found, files are implicitly included and backed up.
4. When a file is backed up, it is bound to the default management class unless it matched an include statement that specified a different management class name, in which case the file is bound to that management class.

The following examples demonstrate *bottom up* processing.

Example 1

Assume that La Pomme is not the startup disk.

```
EXCLUDE /.../*.cpp
INCLUDE "/Volumes/La Pomme/Foo/.../*.cpp"
EXCLUDE "/Volumes/La Pomme/Foo/Junk/*.cpp"
```

The file being processed is: /Volumes/La Pomme/Foo/Dev/test.cpp. Processing follows these steps:

1. Rule 3 (the last include or exclude statement defined) is checked *first* because of bottom-up processing. The pattern /Volumes/La Pomme/Foo/Junk/*.cpp does not match the file name that is being processed.
2. Processing moves to Rule 2 and checks. This time, pattern /Volumes/La Pomme/Foo/.../*.cpp matches the file name that is being processed. Processing stops, the option is checked, and it is included.
3. File /Volumes/La Pomme/Foo/Dev/test.cpp is backed up.

Example 2

Assume that La Pomme is not the startup disk.

```
EXCLUDE /.../*.cpp
INCLUDE "/Volumes/La Pomme/Foo/.../*.cpp"
EXCLUDE "/Volumes/La Pomme/Foo/Junk/*.cpp"
```

The file being processed is: /Volumes/La Pomme/Widget/Sample File. Processing follows these steps:

1. Rule 3 is checked and finds no match.
2. Rule 2 is checked and finds no match.
3. Rule 1 is checked and finds no match.
4. Because a match is not found, Volumes/La Pomme/Widget/Sample File is implicitly included and is backed up.

Example 3

Assume that you defined the following statements for the include and exclude options:

```
exclude *.o
include /home/foo/.../*.o
exclude /home/foo/junk/*.o
```

The file being processed is: /home/foo/dev/test.o. Processing follows these steps:

1. Rule 3 (the last statement defined) is checked *first* because of bottom-up processing. The pattern /home/foo/junk/*.o does not match the file name that is being processed.
2. Processing moves to Rule 2 and checks. This time, pattern /home/foo/.../*.o matches the file name that is being processed. Processing stops, the option is checked, and it is include.
3. File /home/foo/dev/test.o is backed up.

Example 4

Assume that you defined the following statements for the include and exclude options:

```
exclude *.obj
include /home/foo/.../*.o
exclude /home/foo/junk/*.o
```

The file being processed is: /home/widg/copyit.txt . Processing follows these steps:

1. Rule 3 is checked and finds no match.
2. Rule 2 is checked and finds no match.
3. Rule 1 is checked and finds no match.
4. Because a match is not found, file /home/widg/copyit.txt is implicitly included and backed up.

Example 5

Assume that you defined the following statements for the include and exclude options:

```
exclude /.../*.o
include /home/foo/.../*.o
exclude /home/foo/junk/*.o
```

The current file being processed is: /home/lib/objs/printf.o. Processing follows these steps:

1. Rule 3 is checked and finds no match.
2. Rule 2 is checked and finds no match.
3. Rule 1 is checked and a match is found.
4. Processing stops, the option is checked, and it is excluded.
5. File /home/lib/objs/printf.o is not backed up.

Example 6

Assume that you defined the following statements for the include and exclude options:

```
exclude.attribute.symlink /.../*
exclude /.../*.o
include /home/foo/.../*.o
exclude /home/foo/junk/*.o
```

The current file being processed is: /home/lib/objs/printf.o. Processing follows these steps:

1. The exclude.attribute.symlink statement is checked first. If the printf.o file is a symbolic link it is excluded, otherwise proceed to the next step. Note that the exclude.attribute.symlink statements are always processed before the other include-exclude statements, regardless of their position in the include-exclude list.
2. Rule 3 is checked and finds no match.
3. Rule 2 is checked and finds no match.
4. Rule 1 is checked and a match is found.

5. Processing stops, the option is checked, and it is excluded.

6. File /home/lib/objs/printf.o is not backed up.

Related concepts

"Exclude file spaces and directories" on page 94

Chapter 10, "Processing options," on page 241

Related reference

"Exclude options" on page 321

"Query Inclexcl" on page 517

Chapter 3. Getting started

This topic describes some tasks that you need to do before you use the Tivoli Storage Manager Backup-Archive Client.

Perform these tasks before using Tivoli Storage Manager:

Tivoli Storage Manager client authentication

When using the graphical user interface or command line interface of a Tivoli Storage Manager client, you can log on using a node name and password *or* administrative user ID and password.

The client prompts for your user ID and compares it to the configured node name. If they match, the client attempts to authenticate the user ID as a node name. If the authentication fails or if the user ID does not match the configured node name, the client attempts to authenticate the user ID as an administrative user ID.

To use an administrative user ID with any of the backup-archive clients, the user ID must have one of the following authorities:

System privilege
> Authority over the entire system. An administrator with system privilege can perform any administrative task.

Policy privilege
> Authority over the node policy domain. Allows an administrator to manage policy objects, register client nodes, and schedule client operations for client nodes.

Client owner
> Authority over the registered Tivoli Storage Manager client node. You can access the client through the Web client or backup-archive client. You own the data and have a right to physically gain access to the data remotely. You can back up and restore files on the same or different system, and you can delete file spaces or archive data.

Client access
> To use the Web Client to back up and restore files on a remote client system, you must have an administrative user ID with client access authority over the node name for the remote client system. If you do not want Tivoli Storage Manager administrators with client access authority over your node name to be able to back up and restore files on your system, specify the revokeremoteaccess option in your client options file.

Client access authority only allows Tivoli Storage Manager administrators to back up and restore files on remote systems. They do not have physical access to the data. That is, they cannot restore the data belonging to the remote system to their own systems. To restore data belonging to a remote system to your own system, you must possess at least client owner authority.

To determine what authority you have, you can use either of the following methods:

- From the main Tivoli Storage Manager GUI window, select **File** → **Connection Information**.
- Use the Tivoli Storage Manager server QUERY ADMIN command from the administrative command-line client. See the appropriate *Tivoli Storage Manager Administrator's Reference* for more information about the QUERY ADMIN command.

Related reference

"Revokeremoteaccess" on page 402

Starting a Java GUI session

This section explains the different methods to start a GUI session.

The following procedure explains how to start a GUI session on Mac OS X systems:

- Double-click the Tivoli Storage Manager application to start Tivoli Storage Manager without system administrator privileges. Starting Tivoli Storage Manager without system administrator privileges lets you manage files owned by the current user.
- Double-click **TSM Tools for Administrators** and select **Tivoli Storage Manager**. After entering a System Administrator name and password, Tivoli Storage Manager starts with System Administrator privileges. Starting Tivoli Storage Manager with system administrator privileges lets you manage files owned by all users on the system.
- You can also start Tivoli Storage Manager client using the **dsmj** command. The client can run as either a *foreground* or *background* process using this command. The `dsmj` script is installed in the following location: `/Library/Application Support/tivoli/tsm/client/ba/bin`.

On UNIX systems other than Mac OS X, the Tivoli Storage Manager GUI must be run from the X Window System. If you see the Tivoli Storage Manager icon on your desktop, Tivoli Storage Manager is already running. Double-click the icon to open the Tivoli Storage Manager window. If the Tivoli Storage Manager icon does not appear on your desktop, start Tivoli Storage Manager using the **dsmj** (starts the Java GUI) command. Tivoli Storage Manager can run as either a *foreground* or *background* process using this command.

Tivoli Storage Manager locates the client user-options file (dsm.opt) and starts with the options specified in that file.

Related concepts

Chapter 2, "Configure the Tivoli Storage Manager client," on page 47

Tivoli Storage Manager password

Your Tivoli Storage Manager administrator can require you to use a password to connect to the server.

The Tivoli Storage Manager client prompts you for the password if one is required. Contact your Tivoli Storage Manager administrator if you do not know your password.

Related tasks

"Changing your password" on page 113

Setup wizard

When the client GUI starts, it checks to see whether a client options file exists.

If the client options file does not exist (which usually happens after you have installed the client for the first time on your system), the setup wizard automatically starts and guides you through the configuration process.

The client options file is dsm.sys.

Starting a command-line session

You can start a command-line session by invoking the **dsmc** command.

Note: If the /usr/bin directory contains a symbolic link to the Tivoli Storage Manager executable, and all DSM environment variables are set, you can enter the **dsmc** command from any directory. Otherwise, enter the fully qualified path of the command.

Note: On Mac OS X, system administrators can use the **sudo** command to gain additional authority so Tivoli Storage Manager can access files for all users on the system.

On the command line enter **dsmc** followed by the command *(batch mode)*. If the /usr/bin or opt/bin directory contains a symbolic link to the Tivoli Storage Manager installation directory, you can enter the **dsmc** command from any directory. Otherwise you can enter the fully qualified name.

Your Tivoli Storage Manager administrator can require you to use a password to connect to the server. The Tivoli Storage Manager client prompts you for a password, if it is required. Contact your Tivoli Storage Manager administrator if you do not know your password.

Related concepts

"Options in interactive mode" on page 462

"UNIX and Linux client root and authorized user tasks" on page 47

Chapter 11, "Using commands," on page 457

Using batch mode

Use *batch* mode to enter a single client command. When you use batch mode, you must precede the command with **dsmc**.

For example, to issue the **incremental** command, enter the following at the
command prompt:

```
dsmc incremental
```

Some commands require one or more arguments. For example, to archive a file:

```
dsmc archive /home/proj1/file1.txt
```

Depending upon the current setting of your passwordaccess option, Tivoli Storage
Manager might prompt you for your password before the command is processed
in a batch mode session.

When you enter your password, the password is not displayed on your screen.

Related reference

"Passwordaccess" on page 380

Issuing a series of commands by using interactive mode

Use *interactive* mode when you want to issue a series of commands.

Tivoli Storage Manager establishes the connection to the server only once for
interactive mode, so you can process a series of commands more quickly in
interactive mode than in batch mode.

To start a client command session in interactive mode, enter either of the following
commands:
* dsmc
* dsmc loop

The following prompt is displayed on your screen:

```
tsm>
```

When you are in interactive mode, do not precede commands with **dsmc**. For
example, instead of typing **dsmc archive** to archive a file, type only **archive**.

For example, to archive a file, enter the command with the file specification:

```
archive /home/proj1/file1.txt
```

Depending upon the current setting of the passwordaccess option, Tivoli Storage
Manager might prompt you for your password before you are allowed to enter a
command in an interactive session.

When you enter your password, the password is not displayed on your screen.

Specifying input strings that contain blank spaces or quotation marks

You must follow certain rules when you specify an input string that has blanks or
quotation marks.

Follow these rules when you specify an input string that has blank spaces or
quotation marks:
* If the input string has one or more spaces, enclose the string with either single
  or double quotation marks. You can use single or double quotation marks, as
  long as they match.
* If the input string has a single quotation mark, enclose the string within double
  quotation marks, as in this example:

```
-description="Annual backup of the accounting department's monthly reports"
```
- If the input string has a double quotation mark, enclose the string within single quotation marks, as in this example:
```
-description='New translations of "The Odyssey" and "The Iliad"'
```
- If the input string has spaces and quotation marks, enclose the string in quotation marks. The outer quotation marks must not be the same as the quotation marks within the string.

Restriction: An input string that has single and double quotation marks is not a valid input string.

The following rules apply to these types of data:
- Fully qualified names
- The description that you specify in the **archive** command
- Any value for an option value where the character string can include spaces or quotation marks

Important: You cannot use escape characters in input strings. Escape characters are treated the same as any other characters. Here are some examples where escape characters are not recognized:
- If the character string is in an option file
- If the character string is in a list file
- If the character string is entered in interactive mode

Starting: Additional considerations

You can include options as arguments to **dsmj** and **dsmc** commands. For example, you can use options to modify the format that displays dates, times, and numbers, or to include your password so that Tivoli Storage Manager does not prompt for it.

In addition, if you have more than one server defined in dsm.sys and you want to contact a different server for backup-archive services (other than the one specified in your client user-options file dsm.opt), specify the server with the servername option.

For example:
```
dsmj -servername=server_b
```

The Java GUI (dsmj) accepts command-line parameters, such as the Java -X options. Because of this, you can also now modify the Java Heap Size. For example:
```
dsmj -Xmx512M
```

Starting a Web client session

After installing the Web client on your workstation you can use the Web client to perform backup, archive, restore, and retrieve operations from any browser that is running Java Runtime Environment (JRE) 5 or 6.

The Web client facilitates the use of assistive devices for users with disabilities and contains improved keyboard navigation. The native look and feel of the platform running the browser is preserved.

Refer to the software requirements topic for your platform to find the browsers where the Web client runs.

To run the Web Client from Firefox browsers, **Enable JavaScript** must be checked. This setting is enabled by default, but to verify it:

1. Open the Firefox **Edit** menu and select **Client Preferences**.
2. In the Preferences dialog under Category, select **Advanced**, then **Scripts & Plug-ins**.
3. Ensure there is a check mark next to **Enable JavaScript for Navigator**.

If your browser does not have the correct JRE level, the Web client notifies you and if possible, tries to automatically install the correct JRE for you.

You can download and install JRE 5 or 6 (unless otherwise noted), manually from the appropriate URL:
- For AIX:

  `http://www.ibm.com/developerworks/java/jdk/aix/index.html`
- For HP-UX:

  `http://www.hp.com/products1/unix/java/?jumpid=reg_R1002_USEN`
- For Linux:

  `http://www.ibm.com/developerworks/java/jdk/linux/download.html`
- For Solaris:

  `http://java.sun.com/javase/index.jsp`

To use the Web client, specify the URL of the client workstation running the Web client in your Web browser. You also need to specify the HTTP port number defined on the client workstation; the default is 1581. For example:

  `http://myhost.mycompany.com:1581`

Note: Entering a different URL or pressing the browser **Back** button during an operation disconnects the Web client and causes the current operation to end.

Related tasks

"Configuring the Web client" on page 58

User privileges

If you plan to use the Web client, ensure that you were assigned an administrative user ID with system privilege, policy privilege, client access authority, or client owner authority.

When a new node is registered with the server, by default it is given an admin ID of the same node name with client owner authority.

Note: You can use the revokeremoteaccess option to prevent a Tivoli Storage Manager administrator with client access privilege from performing client operations on your workstation through the Web client. However, Tivoli Storage Manager administrators with client owner privilege, system privilege, or policy privilege can still perform client operations on your workstation through the Web client.

Related concepts

"Tivoli Storage Manager client authentication" on page 107

Related reference

"Revokeremoteaccess" on page 402

Start the client scheduler automatically

You can start the client scheduler automatically when you start your workstation.

If the Tivoli Storage Manager administrator has defined schedules for your node, starting the client scheduler permits you to automatically back up your workstation (or perform other scheduled actions).

You can also use the Tivoli Storage Manager Client Acceptor service to manage the scheduler.

Related tasks

"Set the client scheduler process to run as a background task and start automatically at startup" on page 218

Changing your password

Your Tivoli Storage Manager administrator can require you to use a password to connect to the server.

Tivoli Storage Manager prompts you for the password if one is required. Contact your Tivoli Storage Manager administrator if you do not know your password.

Important: The password discussed in this topic is different than the password used for encrypting files.

To change your password from the GUI:
1. On Mac OS X clients, start Tivoli Storage Manager with TSM Tools for Administrators.
2. From the main window, open the **Utilities** menu and select **Change password**.
3. Enter your current and new passwords, and enter your new password again in the **Verify password** field.
4. Click **Change**.

To change your password from the command-line client, enter this command:

For UNIX, Linux, and Windows clients:

```
dsmc set password
```

For Mac OS X clients, enter this command to change your password from the command-line client:

```
sudo dsmc set password
```

Then, enter your old and new passwords when prompted.

A Tivoli Storage Manager password can be up to 63 characters. Valid characters are:

Character	Description

A–Z	Any letter; A through Z, upper or lower case
0–9	Any number; 0 through 9
+	Plus
.	Period
_	Underscore
-	Hyphen
&	Ampersand

A password is not case sensitive.

Related concepts

"Start the client scheduler automatically" on page 113

Related tasks

"Starting: Additional considerations" on page 111

Related reference

"Password" on page 379

"Set Password" on page 555

Sorting file lists using the Tivoli Storage Manager GUI

This topic explains how to work with your files using the Tivoli Storage Manager GUI.

Table 49. Working with your files using the Tivoli Storage Manager GUI

Task	Procedure
Displaying files	To display files in a directory, click the folder icon next to the directory name. The files appear in the File List box on the right.
Sorting the file list	• Click the appropriate column heading in the File List box.
Display active and inactive backup versions	• Click the **Display Active/Inactive Files** option from the **View** menu. • Click the **Display both active and inactive files** tool on the tool bar.
Display only active backup versions	Click the **Display active files only** option from the **View** menu.
Selecting files to restore or retrieve.	• Click the selection box next to the directory or file name that you want to restore or retrieve. • Highlight the files that you want to restore or retrieve and click the Select Items tool on the tool bar. • Highlight the files that you want to restore or retrieve and click the **Select Items** option from the **Edit** menu.
Deselecting files	• Click the checked selection box next to the directory or file name. • Highlight the files that you want to deselect and click the Deselect Items tool on the tool bar. • Highlight the files that you want to deselect and click the **Deselect Items** option from the **Edit** menu.
Displaying file information	• Highlight the file name, and click the **View File Details** button on the tool bar. • Highlight the file name, and select **File Details** from the **View** menu.

Note:

1. Unless otherwise noted, the tasks and procedures in the above table apply to all Tivoli Storage Manager GUIs.
2. Using the Tivoli Storage Manager GUIs, you can sort a list of files by various attributes, such as name, directory, size, or modification date. Sorting files by the last backup date can be useful in determining what date and time to use for the point-in-time function.
3. An *active* file is the most recent backup version of a file that existed on your workstation when you ran your last backup. All other backup versions of that file are *inactive*. Only active backup versions of files are displayed, unless you select the **Display active/inactive files** menu option. If you delete the file from your workstation, the active version becomes inactive the next time you run an incremental backup.

   On the command-line client, you can use **query** commands with the inactive option to display both active and inactive objects. You can use **restore** commands with the pick and inactive options to produce the list of active and inactive backups to choose from.

Related reference

"Inactive" on page 348

"Pick" on page 382

Displaying online help

You can display online help in any of the following ways: On the Tivoli Storage Manager client GUIs, from the Web client, or from the dsmc command line.

- **On Tivoli Storage Manager client GUIs:**
  - Open the **Help** menu.
  - Click the **Help** button in the current window.
  - Click the GUI question mark (**?**) icon, which displays online information about the current operation.
- **From the dsmc command line:** Enter the **help** command. The complete table of contents for the available help text is displayed.

For the GUI-based help, the help pages contain HTML style hyper-links to other topics related to the current operation you are performing. You can select a topic of interest in the **Ask** field and search the online help for matches. Search results are sorted by relevance. Use the left and right arrow keys in the lower right-hand corner to navigate between help pages.

Related reference

"Help" on page 491

Ending a session

You can end a Tivoli Storage Manager client session from the Tivoli Storage Manager client GUI or from the **dsmc** command line.

- **From the Tivoli Storage Manager client GUI:**
  - Open the **File** menu and select **Quit**.
  - Press Command+Q.
  - Open the **File** menu and select **Exit**.
  - Open the **System** menu and select **Close**.
  - For the Web client: Open a different URL or close the browser.
- **From the DSMC command line:**
  - In batch mode, each **dsmc** command you enter is a complete session. Tivoli Storage Manager ends the session when it finishes processing the command.

- To end an interactive session, enter **quit** at the `tsm>` prompt.
- To interrupt a **dsmc** command before Tivoli Storage Manager has finished processing, enter **QQ** on the Tivoli Storage Manager console. In many cases but not all, this interrupts the command. If the command cannot be interrupted, use the UNIX **kill -9** command from an available command prompt. Do not press Ctrl-C because, while it ends the session, it can lead to unexpected results.

Related reference

"Loop" on page 498

Online forums

To participate in user discussions of Tivoli Storage Manager you can subscribe to the ADSM-L list server.

This is a user forum maintained by Marist College. While not officially supported by IBM, Tivoli Storage Manager developers and other IBM support staff also participate on an informal, best-effort basis. Because this is not an official IBM support channel, you should contact IBM Technical Support if you require a response specifically from IBM. Otherwise there is no guarantee that IBM will respond to your question on the list server.

You can subscribe by sending a note to the following e-mail address:

`listserv@vm.marist.edu`

The body of the message must contain the following:

`SUBSCRIBE ADSM-L yourfirstname yourlastname`

The list server will send you a response asking you to confirm the subscription request. Once you confirm your subscription request, the list server will send you further instructions. You will then be able to post messages to the list server by sending e-mail to:

`ADSM-L@vm.marist.edu`

If at a later time you want to unsubscribe from ADSM-L, you can send a note to the following e-mail address:

`listserv@vm.marist.edu`

The body of the message must contain the following:

`SIGNOFF ADSM-L`

You can also read and search the ADSM-L archives, join discussion forums, and access other resources at the following URL:

`http://www.adsm.org`

Other sources of online help

An anonymous FTP server (ftp://public.dhe.ibm.com/storage) is available where you can find maintenance and other Tivoli Storage Manager-related materials.

Three other anonymous servers are unofficially maintained by non-IBM volunteers.

The following are the anonymous servers:

```
ftp.rz.uni-karlsruhe.de (mirror - Germany)
ftp.wu-wien.ac.at (mirror - Austria)
ftp.cac.psu.edu (mirror - Pennsylvania)
```

You can get maintenance information from the Tivoli Storage Manager support page at: http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/ Tivoli_Storage_Manager

Chapter 4. Backing up your data

Use the backup-archive client to store backup versions of your files on the Tivoli Storage Manager server. You can restore these backup versions if the original files are lost or damaged. This section discusses various ways you can back up your data.

All client backup and restore procedures in this topic also apply to the Web client, except the following:
- Estimate
- Preferences editor

The following is a list of primary backup tasks.
- "Planning your backups"
- "Pre-backup considerations (UNIX and Linux)" on page 121
- "Performing an incremental, selective, or incremental-by-date backup (UNIX and Linux)" on page 131
- "Deleting backup data" on page 142
- "Backing up files from one or more file spaces for a group backup (UNIX and Linux)" on page 143
- "Image backup" on page 149
- "Back up NAS file systems using Network Data Management Protocol" on page 157

Planning your backups

If you are a first-time user, or if you only back up files occasionally, you can use the table in this topic as a checklist of preliminary steps to consider before backing up data.

Read the list of tasks to determine whether you are ready to back up your data.
- Decide whether you want to back up files or archive them. See "When to back up and when to archive files" on page 120 for more information.
- See "Pre-backup considerations (UNIX and Linux)" on page 121 for important considerations before you back up your files and directories.
- Do you need to exclude files from backup services? See "Include-exclude options to control processing" on page 123 for more information.

Understanding which files are backed up

When you request a backup, Tivoli Storage Manager backs up a file if some requirements are met.

The following are the requirements that must be met for Tivoli Storage Manager to back up a file.
- The selected management class contains a backup copy group.
- The file meets the serialization requirements defined in the backup copy group. If serialization is *static* or *shared static*, and the file changes during backup, the file will not be backed up.

- The file meets the mode requirements defined in the backup copy group. If the mode is *modified*, the file must have changed since the last backup. If the mode is *absolute*, the file can be backed up even if it does not change.
- The file meets the frequency requirements defined in the backup copy group. The specified minimum number of days since the last backup must elapse before a file is backed up.
- The file is not excluded from backup by an exclude statement.
- The file is not excluded from backup by the operating system. These excluded files can be found in registry subkey HKEY_LOCAL_MACHINE\SYSTEM\ CurrentControlSet\Control\BackupRestore\FilesNotToBackup.

Files that are part of the Windows system state (system objects on Windows XP) are eligible for backup only when performing a system state backup. You can back up the system state only as a single entity because of dependencies among the system state components. You can not back up or restore the files individually. For example, because `C:\windows\system32\ntoskrnl.exe` is part of the Windows system state, it is not backed up during an incremental or selective backup of the `C:` drive.

Related concepts

Chapter 9, "Storage management policies," on page 229

"Management classes and copy groups" on page 230

When to back up and when to archive files

When the backup-archive client backs up or archives a file, it sends a copy of the file and its associated attributes to the server; however, backups and archives have different goals.

Use backups to protect against unforeseen damage to your files, and use archives for maintaining more permanent versions of your files.

Backup data is managed by version, using predetermined policy based rules. Using these rules, the Tivoli Storage Manager administrator can control the following:
- The number of versions
- The number of days each additional backup copy is kept
- What happens to backup data versions when the file is deleted on the client system

Each copy of the file stored on the server is considered to be a separate and unique version of the file.

Archive is a powerful and extremely flexible mechanism for storing long term data. Archive data is kept for a specified number of days. Archive has no concept or support for versions. The user or Tivoli Storage Manager administrator is responsible for determining what files get added to an archive.

Tip: If archive is run multiple times against a file using the same archive description, a new copy of the file is added to the archive each time that archive is run. To simplify retrieve, store only one copy of a file in each archive.

Backups protect against file damage or loss that could occur through accidental deletion, corruption, disk crashes, and so forth. The server maintains one or more

backup versions for each file that you back up. Older versions are deleted as newer versions are made. The number of backup versions the server maintains is set by your administrator.

Archive copies are saved for long-term storage. Your administrator can limit how long archive copies are kept. The server can store an unlimited number of archive versions of a file. Archives are useful if you need to go back to a particular version of your files, or you want to delete a file from your workstation and retrieve it at a later time, if necessary. For example, you might need to save spreadsheets for tax purposes, but because you are not using them, you do not want to leave them on your workstation.

Related concepts

Pre-backup considerations (UNIX and Linux)

This section discusses some items to consider before you back up your data.

LAN-free data movement

LAN-free data movement shifts the movement of client data from the communications network to a storage area network (SAN). This decreases the load on the Tivoli Storage Manager server.

The SAN provides a path that allows you to back up, restore, archive, and retrieve data to and from a SAN-attached storage device. Client data moves over the SAN to the storage device using the Tivoli Storage Manager Storage Agent. The Tivoli Storage Manager Storage Agent must be installed on the same system as the client.

AIX, HP-UX, Linux and Solaris clients support LAN-free data movement.

LAN-free prerequisites

To enable LAN-free support, you must install and configure the Tivoli Storage Manager Managed System for SAN Storage Agent on the client workstation.

The Tivoli Storage Manager Managed System for SAN Storage Agent is a separate product and it is included on its own DVD, with the Tivoli Storage Manager for SAN product.

For more information about installing and configuring the Tivoli Storage Manager Managed System for SAN Storage Agent, refer to the IBM Tivoli Storage Manager Storage Agent User's Guides.

LAN-free data movement options

This topic describes the options that you can use to enable LAN-free data movement, after you have installed and configured the Tivoli Storage Manager Managed System for SAN feature on the client workstation.

The following are the options that you can use to enable LAN-free data movement:

enablelanfree
> Specifies whether to enable an available LAN-free path to a SAN-attached storage device.

lanfreecommmethod
> Specifies a communication protocol between the client and the Storage Agent.

lanfreeshmport
> Specifies the unique number that is used by the client and the storage agent to identify shared memory area used for communications.

lanfreetcpport
> Specifies the TCP/IP port number where the Storage Agent is listening.

lanfreetcpserveraddress
> Specifies the TCP/IP address for the Tivoli Storage Manager storage agent.

Related reference

"Enablelanfree" on page 314

"Lanfreecommmethod" on page 358

"Lanfreeshmport" on page 360

"Lanfreetcpport" on page 361

"Lanfreetcpserveraddress" on page 361

Incremental backups on memory-constrained systems

Incremental backup performance suffers if the system has a low amount of memory available before starting the backup.

If your system is memory constrained, specify the memoryefficientbackup *yes* option in your client options file. This option causes Tivoli Storage Manager to process only one directory at a time, which reduces memory consumption but increases backup time. When you specify *yes*, Tivoli Storage Manager analyzes only one directory at a time for backup consideration. If performance remains poor, check your communication buffer settings and the communication link between your system and the Tivoli Storage Manager server. If your system is not memory constrained, setting the memoryefficientbackup option to *yes* degrades your backup performance.

Related reference

"Memoryefficientbackup" on page 368

Incremental backups on systems with a large number of files

The client can use very large amounts of memory to perform incremental backup operations, especially on file systems that contain large numbers of files.

The term *memory* as used here is the addressable memory available to the client process. Addressable memory is a combination of physical RAM and virtual memory.

On average, the client uses approximately 300 bytes of memory per object (file or directory). Thus for a file system with one million files and directories, the Tivoli Storage Manager client requires, on average, approximately 300 MB of memory. The exact amount of memory used per object varies, depending on the length of the object path and name length, or the nesting depth of directories. The number of bytes of data is not an important factor in determining the Tivoli Storage Manager backup client memory requirement.

The maximum number of files can be determined by dividing the maximum amount of memory available to a process by the average amount of memory needed per object.

The total memory requirement can be reduced by any of the following methods:

- Use the client option memoryefficientbackup *diskcachemethod*. This choice reduces the use of memory to a minimum at the expense of performance and a significant increase in disk space required for the backup. The file description data from the server is stored in a disk-resident temporary database, not in memory. As directories on the workstation are scanned, the database is consulted to determine whether to back up, update, or expire each object. At the completion of the backup, the database file is deleted.
- Use the client option memoryefficientbackup *yes*. The average memory used by the client then becomes 300 bytes times the number of directories plus 300 bytes per file in the directory being processed. Note that for file systems with large numbers (millions) of directories, the client still might not be able to allocate enough memory to perform incremental backup with memoryefficientbackup *yes*.
- UNIX and Linux clients might be able to use the virtualmountpoint client option to define multiple virtual mount points within a single file system, each of which can be backed up independently by the Tivoli Storage Manager client.
- If the client option resourceutilization is set to a value greater than 4, and there are multiple file systems being backed up, then reducing resourceutilization to 4 or lower limits the process to incremental backup of a single file system at a time. This reduces the memory requirement. If backup of multiple file systems in parallel is required for performance reasons, and the combined memory requirements exceed the process limits, then multiple instances of the backup client can be used to back up multiple file systems in parallel. For example, if you want to back up two file systems at the same time but their memory requirements exceed the limits of a single process, then start one instance of the client to back up one of the file systems, and start a second instance of the client to back up the other file system.
- Use the - incrbydate client option to perform an "incremental-by-date" backup.
- Use the Tivoli Storage Manager client exclude.dir option to prevent the client from traversing and backing up directories that do not need to be backed up.
- Except for Mac OS X, use the client image backup function to back up the entire volume. This might actually use less system resources and run faster than incremental backup of some file systems with a large number of small files.
- Reduce the number of files per file system by spreading the data across multiple file systems.

Related reference

Include-exclude options to control processing

You might have files in your file systems that you do not want to back up. These files might be core files, local caches of network file systems, operating system or application files that could be easily recovered by reinstalling the program, or any other files that you could easily rebuild.

An Authorized User on your workstation can use the exclude and include options in your include-exclude options list to specify which files to exclude from backup processing.

Use the include and exclude options in dsm.sys to define which files to include or exclude from incremental or selective backup processing. A file is eligible for backup unless excluded by an exclude option. It is not necessary to use an include option to include specific files for backup unless those files are in a directory containing other files you want to exclude.

Tivoli Storage Manager uses *management classes* to determine how to manage your backups on the server. Every time you back up a file, the file is assigned a management class. The management class is either a default chosen for you, or one you assign to the file using the include option in the include-exclude list. If you assign a management class, it must contain a backup copy group for the file to be backed up.

Related tasks

"Creating an include-exclude list" on page 93

"Set the client scheduler process to run as a background task and start automatically at startup" on page 218

Data encryption during backup or archive operations

For the strongest possible encryption, use 128-bit Advanced Encryption Standard (AES) data encryption, with the encryptiontype option.

The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received.

Important: The include.encrypt option is the only way to enable encryption on the backup-archive client. If no include.encrypt statements are used encryption will not occur.

Use the include and exclude options in dsm.sys to define which files to include or exclude from incremental or selective backup processing. A file is eligible for backup unless excluded by an exclude option. It is not necessary to use an include option to include specific files for backup unless those files are in a directory containing other files you want to exclude.

To encrypt file data, you must select an encryption key password, which Tivoli Storage Manager uses to generate the encryption key for encrypting and decrypting the file data. Store the encryption key password for later use. You can specify whether to save the encryption key password in a file named TSM.PWD by using the encryptkey option.

Tivoli Storage Manager client encryption allows you to enter a value of up to 63 characters in length. This encryption password needs to be confirmed when encrypting the file for backup, and also needs to be entered when performing restores of encrypted files.

While restoring the encrypted file, Tivoli Storage Manager prompts you for the key password to decrypt the file in the following cases:
- If the encryptkey option is set to Prompt.
- If the key supplied by the user in the above case does not match.
- If the encryptkey option is set to Save and the locally saved key password does not match the encrypted file.

File system and ACL support

Special file systems contain dynamic information generated by the operating system; they contain no data or files. The UNIX and Linux clients ignore special file systems and their contents.

Special file systems include the following types:
- the /proc file system on most of the UNIX platforms
- the /dev/fd file system on Solaris
- the /dev/pts on Linux

The Tivoli Storage Manager client can work on specific file system types that are commonly used. Table 50 contains a list of supported file system types.

Table 50. Supported file systems and ACL support

Platform	File System	ACL Support
AIX		
	GPFS	yes
	JFS	yes
	JFS2	yes
	JFS2 NFSV4	yes
	VxFS	yes
HP-UX		
	HFS	no
	VxFS (JFS Veritas)	yes (V3.3.Layout 4)
Linux x86/x86_64		
	XFS	yes
	EXT2	yes
	EXT3	yes
	ReiserFS	yes
	GPFS	yes
	JFS	no
	VxFS	no
	NSS	yes
Linux on POWER	XFS	yes
	EXT2	yes
	EXT3	yes
	ReiserFS	yes
	JFS	no
	GPFS	no
Linux on System z		
	EXT2	yes
	EXT3	yes
	ReiserFS	yes
	JFS	no
	GPFS	yes

Table 50. Supported file systems and ACL support (continued)

Platform	File System	ACL Support
Mac	HFS Standard (HFS)	yes
	HFS Extended (HFS+)	yes
	HFS Extended case-sensitive (HFSX)	yes
	Xsan (XSAN)	yes
	UNIX file system (UFS)	yes
	Universal disk format (UDF)	yes
	ISO9660	yes
Solaris		
	UFS	yes
	VxFS	yes
	QFS	no
	ZFS	yes

With file systems where NFS V4 ACLs are defined and used (Solaris ZFS and AIX JFS2 V2), even if only the standard UNIX permissions or ACLs have changed (such as with the CHMOD command), the file or directory is fully backed up again. With other file systems, this type of change causes only an attribute update on the Tivoli Storage Manager server.

To process all other file systems, use the virtualmountpoint option to enable support for the following items:
- To back up, restore, archive, and retrieve file data
- For basic UNIX and Linux permissions
- For change/access/modification time stamps, and the directory tree structure

No other file-system specific attributes, such as the ACL, are valid. The file system type for such file systems is set to "UNKNOWN".

For example, if the /media/abc/DATA1 file system is not supported by Tivoli Storage Manager, add the following statement to dsm.sys to back up or archive the data in this file system:

```
VIRTUALMOUNTPOINT /media/abc/DATA1
```

This support is only available if the file system can use basic POSIX system calls, such as read or write processing on your system.

Cross-platform backup and restore is not supported. For example, data backed up by an AIX client is not available for restore by a Windows client and vice versa.

Note: Data backed up or archived by the Mac OS X client cannot be restored by any other client. Additionally, the Mac OS X client cannot restore or retrieve data from any other client.

You can use the cross-file system type restore or retrieve method for ACL information if both the original file system and the destination file system support compatible ACLs. For example, on Solaris, the ACL information that is backed up from a VxFS file system is restored to a UFS file system because these file systems support compatible ACLs. The ACL information is not restored during cross-file system restore or retrieve operations if the original file system and the destination file system do not support ACLs,

The standalone package LSCqfs 3.5.0 is the only supported version of QFS. In addition, the following restrictions also apply to the QFS file system:

- Image backup is not supported on QFS file systems.
- The Solaris backup-archive client does not support the combination of QFS and SAM needed to archive files onto tertiary background storage, such as tapes. Instead, it recalls files from tape to disk automatically if it finds migrated files during a backup.
- A QFS file system contains two hidden system files and a system directory that cannot be backed up. This is acceptable because a backup of these files is not needed. They contain internal data to manage the file system. This data is automatically excluded from a backup and recreated automatically by the file system itself if a restore of files in that file system is invoked.

Incremental, selective, filelist back up, archive, restore, and retrieve processing of the Veritas file system and its ACLs on AIX are supported. Restore of a Veritas volume on a Logical Volume Manager volume and vice-versa is allowed provided both have the same file system type.

This information pertains to Mac OS X systems only:
- On Mac OS X systems, the UFS and HFSX file systems are case-sensitive whereas the HFS+ file system is case-insensitive but is case-preserving. Files that you back up from a UFS or HFSX file system (case-sensitive) might not be restored properly to an HFS+ file system (case-insensitive) file system. For example, on a UFS file system, files `Afile` and `afile` are seen as different files. However, on a HFS+ file system the two files are seen as identical.
- If Case-sensitive HFS+ or UFS file systems are used, it is very important that the data from the HFSX or UFS file system is not backed up to a HFS+ file system on the Tivoli Storage Manage server. Either a new name must be used on the system or the existing file space on the Tivoli Storage Manager server must be renamed. For example, consider a system that has a file system named `/Volumes/fs2` and this system is repartitioned with a Case-Sensitive HFS+ file system. Either the `/Volumes/fs2` file system on the Tivoli Storage Manager server needs to be renamed, or a new name must be used on the local system. If this is not done, Tivoli Storage Manager mixes the HFSX case-sensitive data with the HFS+ case-insensitive data that is already stored on the Tivoli Storage Manager server.
- Aliases and symbolic links are backed up. However, Tivoli Storage Manager does not back up the data to which the symbolic links point.
- When files that are backed up from an HFS volume are restored to a UFS volume, the resource forks will not have the correct owner. Correct this problem by using the **chown** command on the resource fork file to change the owner. The resource fork file stores structured data in a file.

On LinuxPPX, Linux390 and LinuxAMD, you must install libacl.so for the Tivoli Storage Manager client to back up ACLs.

Important: If you are running GPFS for AIX or GPFS for Linux x86/x86_64 in a multinode cluster, and all nodes share a mounted GPFS file system, Tivoli Storage Manager processes this file system as a local file system. Tivoli Storage Manager backs up the file system on each node during an incremental backup. To avoid this, you can do *one* of the following:
- Explicitly configure the domain statement in the client user-options file (dsm.opt) to list the file systems you want that node to back up.
- Set the exclude.fs option in dsm.sys to exclude the GPFS file system from backup services.

Maximum file size for operations

The maximum file size depends on the type of a file system. The Tivoli Storage Manager client does not check any file size limit during backup, archive, restore, or retrieve operations.

If the file system allows creation of the file, the Tivoli Storage Manager client backs up or archives the file.

The following table specifies the maximum file sizes for the native file systems on Tivoli Storage Manager UNIX and Linux client platforms.

Table 51. Maximum file size

Platform	Max file size (in bytes)
AIX 5.3 (GPFS)	109 951 162 777 600 (100 TB)
AIX 5.3 (JFS2)	1 099 511 627 264 (1 TB - 512)
AIX 6.1 (JFS2) size limitations	Maximum JFS2 file system size: 32 TB Maximum JFS2 file size: 16 TB Minimum JFS2 file system size: 16 MB
HP-UX	1 099 511 627 775 (1 TB-1)
All Linux clients	9 223 372 036 854 775 807 (8 EB-1)
Mac OS X	HFS - 2 147 485 648 (2GB) HFS+, HFSX, XSAN, and UFS - 9 223 372 036 854 775 808 (8EB)
Solaris	1 099 511 627 775 (1 TB-1)
Solaris (ZFS)	18 446 744 073 709 551 616 (16 EB)

Long user and group names

Tivoli Storage Manager can handle user and group names that are up to 64 characters without any issues. However, names longer than 64 characters require special handling by Tivoli Storage Manager.

Important: Do not exceed the 64 character limit for user and group names. If you do, Tivoli Storage Manager shortens the name to fall within this limit by using the following transformation: Take the first 53 characters, append a forward slash (/), and then the numeric ID as a character string.

An error message is logged that contains both the long name and the resulting shortened string. For most functions, you do not need to be aware of the shortened name. The exceptions are:

- The **set access** command
- The fromowner option
- The users and groups (authorization) options

In each of these cases, when you need to enter a name, you either have to find the error message containing the transformation, or construct the name using the above outlined rule.

Mac OS X volume names

Tivoli Storage Manager backs up volumes based on their UNIX mount point name.

Tivoli Storage Manager maintains each volume name as a separate restore or retrieve volume. These volume names become the names of file spaces on the server.

If you change the name of a volume you have already backed up, Tivoli Storage Manager sees it as a new volume and does not relate it to the previous one. Any backup of the volume backs up the files under the new name. A mismatch might occur if you rename your volumes, or if you access Tivoli Storage Manager from a different workstation than the one from which you backed up the files.

Mac OS X volume naming precautions

Tivoli Storage Manager creates all new Tivoli Storage Manager file spaces on the server with the UNIX mount point of the volume.

If there are two volumes with the names such as `"La Pomme"` and `"la pomme"`, two unique UNIX mount points are created.

The following examples show the two mount points that are created:

```
/Volumes/La Pomme
/Volumes/la pomme
```

If duplicate volumes exist on your desktop, it is possible for the UNIX mount points to be different than the last time Tivoli Storage Manager did a backup. Tivoli Storage Manager might not back up the data to the correct file system on the Tivoli Storage Manager server.

You can check the file system where Tivoli Storage Manager backs up the data:
1. In the TSM Backup window, select a file system.
2. Click **File → Show Info**.

The UNIX mount point is at the top of the Information dialog.

The best way to avoid any potential naming problems is to ensure that the volume names are unique.

Important:
- The Tivoli Storage Manager client continues to use the existing file space names on the Tivoli Storage Manager Server. Only new file spaces use the UNIX mount point for the name.
- Do not specify volumes with periods in the name (...). Tivoli Storage Manager uses the sequence of periods as part of include-exclude processing. Tivoli Storage Manager reports an invalid include-exclude statement if a volume has a sequence of periods in the name. The volume *must* be renamed.

Mac OS X volume naming precautions on dual boot systems

If you have more than one version of Mac OS X that you switch between, it is critical that you understand how Tivoli Storage Manager uses the UNIX mount paths for file space names on the Tivoli Storage Manager server.

For example, consider a dual-boot system that has two volumes, `Panther` and `Tiger`. The finder and the Tivoli Storage Manager GUI displays these as `Panther` and `Tiger`. However, the UNIX mount points depend upon which version of Mac OS is running. If Panther is the startup disk, the UNIX paths are:

```
/
/Volumes/Tiger
```

If Tiger is the startup disk, the UNIX paths are:

```
/
/Volumes/Panther
```

When Tivoli Storage Manager does a backup or archive, the file space names also depend on which version of Mac OS X is running.

Both versions of Mac OS X back up to the / file system on the Tivoli Storage Manager server. When this happens, the system files are intermixed.

To avoid potential problems on dual-boot systems, complete *one* of these tasks:

1. Select one version of Mac OS X on which to install and run Tivoli Storage Manager. This ensures that the UNIX mount points are the same each time Tivoli Storage Manager does a backup.
2. Configure each version of Mac OS X with a unique Tivoli Storage Manager node name. Then exclude the other version of Mac OS X from backup processing with a domain statement in the system options file. For example, if the volume Tiger is the startup disk, add this option to the system options file:

   ```
   DOMAIN -/Volumes/Panther
   ```

   If the volume Panther is the startup disk, add this to the system options file:

   ```
   DOMAIN -/Volumes/Tiger
   ```

Mac OS X Unicode enablement

The Mac OS X client is Unicode enabled. New clients storing data on the server for the first time require no special set up.

The server automatically stores files and directories as Unicode enabled. However, if you are upgrading to the Unicode-enabled client, you need to plan the migration of existing file spaces so they can support Unicode.

Any file spaces that are already on the server must be renamed so Unicode-enabled file spaces can be created. Use the autofsrename option rename existing file spaces.

Related reference

"Autofsrename" on page 274

Mac OS X FileVault home directories

The Mac OS X FileVault feature compresses your home directory into a single file. Tivoli Storage Manager cannot access any user files inside a Mac OS X FileVault home directory.

Tivoli Storage Manager can back up your home directory as a single file. Because of the manner in which FileVault home directories are managed by Mac OS X, back up any Mac OS X FileVault home directories when you are not logged in. When you are logged in, the home directory is mounted in a temporary folder in /Users/.<username>, and the sparse file is backed up again under this temporary folder.

Tivoli Storage Manager can restore your home directory if you are not logged in. This means that you cannot restore your own home directory. Your home directories can be restored by the system administrator.

You must have a primary and secondary system administrator account so that the primary system administrator account can be restored if needed.

Mac OS X Time Machine backup disk

Time Machine is the backup application available with Mac OS X.

Tivoli Storage Manager can be used at the same time as Mac OS X Time Machine application. However, due to the unique nature of how the Mac OS X Time Machine application backs up data, consider the following items before using Tivoli Storage Manager to back up the Mac OS X Time Machine data:

- The Mac OS X Time Machine backup disk makes extensive use of both file and directory hard links to minimize disk usage. For example, if the disk backed up with the Mac OS X Time Machine application is 5 GB, the first backup copies 5 GBs of data to the Mac OS X Time Machine backup disk.

  Subsequent backups only copy the files that have changed since the previous backup. All files and directories that have not changed are hard-linked with the version that was copied during the previous backup.

  The Finder shows each backup as 5 GB, for a total size of 10 GB. However, because of the use of hard links, the total disk usage is only slightly larger than 5 GB.

  All hard-linked objects that are not already on the Tivoli Storage Manager server are backed up.

  For example, 10 GB of data would be sent to the Tivoli Storage Manager server.

- When files that are restored are hard-linked, Tivoli Storage Manager recreates the original hard link. Recreating the original hard link can only be done if *all* files that are hard-linked are restored at the same time. Restoring all the hard-linked files at the same time is not a practical method for a large backup disk that uses the Mac OS X Time Machine application.

- When the Mac OS X Time Machine application copies files to the backup disk, ACLs are added to the files to protect them from deletion. Tivoli Storage Manager can back up and restore files with ACLs. However, any files that are restored must have these restrictive ACLs in place.

Tip: For best results, exclude the Time Machine application backup data. All Time Machine application data is in a directory named `Backups.backupdb`.

Related concepts

"System files to exclude" on page 97

Performing an incremental, selective, or incremental-by-date backup (UNIX and Linux)

Your administrator might have set up schedules to automatically back up files on your workstation. The following sections discuss how to back up files without using a schedule.

There are two types of incremental backup: *full incremental* and *partial incremental*.

"Set the client scheduler process to run as a background task and start automatically at startup" on page 218

Full and partial incremental backup

An incremental backup backs up only new and changed files. Whether you perform a full or partial incremental backup depends on what you select for backup.

If you select entire file systems, you are performing a full incremental backup. If you select a directory tree or individual files, you are performing a partial incremental backup.

The first time you run a full incremental backup, Tivoli Storage Manager backs up all the files and directories that you specify. This process can take a long time if the number of files is large, or if one or more very large files must be backed up. Subsequent full incremental backups only back up new and changed files. This allows the backup server to maintain current versions of your files, without having to waste time or space by backing up files that already exist in server storage.

Depending on your storage management policies, the server might keep more than one version of your files in storage. The most recently backed up files are active backup versions. Older copies of your backed up files are inactive versions. However, if you delete a file from your workstation, the next full incremental backup causes the active backup version of the file to become inactive. If you need to restore a file you have deleted, and if a full incremental backup has been run since you deleted the file, then you need to restore an inactive version of the file (assuming that a version still exists on the server). The number of inactive versions maintained by the server and how long they are retained is governed by the management policies defined by your server administrator. The purpose of the active versions is to represent which files existed on your file system at the time of the backup.

To perform a full or partial incremental backup using the client Java GUI, select the incremental (complete) option from the pull-down menu at the top of the screen on the backup window, or use the **incremental** command in the command-line interface. Specify file systems, directory trees, or individual files to include in the backup.

During an incremental backup, the client queries the server or the journal database to determine the exact state of your files since the last incremental backup. The client uses this information to perform the following tasks:

- Back up new files.
- Back up files whose contents changed since the last backup.
- The client considers a file changed if any of the following attributes changed since the last backup:
  - File size
  - Date or time of last modification
  - Extended Attributes
  - Access Control List

  If *only* the following items change, they are updated without causing the entire file to be backed up to the server:
  - File owner
  - File permissions

- Last access time
- Inode
- Group ID
- Icon location (Mac OS X only)
- Type or creator (Mac OS X only)
- Expire backup versions of files on the server that do not have corresponding files on the workstation. The result is that files which no longer exist on your workstation will not have active backup versions on the server. However, inactive versions are retained according to rules defined by the Tivoli Storage Manager administrator.
- Rebind backup versions if management class assignments change. Only objects that have active backup versions are rebound. Objects for which only inactive backup versions exist are not rebound.

Important: Each directory is also backed up under the following circumstances:
- It has not been backed up previously
- Its permissions have changed since the last backup
- Its Access Control List has changed since the last backup
- Its Extended Attributes have changed since the last backup
- Its time stamp has changed since the last backup

You might be able to use the preservelastaccessdate option during a backup or archive operation to specify whether to reset the last access date to its original value following a backup or archive operation. By default, the Tivoli Storage Manager client *will not* reset the last access date of any backed up or archived files to their original value prior to the backup or archive operation.

Directories are counted in the number of objects backed up. To exclude directories and their contents from backup, use the exclude.dir option.

Related concepts

Chapter 9, "Storage management policies," on page 229

Related reference

"Exclude options" on page 321

"Preservelastaccessdate" on page 389

Journal-based backup on AIX

Journal-based backup is an alternate method of backup that uses a change journal maintained by the Tivoli Storage Manager journal daemon process.

Journal-based backup is supported on JFS, JFS2, VxFs, and HSM-managed JFS2 file systems and the HACMP environment (not supported on GPFS file systems).

To support journal-based backup you must install and configure the Tivoli Storage Manager daemon.

A backup for a particular file system will be journal-based when the Tivoli Storage Manager journal daemon has been installed and configured to journal the particular file system, and a valid journal has been established for the file system.

The primary difference between traditional incremental backup and journal-based backup is the method used for backup and expiration candidates.

Traditional incremental backup obtains the list of backup and expiration candidates by building comprehensive lists of local objects, and lists of active server objects

for the file system being backed up. The local lists are obtained by scanning the entire local file system. The server list is obtained by querying the entire server inventory for all active objects.

The two lists are compared, and candidates are selected according to the following criteria:
- An object is selected as a backup candidate if it exists in the local list. but doesn't exist in the server list. The object is also a backup candidate if it exists in both lists, but differs according to Tivoli Storage Manager incremental criteria (for example, attribute changes, date and size changes).
- An object is selected as an expiration candidate if it exists in the server list, but doesn't exist in the local list.

Journal-based backup obtains the candidates list of objects to backup and expire by querying the Tivoli Storage Manager journal daemon for the contents of the change journal of the file system being backed up.

Change journal entries are cleared (marked as free) after they have been processed by the backup client and committed on the Tivoli Storage Manager server.

Journal-based backup is activated by configuring the journal daemon to monitor specified file systems for change activity.

Journal-based backup is enabled by successfully completing a full incremental backup.

You can use journal-based backup when backing up file systems with small or moderate amounts of change activity between backup cycles. If you have many file changes between backup cycles, you will have very large change journals. Many changes to the journal-based backup file might pose memory and performance problems that can negate the benefits of journal-based backup. For example, creating, deleting, renaming, or moving very large directory trees can also negate the benefit of using journal-based backup instead of normal incremental backup.

Journal-based backup is not intended to be a complete replacement for traditional incremental backup. You should supplement journal-based backup with a full progressive incremental backup on a regular basis. For example, perform journal-based backups on a daily basis, and full incremental backups on a weekly basis.

Here are some limitations of journal-based backup:
- Individual server attributes are not available during a journal-based backup. Certain policy settings such as copy frequency and copy mode might not be enforced.
- Other platform-specific behaviors might prevent objects from being processed properly. Other software that changes the default behavior of the file system might prevent file system changes from being detected.
- If the file system is very active when a journal-based backup is in progress, it is possible that a small number of deleted files will not be expired.

You should perform periodic full incremental backups and more frequent journal backups. Traditional incremental backup compares the entire server inventory of files against the entire local file system. Therefore, incremental backup is always the most comprehensive backup method.

Note:

1. When using antivirus software, there are limitations to journal-based backup. The files that are scanned by the antivirus software could be backed up in the journal again in the next backup after the scan.

2. A journal-based backup might not fall back to the traditional incremental backup if the policy domain of your node is changed on the server. This depends on when the policy set within the domain was last updated and the date of the last incremental backup. In this case, you must force a full traditional incremental backup to rebind the files to the new domain. Use the nojournal option with the **incremental** command to specify that you want to perform a traditional full incremental backup, instead of the default journal-based backup.

Incremental-by-date backup

For a file system to be eligible for incremental-by-date backups, you must have performed at least one full incremental backup of that file system. Running an incremental backup of only a directory branch or individual file will not make the file system eligible for incremental-by-date backups.

The client backs up only those files whose modification date and time is later than the date and time of the last incremental backup of the file system on which the file resides. Files added by the client after the last incremental backup, but with a modification date earlier than the last incremental backup, are not backed up.

Files that were renamed after the last incremental backup, but otherwise remain unchanged, will not be backed up. Renaming a file does not change the modification date and time of the file. However, renaming a file does change the modification date of the directory in which it is located. In this case, the directory is backed up, but not the files it contains.

If you run an incremental-by-date backup of the whole file system, the server updates the date and time of the last incremental backup. If you perform an incremental-by-date backup on only part of a file system, the server does not update the date of the last full incremental backup. In this case, the next incremental-by-date backup backs up these files again.

Note: Unlike incremental backups, incremental-by-date backups do not expire deleted files or rebind backup versions to a new management class if you change the management class.

Related tasks

"Backing up data using the Java GUI" on page 138

Comparing incremental-by-date, journal-based, and NetApp snapshot difference to full incremental and partial incremental backups

Incremental-by-date, journal-based, and NetApp snapshot difference are alternatives to full incremental and partial incremental back methods.

Incremental-by-date backup

An incremental-by-date backup takes less time to process than a full incremental backup and requires less memory.

An incremental-by-date backup might not place exactly the same backup files into server storage because the incremental-by-date backup:

- Does not expire backup versions of files that you delete from the workstation.
- Does not rebind backup versions to a new management class if you change the management class.
- Does not back up files with attributes that change, unless the modification dates and times also change.
- Ignores the copy group frequency attribute of management classes (Journal-based backups also ignore this attribute).

Journal-based backup

Journal-based backup only applies to AIX. The memory requirements for an initial journaling environment are the same as the memory requirements for a full file space incremental, because journal-based backups must complete the full file space incremental in order to set the journal database as valid, and to establish the baseline for journaling.

Journal-based backup has the following limitations:
- Individual server attributes are not available during a journal-based backup. Certain policy settings such as copy frequency and copy mode might not be enforced.
- Other platform-specific behaviors might prevent objects from being processed properly. Other software that changes the default behavior of the file system might prevent file system changes from being detected.
- If the file system is very active when a journal-based backup is in progress, it is possible that a small number of deleted files will not be expired.

The memory requirements for subsequent journal-based backups are much less. Journal backup sessions run in parallel and are governed by the resourceutilization client option in the same manner as normal backup sessions. The size of the journal database file reverts to a minimal size (less than 1 K) when the last entry has been deleted from the journal. Since entries are deleted from the journal as they are processed by the client, the disk size occupied by the journal should be minimal after a complete journal backup. A full incremental backup with journaling active takes less time to process than an incremental-by-date backup.

NetApp snapshot difference

An incremental-by-date backup might not place exactly the same backup files into server storage because the incremental-by-date backup:
- Does not expire backup versions of files that you delete from the workstation.
- Does not rebind backup versions to a new management class if you change the management class.
- Does not back up files with attributes that change, unless the modification dates and times also change.
- Ignores the copy group frequency attribute of management classes (Journal-based backups also ignore this attribute).

For NAS and N-Series file servers that are running ONTAP 7.3.0, or later, you can use the snapdiff option to invoke the snapshot difference backup from NetApp when running a full-volume incremental backup. Using this option reduces memory usage and is faster.

Consider the following restrictions when running a full-volume incremental backup using the snapdiff option, to ensure that data is backed up when it should be.

- A file is excluded due to an exclude rule in the include-exclude file. Tivoli Storage Manager runs a backup of the current snapshot with that exclude rule in effect. This happens when you have not made changes to the file, but you have removed the rule that excluded the file. NetApp will not detect this include-exclude change because it only detects file changes between two snapshots.
- If you added an include statement to the option file, that include option does not take effect unless NetApp detects that the file has changed. Tivoli Storage Manager does not inspect every file on the volume during backup.
- If you used the **dsmc delete backup** command to explicitly delete a file from the Tivoli Storage Manager inventory, NetApp cannot detect that a file was manually deleted from Tivoli Storage Manager. Therefore, the file remains unprotected in Tivoli Storage Manager storage until it is changed on the volume and the change is detected by NetApp, signalling Tivoli Storage Manager to back it up again.
- Policy changes such as changing the policy from *mode=modified* to *mode=absolute* are not detected.
- The entire file space is deleted from the Tivoli Storage Manager inventory. This action causes the snapdiff option to create a new snapshot to use as the source, and a full incremental backup to be run.

The NetApp software determines what is a changed object, not Tivoli Storage Manager.
If you run a full volume backup of an NFS mounted, NetApp or N-Series volume, all the snapshots under the .snapshot directory might also be backed up. To avoid this situation, you can do one of the following actions:

- Run NDMP backups
- Run backups using the snapshotroot option
- Run incremental backups using the snapdiff option

  Tip: If you run an incremental backup using the snapdiff option and you schedule periodic incremental backups, use the createnewbase=yes option with the snapdiff option to create a base snapshot and use it as a source to run an incremental backup.

- Exclude the .snapshot directory from backups

  Note: The .snapshot directory is not backed up for some versions of Red Hat Linux so you are not required to exclude it.

Selective backup

Use a selective backup when you want to back up specific files or directories regardless of whether a current copy of those files exists on the server.

Incremental backups are generally part of an automated system to back up entire file systems. In contrast, selective backups allow you to manually select a set of files to back up regardless of whether they have changed since your last incremental backup.

Unlike incremental backups, a selective backup provides the following:
- Does not cause the server to update the date and time of the last incremental.

- Backs up directory and file entries even if their size, modification timestamp, or permissions have not changed.
- Does not expire deleted files.
- Does not rebind backup versions to a new management class if you change the management class.

Related tasks

"Backing up data using the Java GUI"

Related reference

"Selective" on page 548

Solaris global zone and non-global zones backup

For Solaris zones, the incremental and selective backup of file systems should be performed where the file systems are non-globally mounted.

Each non-global zone should be treated as a separate system with its own Tivoli Storage Manager node name and the backup should be performed within each of the non-global zones.

To perform an incremental or selective backup of the non-global zones from the global zone, the administrator needs to consider which files from the zone directories should be included or excluded in the backup. For example, system files and kernel files would not automatically be excluded.

Saving access permissions

When you back up your files, Tivoli Storage Manager also saves standard UNIX access permissions assigned to the files.

Depending on your operating system, it also saves extended permissions. For example, for files on an AIX workstation, Tivoli Storage Manager saves access control lists.

It is possible for an authorized user to back up files for another user, but this should not cause ownership conflicts. The backup server properly records that the file belongs to the original owner. The authorized user does not need to grant the original owner access to the backup versions.

Setting a virtual mount point

If you are an authorized user and you want to back up files beginning with a specific directory within a file system, you can define that directory as a virtual mount point.

Defining a virtual mount point within a file system provides a direct path to the files you want to back up, saving processing time. It is more efficient than defining the file system with the domain option and then using an exclude option to exclude the files you do not want to back up. It also allows you to store backups and archives for specific directories in separate storage file spaces.

Related reference

"Virtualmountpoint" on page 452

Backing up data using the Java GUI

You can back up specific files, entire directories, or entire file systems from the directory tree.

You can locate the files you want to back up by searching or filtering. Filtering displays only the files matching the filter criteria for your backup.

Use the Tivoli Storage Manager client Java GUI to back up your data as follows:
1. Click **Backup** in the Tivoli Storage Manager window. The Backup window appears.
2. Expand the directory tree if necessary. Click on the selection boxes next to the object or objects you want to back up. To search or filter files, click the **Find** icon on the tool bar.
3. Enter your search criteria in the Find Files (Backup) window.
4. Click the **Search** button. The Matching Files (Backup) window appears.
5. Click the selection boxes next to the files you want to back up and close the Matching Files (Backup) window.
6. Enter your filter criteria in the Find Files (Backup) window.
7. Click the **Filter** button. The Backup window displays the filtered files.
8. Click the selection boxes next to the filtered files or directories you want to back up.
9. Select one of the following backup types from the pull-down menu: (1) To run an incremental backup, click **Incremental (complete)**, (2) To run an incremental-by-date backup, click **Incremental (date only)**, (3) To run a selective backup, click **Always backup**.
10. Click **Backup**. The Backup **Task List** window displays the backup processing status.

Consider the following items when you back up your data using the Java GUI.
- To modify specific backup options, click the **Options** button. The options you select are effective during the current session *only*.
- Tivoli Storage Manager uses management classes to determine how to manage your backups on the server. Every time you back up a file, the file is assigned a management class. The management class used is either a default selected for you, or one that you assign to the file using an *include* option in the include-exclude options list. Select **Utilities** → **View Policy Information** from the backup-archive client Java GUI or Web client GUI to view the backup policies defined by the Tivoli Storage Manager server for your client node.
- To perform an automatic incremental backup of your default domain, select **Actions** → **Backup Domain**. Your default domain is set with the *domain* option in your client user-options file (dsm.opt). If you do not have the *domain* option set, the default domain is *all local file systems*.
- You can use the Preferences editor to exclude file systems in your default domain from backup processing.

Related concepts

Chapter 9, "Storage management policies," on page 229

Related reference

"Domain" on page 303

Backing up data using the command line

You can use the **incremental** or **selective** commands to perform backups.

The following table shows examples of using these commands to perform different tasks.

Table 52. Command-line backup examples

Task	Command	Considerations
Incremental backups		
Perform an incremental backup of your client domain.	`dsmc incremental`	See "Incremental" on page 492 for more information about the **incremental** command.
Back up the /fs1 and /fs2 file systems in addition to the /home, /usr, and /datasave file systems defined in your client domain.	`dsmc incremental -domain="/fs1 /fs2"`	See "Domain" on page 303 for more information about the domain option.
Back up the /Volumes/fs1 and /Volumes/fs2 file systems in addition to the volumes defined in your client domain.	`dsmc incremental -domain="/Volumes/fs1 /Volumes/fs2"`	See "Domain" on page 303 for more information about the domain option.
Back up all local file systems defined in your client domain *except* for the /home file system.	`dsmc incremental -domain="all-local -/home"`	You cannot use the (-) operator in front of the domain keyword all-local. See "Domain" on page 303 for more information. For Windows Server 2003, Windows Server 2008, Windows Vista, and Windows 7 you can also exclude the systemstate domain from backup processing in this way.
Back up *only* the /fs1 and /fs2 file systems.	`dsmc incremental /fs1 /fs2`	None
Back up all files in the /home directory and all its subdirectories.	`dsmc incremental /home/ -subdir=yes`	See "Subdir" on page 432 for more information about the subdir option.
Back up all files in the /Users directory and all its subdirectories.	`dsmc incremental /Users/ -subdir=yes`	See "Subdir" on page 432 for more information about the subdir option.
Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, run an incremental backup of all files and directories under the local snapshot and manage them on the Tivoli Storage Manager server under the file space name /usr.	`dsmc incremental /usr -snapshotroot=/snapshot/day1`	Tivoli Storage Manager considers the snapshotroot value as a file space name. See "Snapshotroot" on page 425 for more information.
Incremental-by-date backup		
Perform an incremental-by-date backup of your default client domain.	`dsmc incremental -incrbydate`	Use the incrbydate option with the **incremental** command to back up new and changed files with a modification date later than the last incremental backup stored at the server. See "Incrbydate" on page 357 for more information about the incrbydate option.
Selective backups		

Table 52. Command-line backup examples (continued)

Task	Command	Considerations
Back up all files in the /home/proj or /Users/van/Documents directory.	dsmc selective /home/proj/ or dsmc selective /Users/van/Documents/	Use the **selective** command to back up specific files or directories regardless of whether they have changed since your last incremental backup. You can use wildcards to back up multiple files at once. See "Selective" on page 548 for more information about the **selective** command.
Back up all files in the /home/proj directory and all its subdirectories.	dsmc selective /home/proj/ -subdir=yes	If you specify -subdir=*yes* when backing up a specific path and file, Tivoli Storage Manager recursively backs up *all* subdirectories under that path, and any instances of the specified file that exist under *any* of those subdirectories. If a subdirectory is a mounted file system, Tivoli Storage Manager does not back up the files in that subdirectory when you use the subdir=*yes* option. See "Subdir" on page 432 for more information about the subdir option.
Back up all files in the /Users/van/Documents directory and all its subdirectories.	dsmc selective /Users/van/Documents/ -subdir=yes	If you specify -subdir=*yes* when backing up a specific path and file, Tivoli Storage Manager recursively backs up *all* subdirectories under that path, and any instances of the specified file that exist under *any* of those subdirectories. If a subdirectory is a mounted file system, Tivoli Storage Manager does not back up the files in that subdirectory when you use the subdir=*yes* option. See "Subdir" on page 432 for more information about the subdir option.
Back up the /home/dir1/h1.doc and /home/dir1/test.doc files.	dsmc selective /home/dir1/h1.doc /home/dir1/test.doc	If you specify the removeoperandlimit option with the **incremental** or **selective** commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits. This allows you to specify more than 20 files on a single command. See "Removeoperandlimit" on page 396 for more information about this option.
Back up the /Users/ann/Documents/ h1.doc and /Users/ann/Documents/ test.doc files.	dsmc selective /Users/ann/Documents/ h1.doc /Users/ann/Documents/test.doc	If you specify the removeoperandlimit option with the **incremental** or **selective** commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits. This allows you to specify more than 20 files on a single command. See "Removeoperandlimit" on page 396 for more information about this option.
Back up a list of files in the /home/filelist.txt file.	selective -filelist=/home/filelist.txt	Use the filelist option to process a list of files. See "Filelist" on page 333 for more information.

Task	Command	Considerations
Back up all files listed in the /Users/filelist.txt file.	`dsmc selective -filelist=/Users/ filelist.txt`	Use the filelist option to process a list of files. See "Filelist" on page 333 for more information.
Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, run a selective backup of the /usr/dir1/sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name /usr.	`dsmc selective /usr/dir1/sub1/ -subdir=yes -snapshotroot=/snapshot/ day1`	Tivoli Storage Manager considers the snapshotroot value as a file space name. See "Snapshotroot" on page 425 for more information.

Related reference

"Incremental" on page 492

"Selective" on page 548

Deleting backup data

If your administrator has given you authority, you can delete individual backup copies from the Tivoli Storage Manager server without deleting the entire file space. To determine if you have this authority, select **File** → **Connection Information** from the Tivoli Storage Manager GUI or Web client main menu. Your authority status is provided in the **Delete Backup Files** field.

Important: When you delete backup files, *you cannot restore them.* Verify that the backup files are no longer needed before you delete them. Tivoli Storage Manager prompts whether you want to continue with the delete. If you specify *yes*, the specified backup files are immediately deleted and removed from Tivoli Storage Manager server storage.

To delete backup copies using the Tivoli Storage Manager GUI or Web client:

1. Select **Delete Backup Data** from the **Utilities** menu. The Backup Delete window appears.
2. Expand the Directory tree by clicking the plus sign **(+)** or folder icon next to the object you want to expand.
3. Click the selection boxes next to objects that you want to delete.
4. Select an item from the drop-down list near the top of the **Backup Delete** window to specify the type of backup delete to perform. You can delete active backup versions, inactive backup versions, or all objects that you have selected in the tree.

Note:

1. A directory is deleted only if you select **Delete All Objects**.
2. To delete backup copies using the Tivoli Storage Manager command line client, use the **delete backup** command.

Related reference

"Delete Backup" on page 484

Deleting file spaces

If your Tivoli Storage Manager administrator gives you authority, you can delete entire file spaces from the server. When you delete a file space, you delete all the files and images, both backup versions and archive copies, that are contained within the file space. For example, if you delete the /tmp file space, you are deleting every backup for every file in that file system and every file you archived from that file system. **Carefully consider whether you want to delete a file space**.

You can also delete a file space using the **delete filespace** command. Use the class option with the **delete filespace** command to delete NAS file spaces.

You can delete individual backup versions by using the **delete backup** command.

You can delete file spaces using the Tivoli Storage Manager GUI or command line clients. To delete NAS file spaces, use the Web client or command line client.

To delete a file space using the GUI, perform the following steps:

1. Select **Utilities→ Delete Filespaces** from the main window.
2. Click the selection boxes next to the file spaces you want to delete.
3. Click the **Delete** button. Tivoli Storage Manager prompts you for confirmation before deleting the file space.

Related reference

"Class" on page 279

"Delete Backup" on page 484

"Delete Filespace" on page 487

Backing up files from one or more file spaces for a group backup (UNIX and Linux)

You can use the **backup group** command to create and back up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.

Restriction: The **backup group** command does not apply to Mac OS X.

A *group backup* allows you to create a consistent point-in-time backup of a group of files that is managed as a single logical entity:

- All objects in the group are assigned to the same management class.
- Existing *exclude* statements for any files in the group are ignored.
- All objects in the group are exported together.
- All objects in the group are expired together as specified in the management class. No objects in a group are expired until all other objects in the group are expired, even when another group they belong to gets expired.

A group backup can be added to a backup set.

You can perform a full or differential backup using the mode option.

For example, to perform a full backup of all the files named in the /home/dir1/filelist1 file to the virtual file space /virtfs containing the group leader /home/group1 file, enter:

```
dsmc backup group -filelist=/home/dir1/filelist1 -groupname=group1 -virtualfsname=
/virtfs -mode=full
```

Related concepts

"Restore data from a backup set" on page 179

Related reference

"Backup Group" on page 471

"Include options" on page 351

"Mode" on page 369

Backing up data with client-node proxy support (UNIX and Linux)

Backups of multiple nodes that share storage can be consolidated to a common target node name on the Tivoli Storage Manager server.

This is useful when the workstation responsible for performing the backup can change over time, such as with a cluster.

The asnodename option also allows data to be restored from a different system than the one which performed the backup.

An agent node is a client node which has been granted authority to perform client operations on behalf of a target node.

A target node is a client node which grants authority to one or more agent nodes to perform client operations on its behalf.

Use the asnodename option with the appropriate command to back up, archive, restore, and retrieve data under the target node name on the Tivoli Storage Manager server. This support is only available with Tivoli Storage Manager Version 5.3 and higher server and client.

To enable this option, follow these steps:
1. Install the backup-archive client on all nodes in a shared data environment.
2. Register each node with the Tivoli Storage Manager server, if it does not exist. Register the common target node name to be shared by each of the agent nodes used in your shared data environment.
3. Register each of the nodes in the shared data environment with the Tivoli Storage Manager server. This is the agent node name that is used for authentication purposes. Data is not stored using the node name when the asnodename option is used.
4. Grant proxy authority to all nodes in the shared environment to access the target node name on the Tivoli Storage Manager server, using the GRANT PROXYNODE command (Tivoli Storage Manager administrator).
5. Use the QUERY PROXYNODE administrative client command to display the client nodes of the authorized user, granted by the GRANT PROXYNODE command.

Suggestions:
- All agent nodes in the multiple node environment should be of the same platform type.

- Do not use target nodes as traditional nodes, especially if you encrypt your files before backing them up to the server.

Restrictions enforced within a proxied session:
- You cannot access another node (either from GUI drop down or use of the fromnode option).
- You cannot perform NAS backup or restore.

Related reference

"Asnodename" on page 268

Enabling multiple node operations from the GUI

This topic lists the steps that you must follow to enable multiple node operation from the GUI.

1. Verify that the client node has proxy authority to a target node (or authorized to act as the target node) using the QUERY PROXYNODE administrative client command.
2. Select **Edit → Preferences** to open the preferences window.
3. Select the **General** tab and fill in the **As Node Name** field with the name of the target node.
4. Click **Apply** and then **OK** to close the preferences window.

Perform one of the following steps to verify that your client node is now accessing the server as the target node:
- Open the tree window and check that the target node name specified by the **As Node Name** field appears.
- Verify the target node name in the **Accessing As Node** field in the **Connection Information** window.

To return to single node operation, delete the **As Node Name** from the **Accessing As Node** field in the **General → Preferences** tab.

Setting up encryption

This topic lists the steps that you must follow to set up encryption with the encryptkey option.

1. Specify *encryptkey=save* in the options file.
2. Back up at least one file with *asnode=ProxyNodeName* to create a local encryption key on each agent node in the multiple node environment.

Follow these steps to set up encryption with the *encryptkey=prompt* option:

1. Specify *encryptkey=prompt* in the options file.
2. Ensure that users of the agent nodes in the multiple node environment are using the same encryption key.

Important:
- If you change the encryption key, you must repeat the previous steps.
- Use the same encryption key for all files backed up in the shared node environment.

Scheduling backups with client-node proxy support

Multiple nodes can be used to perform backup operations using the scheduler.

When you grant proxy authority to the agent nodes, they perform scheduled backup operations on behalf of the target node. Each agent node must use the asnodename option within their schedule to perform multiple node backup for the agent node.

Start the schedules using `dsmc sched` client command:

The following examples show the administrative client-server commands using the scheduler on multiple nodes.

- The administrator registers all of the nodes to be used: (1) `register node NODE-A`, (2) `register node NODE-B`, (3) `register node NODE-C`.
- The administrator grants proxy authority to each agent node: (1) `grant proxynode target=NODE-Z agent=NODE-A`, (2) `grant proxynode target=NODE-Z agent=NODE-B`, (3) `grant proxynode target=NODE-Z agent=NODE-C`.
- The administrator defines the schedules: (1)`define schedule standard proxy1 description="NODE-A proxy schedule" action=incremental options="-asnode=NODE-Z" objects=/Volumes/Xsan1 startdate=05/21/2005 starttime=01:00`, (2) `define schedule standard proxy2 description="NODE-B proxy schedule" action=incremental options="-asnode=NODE-Z" objects=/Volumes/Xsan2    startdate=05/21/2005 starttime=01:00`, (3) `define schedule standard proxy3 description="NODE-C proxy schedule" action=incremental options="-asnode=NODE-Z" objects=/Volumes/Xsan3 startdate=05/21/2005 starttime=01:00`.

Note: Place the *asnode* option in the schedule definition only. Do not place it in the client options file, on the command line, or in any other location.

For more information about the server scheduler commands, see *IBM Tivoli Storage Manager Administrator Command Reference*.

You can also use the client acceptor daemon (dsmcad), with managedservices set to schedule in the systems options file.

Note:
1. Each schedule can be started from a different workstation or LPAR.
2. After running the schedules, any proxied client can query and restore all of the backed up data.

Examples of how to schedule a backup of an HACMP cluster
This section shows lists some examples of how to back up an HACMP cluster.

Perform the following steps to enable scheduling of multiple nodes:
1. Ensure that all agent nodes must have proxy authority over the common target node
2. Ensure that all agent nodes must have a schedule defined on the server:
   ```
   def sched domain_name sched_name options='-asnode=target'
   ```
3. Ensure that each agent node must have its schedule associated with a node:
   ```
   def association domain_name schedule_name <agentnodename>
   ```

In the following examples, HACMP is configured for two AIX hosts, `host_a` and `host_b`. Along with their own local data, the hosts are sharing disk storage which has two file spaces: `/disk1` and `/disk2`.

The CLUSTERNODE example shows how the clusternode option is used in a current HACMP environment.

- The administrator defines 3 nodes on the Tivoli Storage Manager server: `host_a`, `host_b`, `cluster_group`, using the following commands: (1) REGISTER NODE `host_a mysecretpa5s`, (2) REGISTER NODE `host_b mysecretpa5s`, (3) REGISTER NODE `cluster_group mysecretpa5s`.
- The administrator defines a dsm.opt file on `host_a` and `host_b` (note that the opt files are different on each host), using the following commands: (1) NODENAME `host_a (option can be left as default)`, (2) DOMAIN `/home /usr ... etc.`.
- The administrator defines a dsm.opt file located somewhere on one of the cluster disk groups, for example, `/disk1/tsm/dsm.opt`, using the following commands: (1) NODENAME `cluster_group`, (2) DOMAIN `/disk1 /disk2`, (3) CLUSTERNODE YES.
- The administrator defines a schedule on the Tivoli Storage Manager server, using the following command: DEFINE SCHEDULE STANDARD CLUSTER_BACKUP.
- The administrator defines associations for each of the 3 nodes, using the following command: DEFINE ASSOC STANDARD CLUSTER_BACKUP `host_a,host_b,cluster_group`. At any one time, there are three instances of the Tivoli Storage Manager Backup-Archive client schedule running (with the scheduler for `cluster_group` being part of the cluster resources that failover whenever the cluster group disk resources failover. Thus, it would be running on either `host_a` or `host_b` but not both simultaneously).
- All three node names contain data on the Tivoli Storage Manager server.

The ASNODE example shows a generic solution which could be applied to UNIX cluster solutions to which we do not have support, for example: Veritas Cluster Server for Solaris.

- The administrator defines 3 nodes on the Tivoli Storage Manager server `host_a`, `host_b`, `cluster_group`:
```
REGISTER NODE host_a mysecretpa5s
REGISTER NODE host_b mysecretpa5s
REGISTER NODE cluster_group mysecretpa5s
```
- The administrator defines a proxy node relationship between `host_a` and `host_b` to `hacmp_cluster`
```
GRANT PROXYNODE TARGET=cluster_group AGENT=host_a,host_b
```
- The administrator defines a dsm.opt file on `host_a` and `host_b` to handle the local file systems:
```
NODENAME  host_a (option can be left as default)
DOMAIN    /home /usr ... etc.

NODENAME  host_b (option can be left as default)
DOMAIN    /home /usr ... etc.
```
- The administrator defines a dsm.opt file on the cluster resource to handle the backup of the clustered resources, e.g. `/disk1/tsm/dsmcluster.opt` (the nodename is the default nodename, which is either `host_a` or `host_b`, depending on which workstation contains the cluster group at any given time):
```
DOMAIN    /disk1 /disk2
ASNODE    cluster_group
```
- The administrator defines a schedule on the Tivoli Storage Manager server:
```
DEFINE SCHEDULE STANDARD CLUSTER_BACKUP
```
- The administrator defines associations for each one of the 3 nodes.
```
DEFINE ASSOC STANDARD CLUSTER_BACKUP host_a,host_b,cluster_group
```
- At any one time, there are three instances of the Tivoli Storage Manager Backup-Archive client schedule running with the scheduler for node

hacmp_cluster running on either host_a or host_b but not both (it is included in the cluster resources that would failover). This scheduler would point to the dsmcluster.opt that is defined on each host. The three instances would be started as:

```
[host_a]          dsmc sched
[host_b]        dsmc sched
[cluster_group] dsmc sched -optfile=/disk/tsm/dsmcluster.opt
```

- All three node names contain data on the Tivoli Storage Manager server.

For more information about the server scheduler commands, see *IBM Tivoli Storage Manager Administrator Command Reference*.

Examples of how to schedule a backup of a GPFS file system

The following examples show the administrative client-server commands using the scheduler to back up a GPFS file system, using three nodes in the GPFS cluster which participate in the backup.

- The administrator defines four nodes on the Tivoli Storage Manager server, using the following commands: (1) REGISTER NODE node_1 mysecretpa5s, (2) REGISTER NODE node_2 mysecretpa5s, (3) REGISTER NODE node_3 mysecretpa5s, (4) REGISTER NODE node_gpfs mysecretpa5s. node_1, node_2, node_3 and node_gpfs. node_1, node_2, and node_3 are only used for authentication; all file spaces are stored with node_gpfs.

- The administrator defines a proxynode relationship between the nodes, using the following commands: GRANT PROXYNODE TARGET=node_gpfs AGENT=node_1, node_2, node_3.

- The administrator defines the node name and asnodename for each of the workstations in their respective dsm.sys files, using the following commands: (1) nodename node_1, (2) asnodename node_gpfs.

- The administrator defines a schedule for only node_1 to do the work, using the following commands: (1) DEFINE SCHEDULE STANDARD GPFS_SCHEDULE ACTION=MACRO OBJECTS="gpfs_script", (2) DEFINE ASSOCIATION STANDARD GPFS node_gpfs.

- To execute the schedule on node node_gpfs, enter the client command: DSMC SCHED.

Another way to back up GPFS is to use Tivoli Storage Manager to look for the incremental changes. The GPFS file system can be divided into three branches and each branch can be statically assigned to each node using the virtualmountpoint option. In the following example, you have a file system called /gpfs with three branches: /gpfs/branch_1, /gpfs/branch_2, and /gpfs/branch_3.

- The administrator defines four nodes on the Tivoli Storage Manager server: node_1, node_2, node_3 and node_gpfs. node_1, node_2 and node_3 are only used for authentication; all file spaces are stored with node_gpfs.

```
REGISTER NODE node_1 mysecretpa5s
REGISTER NODE node_2 mysecretpa5s
REGISTER NODE node_3 mysecretpa5s
REGISTER NODE node_gpfs mysecretpa5s
```

- The administrator defines a proxynode relationship between the nodes:

```
GRANT PROXYNODE TARGET=node_gpfs AGENT=node_1,node_2,node_3
```

- The administrator defines the node name, virtualmountpoint and domain for each of the three workstations in their respective dsm.sys files:

```
nodename        node_1
virtualmountpoint /gpfs/branch_1
domain          /gpfs/branch_1
```

Note: The user does not want to define asnodename in the options file. In this case the asnodename must be on the schedule so that each one of the nodes can have its own schedule associated with its real node name.

- The administrator defines a schedule for all nodes: node_1, node_2 and node_3

```
DEFINE SCHEDULE STANDARD GPFS_SCHEDULE OPTIONS="-asnode=node_gpfs"
DEFINE ASSOCIATION STANDARD GPFS node_1,node_2,node_3
```

- To start the scheduler on the three nodes, enter the client command:

```
DSMC SCHED
```

Associate a local snapshot with a server file space (UNIX and Linux)

Use the snapshotroot option with the **incremental** and **selective** commands in conjunction with an independent software vendor application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server.

The snapshotroot option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot.

Related reference

"Snapshotroot" on page 425

Image backup

From your local workstation, you can back up a logical volume as a single object (image backup) on your system.

The traditional *static* image backup prevents write access to the volume by other system applications during the operation.

You must be a root user to perform this task, and image backup does not apply to Mac OS X.

An image backup provides the following benefits:

- Backs up file systems containing a large number of files faster than a full file system incremental backup.
- Improves the speed with which Tivoli Storage Manager restores file systems containing many small files.
- Conserves resources on the server during backups since only one entry is required for the image.
- Provides a point-in-time picture of your logical volume, which might be useful if your enterprise needs to recall that information.
- Restores a corrupt file system or raw logical volume. Data is restored to the same state it was when the last logical volume backup was performed.

The traditional *static* image backup prevents write access to the volume by other system applications during the operation. Use the dynamicimage option to back up the volume *as is* without remounting it read-only. Corruption of the backup can occur if applications continue to write to the volume while the backup is in progress, which can result in inconsistent data and data loss after the restore. This option overrides the copy serialization value in the management class to perform an image backup. After restoring an image backup taken with the dynamicimage option, always run the `fsck` utility.

To restore an image backup of a volume, the Tivoli Storage Manager client must be able to obtain an exclusive lock on the volume being restored.

Restriction: Do not use Tivoli Storage Manager dynamic image backups for file systems, because the file system might provide inconsistent data even when there is no write activity. Also, dynamic image backup might result in a fuzzy image, which might not be valid or complete when restored.

If the Tivoli Storage Manager client fails to mount the file system after restoring an image, you can run `fsck`. However, running `fsck` can affect the integrity of large amounts of data. Do *not* use dynamic image backup for AIX JFS2 file systems. The Tivoli Storage Manager client does not allow dynamic image backup for AIX JFS2 file systems. If you specify dynamicimage=*yes* for a JFS2 file system, the Tivoli Storage Manager client performs a snapshot-based image backup. If the snapshot cannot be created for some reason, the Tivoli Storage Manager client instead performs a static image backup.

Attention: To prevent data loss, avoid using the dynamicimage option, and ensure that there is no write activity on the volume while the backup is in progress.

For AIX JFS2 file systems, the amount of data backed up to the Tivoli Storage Manager server during static or snapshot image backup is reduced by backing up only those blocks used by the file system or smaller than the imagegapsize option. This method of backing up your data improves the performance of image backup. Refer to "Imagegapsize" on page 347 for more information.

For AIX clients *only*: By default, Tivoli Storage Manager performs an online snapshot image backup of JFS2 file systems, during which the volume is available to other system applications.

For Linux clients *only*: By default, Tivoli Storage Manager performs a *snapshot* image backup of file systems residing on a logical volume created by the Linux Logical Volume Manager, during which the volume is available to other system applications.

Attention: File systems managed by Tivoli Storage Manager for space management are not enabled for image backup.

Related tasks

"Snapshot-based file backup and archive and snapshot-based image backup" on page 156

Performing prerequisite tasks before creating an image backup

This topic lists some items to consider before you perform an image backup.

The following items are the image backup considerations.
- Ensure that no other application is using the volume when you run a static image backup. To ensure a consistent image during backup processing, if a file space is detected on the volume the client unmounts and remounts the volume as read only, so that no other applications can write to it. If the volume is in use when the client attempts to unmount, the backup fails. If the client cannot unmount and remount the volume as read only because it is in use, and snapshot image backup is not available, you can use the dynamicimage option to force the client to perform an image backup without unmounting and

remounting the volume in read-only mode. Set the dynamicimage option in an include.image statement or from the command line. The backup can be corrupted if applications write to the volume while the backup is in progress. This can be corrected by running `fsck` after a restore to fix any corrupted blocks.

If no file system is detected on the volume being backed up, ensure that all applications writing to the volumes are quiesced. The backup-archive client uses the file system table and mount table to detect the supported file systems.

Do not include system files in an image backup because file systems being actively used cannot be unmounted.

For AIX and Linux only: If you perform an image backup of a mounted file system which is mounted to another mount point and specified in the file system table, then after completing the image backup, all mount options for this file system, except read or write state, is lost.

Important: If a mounted file system has nested mount points, unmount them before attempting a backup. Otherwise, Tivoli Storage Manager is unable to unmount the volume. The file system is rendered *busy* if it contains any mounts.

- Use the include.image option to assign a management class to the volume image. If you do not assign a management class, the default management class is used for the image.

- You can exclude a volume from image backup using the exclude.image option.

- You must use the mount point for the file system volume on which you want to perform an image backup. Tivoli Storage Manager will not back up a file system volume without the use of a mount point. Back up file systems using the mounted name. For example, if `/dev/lv01` is formatted as a file system mounted on `/home`, enter this command to perform an image backup of this volume:

  ```
  dsmc backup image /home
  ```

  Back up raw volumes using the device name. For example, if `/dev/lv02` is a raw volume, enter this command to perform an image backup of this volume:

  ```
  dsmc  backup image /dev/lv02
  ```

  If you back up a raw volume which is formatted as a file system, ensure that the file system is not mounted and does not have an entry in `/etc/filesystems`.

Related concepts

Chapter 9, "Storage management policies," on page 229

Related reference

"Exclude options" on page 321

"Include options" on page 351

Volume device type support for an image backup

This topic lists several devices that are supported by the **backup image** command.

The following table lists the devices supported by the **backup image** command. A raw device might be a disk slice, a partition, or a logical volume.

Table 53. Volume device-type support for an image backup

Logical volume manager	Raw device types	Sample device name	Backup image command support
AIX Logical Volume Mgr	Logical Volumes	/dev/lv00	AIX
Sun Solstice DiskSuite Volume Manager	Metadevices	/dev/md/dsk/dl	Solaris
Solaris Volume Manager	Metadevices	/dev/md/dsk/dl	Solaris

Table 53. Volume device-type support for an image backup (continued)

Logical volume manager	Raw device types	Sample device name	Backup image command support

Veritas Volume Mgr	Logical Volumes	/dev/vx/dsk/rootdg/vol01
		   - AIX
		/dev/vg00/lvol01
		   - Solaris
			Solaris SPARC
			AIX

Raw Disk	Partitions	/dev/hda1, /dev/sda3	Linux x86/x86_64, Linux on POWER
Raw Disk	Partitions	/dev/dasdxx	Linux on System z
Linux Logical Volume Mgr	Logical Volumes	/dev/myvolgroup/ myvolume	All Linux
Raw Disk	Disk Slices	/dev/dsk/c0tld0s0	Solaris SPARC/x86
HP-UX Logical Volume Mgr	Logical Volumes	/dev/dsk/c0t0d1 /dev/vg00/lvol01	HP-UX, HP-UX Itanium 2

The client must support the raw device type on the specific platform in order to perform an image backup of a raw device. If you want to perform an image backup for a file system mounted on a raw device, the raw device must be supported. Remember to specify raw devices by their block device name.

For the Linux clients, image backup is only supported on partitions with id 0x83 or logical volumes created with the Linux Logical Volume Manager. Backing up other partitions, such as extended partitions that contain mounted file systems or database data, might produce inconsistent backup data if the data changes during the image backup operation.

On HP-UX, image backup for raw volumes is only supported for devices that are created using HP-UX LVM logical volumes, such as /dev/dsk/c0t0d1. HP-UX LVM logical volume devices usually take the form /dev/vgXY/lvolAB.

For AIX and Solaris: You can perform image operations on volumes created using Veritas Volume Manager. Tivoli Storage Manager initially supports static (default) and dynamic image type for backup.

For Solaris 10 clients, only use image backup for file systems that are assigned from the global zone to the non-global zone by exporting the device, specifying add device and set match. Do not use image backup for other file systems in the non-global zones because the non-global zone does not have the authority to mount or unmount the file system. Also, for Solaris 10 clients, do not use the overlap device of the root disk (c0t0d0s2) for raw device backup. Avoid using this feature on disks or slices that are used as swapping devices.

Meta devices created by the Veritas Volume Manager must be listed, including the disk group in /etc/vfstab, to be recognized by the Tivoli Storage Manager Backup-Archive client for an image backup of file systems. The file systems should be unmounted. Raw devices should not be listed in /etc/vfstab. For example, the following is the correct meta device name to be used in the /etc/vfstab file:

```
/dev/vx/dsk/<disk group>/<meta device name>
```

Specifying /dev/vx/dsk/ would not be recognized correctly, and you would receive an error (ANS1134E).

Disk slices containing cylinder 0 should not be backed up or restored. In this case the VTOC is overwritten. If you need to back up the first disk slice, exclude cylinder 0 by starting the disk slice from cylinder 1 (use the format utility). The Tivoli Storage Manager Backup-Archive client does not check whether cylinder 0 is contained in the device that is overwritten during a restore.

Utilizing image backups to perform file system incremental backups

This topic lists the methods and steps to use image backups to perform efficient incremental backups of your file system.

These backup methods allow you to perform a point-in-time restore of your file systems and improve backup and restore performance. You can perform the backup only on formatted volumes; not on raw logical volumes.

You can use one of the following methods to perform image backups of volumes with mounted file systems.

Method 1: Using image backups with file system incremental backups

This topic lists the steps to perform image backups with file system incremental backup.

1. Perform a full incremental backup of the file system. This establishes a baseline for future incremental backups.
2. Perform an image backup of the same file system to make image restores possible.
3. Perform incremental backups of the file system periodically to ensure that the server records additions and deletions accurately.
4. Perform an image backup periodically to ensure faster restore.
5. Restore your data by performing an incremental restore. Ensure that you select the **Image plus incremental directories and files** and **Delete inactive files from local** options in the Restore Options window before beginning the restore. During the restore, the client does the following:
- Restores the most recent image on the server.
- Deletes all of the files restored in the previous step which are inactive on the server. These are files which existed at the time of the image backup, but were subsequently deleted and recorded by a later incremental backup.
- Restores new and changed files from the incremental backups.

Note: If an incremental backup is performed several times after backing up an image, make sure that the backup copy group of the Tivoli Storage Manager server has enough versions for existing and deleted files on the server so that the subsequent restore image with incremental and deletefiles options can delete files correctly.

Method 2: Using image backups with incremental-by-date image backups

This topic lists the steps to perform image backups with incremental-by-date image backup.

1. Perform an image backup of the file system.
2. Perform an incremental-by-date image backup of the file system. This sends only those files that were added or changed since the last image backup to the server.
3. Periodically, perform full image backups.
4. Restore your volume by performing an incremental restore. Ensure that you select the **Image plus incremental directories and files** option in the Restore Options window before beginning the restore. This first restores the most recent image and then restores all of the incremental backups performed since that date.

Note: You should perform full image backups periodically in the following cases:
- When a file system changes substantially (more than 40%), as indicated in step 4 of method 1 and step 3 of method 2. On restore, this would provide a file system image close to what existed at the time of the last incremental-by-date image backup and it also improves restore time.
- As appropriate for your environment.

This improves restore time because fewer changes are applied from incremental backups.

The following restrictions apply when using method 2:
- The file system can have no previous full incremental backups.
- Incremental-by-date image backup does not inactivate files on the server; therefore, when you restore an image with the incremental option, files deleted after the original image backup is present after the restore.
- If this is the first image backup for the file system, a full image backup is performed.
- If file systems are running at or near capacity, an out-of-space condition could result during the restore.

Comparing methods 1 and 2

This topic shows a comparison of methods 1 and 2: (1) Using image backup with file system incremental or (2) Using image backup with incremental-by-date image backup.

To help you decide which method is appropriate for your environment, the following table is a comparison of methods 1 and 2.

Table 54. Comparing incremental image backup methods

Method 1: Using image backup with file system incremental	Method 2: Using image backup with incremental-by-date image backup
Files are expired on the server when they are deleted from the file system. On restore, you have the option to delete files which are expired on server from image.	Files are not expired on server. After the image incremental restore completes, all of the files that are deleted on the file system after the image backup are present after the restore. If file systems are running at or near capacity, an out-of-space condition could result.
Incremental backup time is the same as regular incremental backups.	Incremental image backup is faster because the client does not query the server for each file that is copied.
Restore is much faster compared to a full incremental file system restore.	Restore is much faster compared to a full incremental file system restore.
Directories deleted from the file system after the last image backup are not expired.	Directories and files deleted from the file system after the last full image backup are not expired.

Performing an image backup using the GUI

If the Tivoli Storage Manager image feature is configured, you can create an image backup where the real volume is available to other system applications.

A consistent image of the volume is maintained during the image backup.

When you perform an image backup using the client Java GUI image backup option, Tivoli Storage Manager honors the setting of the snapshotproviderimage option. The snapshotproviderimage option defaults to an AIX JFS2 snapshot for AIX and a Linux LVM snapshot for Linux. You can override the default by using the Preferences editor Snapshot tab and the Image Snapshot Preferences.

For HP-UX and Solaris clients, selecting the image backup option performs a static image backup by default. For static image backup, the client unmounts and remounts the volume as read-only, so that no other applications can access it. You can override the default value by using the include.image option and selecting dynamicimage *yes*. For dynamic image backup, the client performs the image backup without making the file system read-only during the backup.

To create an image backup of your file system or raw logical volume, perform the following steps:

1. Click on the **Backup** button in the Tivoli Storage Manager main window. The Backup window appears.
2. Expand the directory tree and select the objects you want to back up. To back up a raw logical volume, locate and expand the **RAW** directory tree object.
3. Click **Backup**. The Backup **Task List** window displays the backup processing status. The Backup Report window displays a detailed status report.
- To perform a static image backup, select **Image Backup** from the drop-down list.
- For AIX and Linux clients *only*: To perform a snapshot image backup, use the snapshotproviderimage option.

- To perform an incremental-by-date image backup, select **Incremental image (date only)** from the drop-down list.

The following are some items to consider when you perform an snapshot-based image backup:

- To modify specific backup options, click the **Options** button. The options you select are effective during the current session *only*.
- To modify specific backup options, click the **Options** button. The options you select are effective during the current session *only*.

Linux only: The Tivoli Storage Manager Version 5.4 client will not recognize any LVM1 volumes for image operations. However, it allows prior image backups of LVM1 volumes to be restored on LVM2 volumes. Table 55 shows the combinations involving the old and new client levels handling LVM1 and LVM2 volumes for different image operations.

Table 55. LVM1 and LVM2 image operation comparisons

Tivoli Storage Manager client version	LVM1 Backup and Restore	LVM2 Backup and Restore	Mixed Volumes	
			Backup: LVM1, Restore: LVM2	Backup: LVM2, Restore: LVM1
V5.3 and prior	YES	Only static image for file system	NO	NO - raw volumes are not supported
V5.4 and beyond	NO Error msg ANS1090E displayed	YES	YES LVM1 vol must have been backed up using prior client	NO Restore to LVM1 vol fails

Related reference

"Snapshotproviderimage" on page 424

Performing an image backup using the command line

Use the **backup image** and **restore image** commands to perform image backup and restore operations on a single volume.

Use the mode option with the **backup image** command to perform an incremental-by-date image backup that backs up only new and changed files after the last full image backup. However, this only backs up files with a changed date, not files with changed permissions.

Related reference

"Backup Image" on page 473

"Mode" on page 369

"Restore Image" on page 539

Snapshot-based file backup and archive and snapshot-based image backup

For Tivoli Storage Manager clients running on AIX 5.3 or later JFS2 file systems as root user, snapshot-based image backup is created using snapshots by default.

Optionally, you can enable snapshot-based file level backup and archive operations by specifying the snapshotproviderfs option. If for some reason a snapshot cannot be taken, Tivoli Storage Manager attempts to perform a static image backup or regular file backup.

If you want to specify snapshot-based file backup and archive, set the option snapshotproviderfs to JFS2. This is applicable to all JFS2 file systems for that client.

Important: Use snapshot-based file backup and archive and snapshot-based image backup for all of your AIX JFS2 file systems.

For example, to turn *on* snapshot-based file backup and archive for all JFS2 file systems on the client, specify the following in the server stanza in the dsm.sys file:

```
snapshotproviderfs JFS2
```

To explicitly turn *off* snapshot-based file backup and archive for all JFS2 file systems on the client, specify the following in the server stanza in the dsm.sys file:

```
snapshotproviderfs   NONE
```

To turn *on* snapshot-based file backup and archive for only one specific JFS2 file system on the client, specify the following in the server stanza in the dsm.sys file:

```
snapshotproviderfs   NONE

include.fs   /kalafs1   snapshotproviderfs=JFS2
```

To turn *off* snapshot-based file backup and archive for only one specific JFS2 file system on the client, specify the following in the server stanza in the dsm.sys file:

```
snapshotproviderfs   JFS2

include.fs   /kalafs2   snapshotproviderfs=NONE
```

To turn *on* snapshot-based file backup and archive for only one specific operation on the client, specify the following on the command line:

```
dsmc incr   -snapshotproviderfs=JFS2   /kalafs1
```

To turn *off* snapshot-based file backup and archive for only one specific operation on the client, specify the following in the server stanza in the dsm.sys file:

```
snapshotproviderfs   JFS2
```

Then perform the backup command. For example:

```
dsmc incr -snapshotproviderfs=NONE /kalafs2
```

The snapshotproviderfs option does not apply to **incremental** with the snapdiff and diffsnapshot options.

Related reference

"Snapshotproviderfs" on page 423

Back up NAS file systems using Network Data Management Protocol

Through support of Network Data Management Protocol (NDMP), Tivoli Storage Manager Windows, AIX, and Solaris backup-archive clients can efficiently back up and restore network attached storage (NAS) file system images to automated tape drives or libraries that are locally attached to Network Appliance and EMC Celerra NAS file servers.

NDMP support is available only on IBM Tivoli Storage Manager Extended Edition.

Incremental backup can also be used to back up NAS file system snapshots. See the **incremental** command and snapshotroot, snapdiff, createnewbase, and diffsnapshot options for more information.

After configuring NDMP support, the server connects to the NAS device and uses NDMP to initiate, control, and monitor each backup and restore operation. The NAS device performs outboard data transfer to and from the NAS file system to a locally attached library.

Filer to server data transfer is available for NAS devices that support NDMP Version 4.

The benefits of performing backups using NDMP include the following:
- LAN-free data transfer.
- High performance and scalable backups and restores.
- Backup to local tape devices without network traffic.

The following support is provided:
- Full file system image backup of all files within a NAS file system.
- Differential file system image backup of all files that have changed since the last full image backup.
- Parallel backup and restore operations when processing multiple NAS file systems.
- Choice of interfaces to initiate, monitor, or cancel backup and restore operations:
  - Web client
  - Backup-archive client command interface
  - Administrative client command line interface (backup and restore operations can be scheduled using the administrative command scheduler)
  - Administrative Web client

The following functions are *not* supported:
- Archive and retrieve
- Client scheduling. Use server commands to schedule a NAS backup.
- Detection of damaged files.
- Data-transfer operations for NAS data stored by Tivoli Storage Manager:
  - Migration
  - Reclamation
  - Export
  - Backup set generation

Related concepts

"NDMP support requirements (Extended Edition only)" on page 14

Related reference

"Diffsnapshot" on page 297

"Incremental" on page 492

"Snapdiff" on page 418

"Snapshotroot" on page 425

Backing up NAS file systems with the Web client GUI using NDMP protocol

For both the Web client GUI and the client command line interface, you must specify *passwordaccess*=generate (which is a current Web client restriction for the client node) and the **authentication**=*on* must be specified at the server.

You are always prompted for a user ID and password. To display NAS nodes and perform NAS functions, you must enter an authorized administrative user ID and password. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web.

You can use the toc option with the include.fs.nas option in the client options file to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup. If you save TOC information, you can use Tivoli Storage Manager for Windows Web client to examine the entire file system tree and select files and directories to restore. Creation of a TOC requires that you define the TOCDESTINATION attribute in the backup copy group for the management class to which this backup image is bound. Note that TOC creation requires additional processing, network resources, storage pool space, and possibly a mount point during the backup operation.

To back up NAS file systems using the Web client GUI:

1. Click **Backup** from the main window. The Backup window is displayed.
2. Expand the directory tree if necessary.

   Note:
   a. The root node called **Nodes** is not selectable. This node only appears if a NAS plug-in is present on the client workstation.
   b. NAS nodes display on the same level as the client workstation node. Only nodes for which the administrator has authority appear.
   c. You can expand NAS nodes to reveal file spaces, but no further expansion is available (no file names).
3. Click the selection boxes next to the nodes or file systems you want to back up.
4. Click the type of backup you want to perform in the backup type pull-down menu. The NAS backup type list is active only when you first select NAS backup objects. **Full backup** backs up the entire file system. **Differential** backs up the changes since the most recent full backup.
5. Click **Backup**. The NAS Backup **Task List** window displays the backup processing status and progress bar. The number next to the progress bar indicates the number of bytes backed up so far. After the backup completes, the NAS Backup Report window displays processing details, including the *actual* size of the backup including the total bytes backed up.

Note: If it is necessary to close the Web browser session, current NAS operations continue after disconnect. You can use the **Dismiss** button on the NAS Backup **Task List** window to quit monitoring processing without ending the current operation.

6. (Optional) To monitor processing of an operation from the GUI main window, open the **Actions** menu and select **TSM Activities**. During a backup, the status bar indicates processing status. A percentage estimate is not displayed for differential backups.

Here are some items to consider when you back up NAS file systems using the Web client GUI:

- Workstation and remote (NAS) backups are mutually exclusive in a Backup window. After selecting an item for backup, the next item you select must be of the same type (either NAS or non NAS).
- Details will not appear in the right-frame of the Backup window for NAS nodes or file systems. To view information about objects in a NAS node, highlight the object and select **View** → **File Details** from the menu.
- To delete NAS file spaces, select **Utilities** → **Delete Filespaces**.
- Backup options do not apply to NAS file spaces and are ignored during a NAS backup operation.

Related concepts

"Restore NAS file systems" on page 190

Related tasks

"Configuring the Web client" on page 58

Related reference

"Toc" on page 444

Back up NAS file systems using the command line

You can use the command line to back up NAS file system images.

Table 56 lists the commands and options that you can use to back up NAS file system images from the command line.

Table 56. NAS options and commands

Option or command	Definition	Page
domain.nas	Use the domain.nas option to specify the volumes to include in your default domain for NAS backups.	"Domain.nas" on page 308
exclude.fs.nas	Use the exclude.fs.nas option to exclude file systems on the NAS file server from an image backup when used with the **backup nas** command. This option is for AIX and Solaris clients *only*.	"Exclude options" on page 321

Table 56. NAS options and commands (continued)

Option or command	Definition	Page
include.fs.nas	Use the include.fs.nas option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system image backup, using the *toc* option with the include.fs.nas option in your client options file..	"Include options" on page 351
	This option is for AIX and Solaris clients *only*.	
query node	Use the **query node** command to display all the nodes for which a particular administrative user ID has authority to perform operations. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using.	"Query Node" on page 518
backup nas	Use the **backup nas** command to create an image backup of one or more file systems that belong to a Network Attached Storage (NAS) file server.	"Backup NAS" on page 479
toc	Use the toc option with the **backup nas** command or the include.fs.nas option to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup.	"Toc" on page 444
monitor process	Use the **monitor process** command to display current backup and restore processes for all NAS nodes for which an administrative user has authority. The administrative user can then select one process to monitor.	"Monitor Process" on page 500
cancel process	Use the **cancel process** command to display current backup and restore processes for all NAS nodes for which an administrative user has authority. From the display, the administrative user can select one process to cancel.	"Cancel Process" on page 481
query backup	Use the **query backup** command with the class option to display information about file system images backed up for a NAS file server.	"Query Backup" on page 506
query filespace	Use the **query filespace** command with the class option to display a list of file spaces belonging to a NAS node.	"Query Filespace" on page 512
delete filespace	Use the **delete filespace** command with the class option to display a list of file spaces belonging to a NAS node so that you can choose one to delete.	"Delete Filespace" on page 487

- NAS nodes represent a new node type. The NAS node name uniquely identifies a NAS file server and its data to Tivoli Storage Manager. You can prefix the NAS

node name to the file specification to specify the file server to which the include statement applies. If you do not specify a NAS node name, the file system you specify applies to all NAS file servers.
- Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.

Note: When you initiate a NAS backup operation using the client command line interface, client GUI, or Web client the server starts a process to initiate, control, and monitor the operation. It might take several moments before you notice progress at the client command line interface because the server must perform a mount and other necessary tasks before data movement occurs.

Related reference

Back up NFS file systems

The Tivoli Storage Manager Backup-Archive client can be configured to protect files which are accessed with the Network File System (NFS) protocol.

Backup performance is better when you install the backup-archive client where the file system physically resides, but sometimes it is necessary to access file systems using NFS for purposes of backup and recovery. The Tivoli Storage Manager UNIX and Linux Backup-Archive client can back up, archive, restore and retrieve file data using an NFS mount. This includes all versions of the NFS protocol, including NFS version 2, NFS version 3, and NFS version 4.

The NFS version 4 protocol allows for the export of access control lists (ACL) with NFS. As shown in the following table, some of the Backup-Archive clients provide support to back up and restore ACL information on data accessed using NFS. Note that the NFS version 4 protocol must be installed on both the NFS client and NFS server in order to backup and restore ACL data.

Table 57. Network File System (NFS) ACL support

Tivoli Storage Manager Backup-Archive Client platform	ACL Support for data accessed using NFS version 4
AIX 5.3	yes
AIX 6.1 or later	yes

On other operating systems the Backup-Archive client cannot protect native ACL on data accessed using NFS.

The following restrictions apply when using the Backup-Archive client with data accessed using NFS:
- The Tivoli Storage Manager Backup-Archive clients cannot perform image-based backups on data accessed using NFS.
- The Tivoli Storage Manager Backup-Archive AIX client cannot perform snapshot-based file backup and archive on data accessed using NFS.
- The Tivoli Storage Manager Backup-Archive AIX client cannot perform journal-based backup on data accessed using NFS.
- The Tivoli Storage Manager Backup-Archive Linux client cannot perform snapshot backup of file systems nested under NFS V4.

Back up NFS file systems with the global namespace feature

NFS V4 clients can back up NFS file systems mounted using the global namespace feature, which is called a referral. All file systems in the global namespace are backed up under a single file space.

Here are some examples of the file systems in the global namespace that are backed up under a single file space:

```
server 'publications' has /doc file system
server 'projects' has /projects file system
server 'data'  has /data file system
```

The server account1 is the main NFS server that exports all these files systems using referral, and it is the server that all of the clients recognize. /etc/exports directory on account1 looks like this:

```
/doc -vers=4,refer=/doc@publications
/projects -vers=4,refer=/projects@projects
/data -vers=4,refer=/data@data
```

The client payroll mounts directories from the account1 server and can access all three file systems:

```
payroll:/#mount -o vers=4 account1:/ /mnt
payroll:/#ls /mnt
doc/ projects/ data/
```

The Client payroll can back up the /mnt file as one NFS file system, which backs up all other file systems.

Important: Using the virtualmountpoint option can improve system performance when backing up NFSV4 file systems using the global namespace. Add the following entries in the dsm.sys stanza to back up each mounted directory as a separate file space:

```
VIRTUALMOUNTPOINT   /doc
VIRTUALMOUNTPOINT   /projects
VIRTUALMOUNTPOINT   /data
```

Back up AIX workload partition file systems

Using the Tivoli Storage Manager Backup-Archive AIX client, you can back up and restore local partition file data within the global partition using the local partition name space available within the global partition.

Each workload partition (WPAR) has its own security domain, so only the global root user is guaranteed to have access to all of the data.

The WPARs are partitions created entirely in software within a single AIX system image, with the following attributes:

- Usually the WPAR appears to be a complete standalone AIX system
- There is no hardware assist or configuration

Workload partitions provide a secure and isolated environment for enterprise applications in terms of process, signal and file system space. Software running within the context of a workload partition appears to have its own separate instance of AIX.

The following example shows a WPAR configuration from within the global WPAR:

Global partition:

system name: `shimla`

file system: `/home /opt`

WPAR #1 configuration:

name: `wpar1`

file system: `/home`; name in global WPAR: `/wpars/wpar1/home`

WPAR #2 configuration:

name: `wpar2`

file system: `/data`; name in global WPAR: `/wpars/wpar2/data`

There are two ways to back up WPAR data, as follows:

- Back up all WPAR file systems as the file spaces within the global partition. The file space name must be used to identify the WPAR to which it belongs. All of the data is managed on one node using one schedule. Using the example configuration mentioned previously, here is a sample dsm.sys file with one server stanza for all file systems, both global and local:

```
SErvername  shimla
   TCPPort           1500
   TCPServeraddress  clinton.storage.sanjose.ibm.com
   nodename          shimla
   PasswordAccess    generate
   Domain            /wpars/wpar1/home /wpars/wpar2/data /home /opt
```

- Back up each WPAR file system under a different node name. This provides file space name segregation for each WPAR. Each WPAR must have a separate node name and a scheduler running within the global partition. Also, three scheduler services must be set up, each using a different dsm.opt file corresponding to the server stanza name. This method allows each WPAR backup operation to be managed independently of the others. Using the example configuration mentioned previously, here is a sample dsm.sys file with three server stanzas: one for `wpar1`, one for `wpar2`, and one for global partition `shimla`:

```
SErvername  shimla_wpar1
   TCPPort           1500
   TCPServeraddress  clinton.storage.sanjose.ibm.com
   nodename          wpar1
   PasswordAccess    generate
   Domain            /wpars/wpar1/home

SErvername  shimla_wpar2
   TCPPort           1500
   TCPServeraddress  clinton.storage.sanjose.ibm.com
   nodename          wpar2
   PasswordAccess    generate
   Domain            /wpars/wpar2/data

SErvername  shimla
   TCPPort           1500
   TCPServeraddress  clinton.storage.sanjose.ibm.com
   nodename          shimla
   PasswordAccess    generate
   Domain            /home /opt
```

Create and back up Solaris Zettabyte file systems

On Solaris SPARC and x86 systems, you can create and back up Zettabyte file systems (ZFS).

On Solaris SPARC and x86 systems, Zettabyte file systems (ZFS) is a transactional file system, which means that the file system state is always consistent on disk. ZFS also uses the concept of storage pools to manage physical storage.

You can create a ZFS snapshot using the following command:

```
zfs snapshot tank/myZFS@mySnapshot
```

In this example, the ZFS pool name is `tank` and the ZFS file system name is `myZFS`. Files which belong to this ZFS snapshot can be found in the subdirectory `tank/myZFS/.zfs/snapshot/mySnapshot/`. This directory is read-only and for this reason it is excluded by default backup or archive operations

There are two possible ways to back up this ZFS snapshot:
- Back up each file of the snapshot by using the snapshotroot option, as in the following example. This option allows the administrator to replace the current snapshot path with the ZFS file system path: `dsmc inc -snapshotroot=/tank/myZFS/.zfs/snapshot/mySnapshot /tank/myZFS`. Using the snapshotroot option, the files can be directly restored in the file system. If you have many changes in the file system during a backup, a snapshot can be made first, and then the Tivoli Storage Manager Backup-Archive Client can make a backup of the static snapshot.
- Another way is to back up the complete snapshot, using the following command: `zfs send tank/myZFS@mySnapshot > /tmpdir/mySnapshotFile`. Using this command, the administrator can make a copy of the complete snapshot data in a single file. This file can then be backed up by the Tivoli Storage Manager Backup-Archive Client to the Tivoli Storage Manager Server, using the following command: `dsmc sel /tmpdir/mySnapshotfile`. The advantage of backing up the complete snapshot is that the full file system can be restored in a disaster recovery case. The following command creates and restores the complete file system with the data of the snapshot: `zfs receive tank/myZFS@mySnapshot < /tmpdir/mySnapshotFile`.

Related reference

AIX JFS2 encrypted file system backup

Use AIX JFS2 Encrypted File System (EFS) to back up files either in clear text or raw format. With clear text format, the file is decrypted by EFS as it is read. With raw format, the data is not decrypted. The default is raw format, but when you set the *efsdecrypt* option to *yes*, you get clear text backups.

Important: Whenever you run a backup that includes any files encrypted on an EFS, you must ensure that you use the correct specification of the *efsdecrypt* option. If the *efsdecrypt* option value changes between two incremental backups, all encrypted files on EFS file systems are backed up again, even if they have not changed since the last backup. For example, if you are running an incremental backup of encrypted files that were previously backed up as raw, then ensure that *efsdecrypt* is specified as *no*. If you change *efsdecrypt* to *yes*, all of the files are backed up again in clear text even if they are unchanged, so ensure that you use this option carefully.

If you attempt to restore an encrypted file to either a work station that does not support EFS, or a file system where EFS is not active, an error message is written and the file is skipped.

Here are some reasons to back up EFS using clear text encryption:
- This type of decryption is useful if you want to use the Tivoli Storage Manager Backup-Archive Client encryption or another type of hardware encryption (for tape systems, for example).
- You can use clear text for long term archival of data, because the data is stored independent of the platform or encryption scheme.

Here are some things to consider when backing up a file in clear text:
- The user who invoked Tivoli Storage Manager must be able to decrypt it
- The user can have read access to a file, but not have access to the key

In the following scenarios an error message is issued:
1. The user is running in root guard mode, and EFS has the concept of two types of root. Root admin is the traditional mode. A root in guard mode will not have access to the unencrypted data, unless the user is the owner or a member of the file group.
2. The user is running with a non-root user ID and attempting an archive of a file to which they have read access, but the user is not the owner or member of the file group. EFS will not allow the data to be decrypted.

Here are some considerations when backing up EFS raw data:
- Tivoli Storage Manager will *not* honor the client encryption setting, which prevents double encryption, but only at the client. The server has no knowledge that the data is encrypted so any encryption done by a tape drive, for example, still occurs.
- Tivoli Storage Manager will *not* honor the compression setting, so the client will not even try to compress the data.
- Tivoli Storage Manager does not automatically back up or restore the keystore files. When you are restoring encrypted files, you might also have to restore keystores in order to decrypt the data.
  Tips:
  1. To protect the keystore, make sure the contents of /var/efs are included in your periodic backups.
  2. For the keystore data, use Tivoli Storage Manager storage policy with an unlimited number of versions.
- Encrypted file system (EFS) files backed up in raw mode (default) cannot be restored by a Tivoli Storage Manager Client prior to Version 5.5, or by a client on another UNIX platform.

Back up AIX JFS2 extended attributes

AIX Enhanced Journal File System (JFS2) provides backup processing for named extended attributes for all file systems that support named extended attributes.

These extended attributes are automatically backed up with each object that has extended attributes data, and no additional action is required.

When the file system is defined with the v2 format, the only file system that supports named extended attributes is JFS2. You can use JFS2 for extended attributes for files and directories, but you cannot use JFS2 for extended attributes on symbolic links.

Back up and archive Tivoli Storage Manager FastBack data

Use IBM Tivoli Storage Manager FastBack to back up and archive the latest snapshots for short-term retention.

Use the **archive fastback** and **backup fastback** commands to archive and back up volumes that are specified by the fbpolicyname, fbclientname and fbvolumename options for short-term retention.

Related concepts

"Installation requirements for backing up and archiving Tivoli Storage Manager FastBack client data" on page 15

"Configuring the client to back up and archive Tivoli Storage Manager FastBack data" on page 80

Related reference

"Fbclientname" on page 326

"Fbpolicyname" on page 328

"Fbvolumename" on page 332

Display backup processing status

During a backup, by default Tivoli Storage Manager displays the status of each file it attempts to back up.

Tivoli Storage Manager reports the size, path, file name, total number of bytes transferred, and whether the backup attempt was successful for the file. These are also recorded in the dsmsched.log file for scheduled commands.

The Web client and backup-archive client Java GUI provide a **Task List** window that displays information about files during processing. When a task completes, a **Backup Report** window displays processing details. Click the **Help** button in the Backup Report window for context help.

On the backup-archive command line the name of each file is displayed after it is sent to the server. The progress indicator shows overall progress.

Table 58 lists some informational messages and meanings.

Table 58. Client command line informational messages

Informational message	Meaning
Directory-->	Indicates the directory that you back up.
Normal File-->.	Any file that is not a directory, symbolic link or special file.

Table 58. Client command line informational messages (continued)

Informational message	Meaning
`Special File-->`	Special files define devices for the system or temporary files created by processes. There are three basic types of special files: FIFO (first-in, first-out), block, and character. FIFO files are also called pipes. Pipes are created by one process to temporarily allow communication with another process. These files cease to exist when the first process finishes. Block and character files define devices. Tivoli Storage Manager processes only device and named pipe special files. Socket special files are not processed.
`Symbolic Link-->`	Indicates that Tivoli Storage Manager backs up a symbolic link.
`Updating-->`	Indicates that only the file meta data is sent, not the data itself.
`Expiring-->`	Indicates an object (file or directory) on the server that no longer exists on the client is expired and made inactive on the server.
`Total number of objects inspected:`	As indicated. When using journal-based backup (AIX and Windows only), the number of objects inspected might be less than the number of objects backed up. When using snapshot difference incremental backup, the number of objects inspected is zero. The number is zero because the client performs an incremental backup of the files that NetApp reported as changed. The client does not scan the volume looking for files that have changed.
`Total number of objects backed up:`	As indicated.
`Total number of objects updated:`	These are files whose attributes, such as file owner or file permissions, have changed.
`Total number of objects rebound:`	See "Bind management classes to files" on page 237 for more information.
`Total number of objects deleted:`	This is a count of the objects deleted from the client workstation after being successfully archived on the server. The count is zero for all backup commands.
`Total number of objects expired:`	See the section about full and partial incremental backup for more information.
`Total number of objects failed:`	Objects can fail for several reasons. Check the `dsmerror.log` for details.
`Total objects deduplicated:`	Specifies the number of files that are deduplicated.
`Total number of bytes inspected:`	Specifies the sum of the sizes of the files that are selected for the operation. For example, the total number of bytes that are inspected for this command is the number of bytes that are used on the volume `/Volumes/BUILD`: `dsmc INCREMENTAL /Volumes/BUILD/* -SU=Yes`
`Total number of bytes processed:`	Specifies the sum of the sizes of the files that are processed for the operation.
`Data transfer time:`	The total time to transfer data across the network. Transfer statistics might not match the file statistics if the operation was retried due to a communications failure or session loss. The transfer statistics display the bytes attempted to be transferred across all command attempts.
`Network data transfer rate:`	The average rate at which the network transfers data between the client and the server. This is calculated by dividing the total number of bytes transferred by the time to transfer the data over the network. The time it takes to process objects is not included in the network transfer rate. Therefore, the network transfer rate is higher than the aggregate transfer rate.

Table 58. Client command line informational messages (continued)

Informational message	Meaning
`Aggregate data transfer rate:`	The average rate at which Tivoli Storage Manager and the network transfer data between the client and the server. This is calculated by dividing the total number of bytes transferred by the time that elapses from the beginning to the end of the process. Both Tivoli Storage Manager processing and network time are included in the aggregate transfer rate. Therefore, the aggregate transfer rate is lower than the network transfer rate. **Note:** On occasion, the aggregate data transfer rate might be higher than the network data transfer rate. This is because the backup-archive client can have multiple simultaneous sessions with the backup server. If you set the resourceutilization option, Tivoli Storage Manager attempts to improve performance and load balancing by using multiple sessions when it backs up a file space or other set of files. When multiple sessions are open during backup, the data transfer time represents the sum of the times reported by all sessions. In this case, aggregate data transfer time is incorrectly reported as higher. However, when running with a single session, the aggregate data transfer rate should always be reported as lower than the network data transfer rate.
`Objects compressed by:`	Specifies the percentage of data sent over the network divided by the original size of the file on disk. For example, if the net data-bytes are 10K and the file is 100K, then Objects compressed by: == (1 - (10240/102400)) x 100 == 90%.
`Deduplication reduction:`	Specifies the size of the duplicate extents that were found, divided by the initial file or data size. For example, if the initial object size is 100 MB, after deduplication it is 25 MB. The reduction would be: (1 - 25/100) * 100 = 75%.
`Total data reduction ratio:`	Adds incremental and compression effects. For example, if the bytes inspected are 100 MB and the bytes sent are 10 MB, the reduction would be: (1 - 10/100) * 100 = 90%
`Elapsed processing time:`	The active processing time to complete a command. This is calculated by subtracting the starting time of a command process from the ending time of the completed command process.
`Total number of bytes transferred:`	As indicated.
`LanFree bytes transferred:`	The total number of data bytes transferred during a lan-free operation. If the enablelanfree option is set to *no*, this line will not appear.
`Total number of bytes inspected:`	A sum of sizes of files selected for the operation.
`Total objects deduplicated:`	The number of files that have been deduplicated.

Backup (UNIX and Linux): Additional considerations

There are some special situations that you need to consider before you back up your data.

Stored files

When you back up and archive files, Tivoli Storage Manager stores the backups and archives in a file space in storage that has the same name as the file system or virtual mount point from which the files originated.

For example, if you have a file system named /home, and you back up a file named doc1 in the /home/monnett directory, Tivoli Storage Manager stores the file in a file space named /home. If an Authorized User later defines /home/monnett as a virtual

mount point, any files you back up from the /home/monnett directory, such as doc2, are stored in a file space named /home/monnett. If you enter this command:

```
dsmc query backup "/home/monnett/*"
```

Tivoli Storage Manager looks for files in the /home/monnett file space. It always looks for a file in the file space with the longest name that matches the file specification you include in a command. It locates the file named doc2 that was backed up after the virtual mount point was defined. However, it does not locate the file named doc1 because that file was backed up before the virtual mount point was defined and the backup was stored in the /home file space.

To list or restore the doc1 file using a command, you must explicitly specify the file space name by enclosing it in braces. For example:

```
dsmc query backup "{/home}/monnett/*"
dsmc restore {/home}/monnett/doc1
```

If the authorized user subsequently removes the /home/monnett virtual mount point, and you then back up additional files in the /home/monnett directory, the backups are once again stored in the /home file space. For example, if you now back up a file named doc3 in the /home/monnett directory, it is stored in the /home file space. It is not stored in the existing /home/monnett file space.

However, because the /home/monnett file space already exists, when you try to query or restore the doc3 file, Tivoli Storage Manager looks for the file in the /home/monnett file space unless you specify the correct file space name. For example:

```
dsmc query backup "{/home}/monnett/*"
dsmc restore {/home}/monnett/doc2
```

Note: You must explicitly specify the file space name only when there can be more than one resolution to the file specification.

For example, if the following file spaces exist in storage:

```
/home
/home/monnett
/home/monnett/project1
/home/monnett/project1/planning
```

then enter:

```
dsmc query backup "/home/monnett/project1/planning/*"
```

Tivoli Storage Manager looks for files only in the /home/monnett/project1/planning file space, even if one or more of the other file spaces contains a path with the same name. But, when you enter one of the following:

```
dsmc query backup "{/home}/monnett/project1/planning/*"
dsmc query backup "{/home/monnett}/project1/planning/*"
dsmc query backup "{/home/monnett/project1}/planning/*"
```

Tivoli Storage Manager looks for files only in the /home file space, the /home/monnett file space, or the /home/monnett/project1 file space, depending on which form you use.

Special file systems

Special file systems contain dynamic information generated by the operating system; they contain no data or files. The Tivoli Storage Manager client ignores special file systems and their contents.

Special file systems include the following:
- the /proc file system on most of the UNIX platforms
- the /dev/fd file system on Solaris
- the /dev/pts on Linux

NFS or virtual mount points

When files are backed up and archived from a file system or virtual mount point, Tivoli Storage Manager does not follow the nested NFS or virtual mount points (if any are defined on a file system). The nested NFS or virtual mount points will not be backed up or archived.

Management classes

Tivoli Storage Manager uses management classes to determine how to manage your backups on the server.

Every time you back up a file, the file is assigned a management class. The management class used is either a default selected for you, or one assigned to the file by an Authorized User with an include option in the include-exclude options list. The selected management class must contain a backup copy group in order for the file to be backed up.

Select **Utilities** → **View Policy Information** from the Java or Web client GUI to view the backup policies defined by the Tivoli Storage Manager server for your client node.

Related concepts

Chapter 9, "Storage management policies," on page 229

Back up symbolic links

Tivoli Storage Manager backs up symbolic links differently than it does regular files and directories.

The way that the client backs up symbolic links depends on options settings, whether the target directory is accessible, and the way you specify objects.

A *UNIX symbolic link* is a file that contains a pointer to another file or directory. The object the symbolic link points to is called the target object.

A symbolic link can be backed up as path information to a target directory, or it can be backed up as a directory. If the symbolic link is backed up as a directory, the files and folders in the target directory can also be backed up.

Note: Symbolic link processing as described here does not apply to Mac OS X. Symbolic links are always backed up as files and are never followed.

Related reference

"Archsymlinkasfile" on page 268

"Followsymbolic" on page 336

"Virtualmountpoint" on page 452

Examples: Incremental or selective backup of symbolic links

How the client backs up a symbolic link depends on whether the target of the symbolic link is a file or a directory, and how you specify the symbolic link on the incremental or selective backup command.

If a symbolic link points to a file, the client only backs up the path information. The client does not back up a file that is the target of a symbolic link.

If a symbolic link points to a directory, the backup depends on how the directory is specified on the command.

If a directory is specified with a trailing slash on a selective or incremental backup command, the client saves the symbolic link as a directory, and backs up the contents of the target directory.

If the symbolic link is entered without a trailing slash, or if a symbolic link is not explicitly stated in a backup file specification, the client backs up only the path information to the target directory. The contents of the target directory are not backed up.

In the following examples, assume that `symdir` is a symbolic link to target directory `/fs1/guest/`. `/fs1/guest/` contains these objects:

`/fs1/guest/file` (a file)

`/fs1/guest/dir1` (a directory)

`/fs1/guest/dir1/file1` (a file)

Example 1
        dsmc incr /home/gillis/symdir/

In this example, the client backs up the symbolic link as a directory, and backs up the contents of the target directory `/fs1/guest/`. If you specify the subdir=yes option, the client backs up subdirectories of `/fs1/guest/`.

Example 2
        dsmc incr /home/gillis/symdir/dir1

Example 3
        dsmc incr /home/gillis/symdir/dir1/

In examples 2 and 3, the client backs up the symbolic link as a directory, and backs up the `/dir1/` subdirectory of the target directory. the trailing slash is relevant only for the symbolic link; it is not relevant for subdirectories of the symbolic link. If you specify the subdir=yes option, the client backs up subdirectories of `/fs1/guest/dir1`. Backup copies that are stored on the Tivoli Storage Managerr server have a path like `/home/gillis/symdir/dir1/file1`.

Example 4
        dsmc incr /home/gillis/symdir

In example 4, because there is no trailing slash after the symbolic link, the client backs up only the path to the target directory. The client does not back up the symbolic link as a directory, and does not back up files nor folders in the target directory.

Example 5
        dsmc incr /home/gillis/

In example 5, because the symbolic link is not explicitly stated in the backup file specification, the client backs up only the path to the target directory. The client does not back up the symbolic link as a directory, and does not back up files nor folders in the target directory.

Restriction: If you back up a symbolic link as a directory, a future incremental backup that does not back up that symbolic link as a directory expires that symbolic link as a directory, and expires the files and directories in that directory.

For example, assume that you first back up the symbolic link `symdir` as a directory, and back up the contents of the target directory. The command in example 1 does this. The client creates backup copies with a high-level path `/home/gillis/symdir/`. In this example, the client creates backup copies with these paths:

```
/home/gillis/symdir/
/home/gillis/symdir/file
/home/gillis/symdir/dir1
/home/gillis/symdir/dir1/file1
```

The contents of `/home/gillis` are backed up using the following command:

```
dsmc inc /home/gillis/ -subdir=yes
```

This command processes the value `symdir` as a symbolic link and does not process any objects that the symbolic link points to. Hence, the client expires backup copies in the `/home/gillis/symdir/` directory that were created in example 1.

Incremental backup of a domain only

The client backs up a symbolic link during an incremental backup of the domain, if the symbolic link is defined as a virtual mount point and the followsymbolic option is set to yes.

The client backs up a symbolic link and the target directory when all of the following conditions are true:

- The client performs an incremental backup of the domain.
- The symbolic link is defined as a virtual mount point using the virtualmountpoint option.
- `followsymbolic=yes`

The virtualmountpoint and followsymbolic options add the symbolic link to the domain. The **Incremental** command backs up the domain, which includes the symbolic link target.

Related reference

"Followsymbolic" on page 336

"Virtualmountpoint" on page 452

Hard links

When you back up files that are hard-linked, Tivoli Storage Manager backs up each instance of the linked file.

For example, if you back up two files that are hard-linked, Tivoli Storage Manager backs up the file data twice.

When you restore hard-linked files, Tivoli Storage Manager attempts to reestablish the links. For example, if you had a hard-linked pair of files, and only one of the hard-linked files is on your workstation, when you restore both files, they are hard-linked. The files are also hard-linked even if neither of the files exists at the time of restore, if both of the files are restored together in a single command. The one exception to this procedure occurs if you back up two files that are hard-linked and then break the connection between them on your workstation. If

you restore the two files from the server using the standard (or classic) restore process, Tivoli Storage Manager respects the current file system and not re-establish the hard link.

Important: If you do not back up and restore all files that are hard-linked at the same time, problems occur. To ensure that hard-linked files remain synchronized, back up all hard links at the same time and restore those same files together.

Sparse files

Sparse files do not have disk space allocated for every block in the whole address space, leading to holes within the file. Holes are detected by their content, which is always zeros, and these zeros take up space.

The default is to restore the sparse file without the holes, which would leave more free disk space. The Tivoli Storage Manager client detects sparse files during a backup operation and marks them as sparse on the Tivoli Storage Manager server.

Note: Sparse files do not apply to Mac OS X.

Tivoli Storage Manager backs up a sparse file as a regular file if client compression is off.

Related reference

"Compression" on page 285

"Makesparsefile" on page 364

NFS hard and soft mounts

When Tivoli Storage Manager connects a backup-archive client to an NFS file system, you can use either a hard mount or a soft mount.

Tivoli Storage Manager uses the nfstimeout option setting to determine how long to wait for an NFS system call to respond before timing out; this applies to hard and soft mounts. The default is 0 seconds. This means that Tivoli Storage Manager uses the default behavior of NFS system calls.

You should be aware of the consequences of hard and soft mounts if the mount becomes stale (for example, if the server for the file system is not available).

Hard mount
— If the NFS file system is hard mounted, the NFS daemons try repeatedly to contact the server. The NFS daemon retries will not time out, they affect system performance, and you cannot interrupt them, but control returns to Tivoli Storage Manager when the nfstimeout value is reached.

Soft mount
— If the NFS file system is soft mounted, NFS tries repeatedly to contact the server until either:
- A connection is established
- The NFS retry threshold is met
- The nfstimeout value is reached

When one of these events occurs, control returns to the calling program.

Note: On Solaris and HP-UX, the nfstimeout option can fail if the NFS mount is hard. If a hang occurs, deactivate the nfstimeout option and mount the NFS file system soft mounted, as follows:

```
mount -o soft,timeo=5,retry=5 machine:/filesystem /mountpoint
```

The parameters are defined as follows:

soft   Generates a soft mount of the NFS file system. If an error occurs, the stat() function returns with an error. If the option hard is used, stat() does not return until the file system is available.

timeo=n
Sets the time out for a soft mount error to *n* seconds

retry=n
Set the internal retries and the mount retries to *n*, the default is 10000.

Deleted file systems

When a file system or drive has been deleted, or it is no longer backed up by Tivoli Storage Manager, the existing backup versions for each file are managed according to the following policy attributes: Number of days to keep inactive backup versions, and number of days to keep the last backup version (if there is no active version)

If you do nothing else, active backup versions remain indefinitely. If you do not need to keep the active versions indefinitely, use the **expire** command to inactive the active versions.

If you do not need to keep any of the backup versions, use the **delete backup** command to delete all backup versions in the file space. Your Tivoli Storage Manager server administrator must give you the authority to use this command. Use the **query session** command to determine whether you have "delete backup" authority. Alternatively, you can ask your Tivoli Storage Manager server administrator to delete the file space for you.

Related concepts

Chapter 9, "Storage management policies," on page 229

Opened files

Tivoli Storage Manager looks for files that have changed between the start and the completion of the backup of the file.

Some files on your system might be in use, or open, when you try to back them up. Because an open file can change, a backup action might not reflect the correct contents of the file at a given time.

Consider whether the file is important, and whether you can build the file again. If the file is not important, you might not want to back it up. Or, if the file is important, a root user on your workstation can ensure the file is closed before backup.

If your backups run on a schedule, a root user can use the preschedulecmd option to enter a command to close the file. For example, if the open file is a database, use the **quiesce** command of the database to shut down the database. A root user can use the postschedulecmd option to restart the application that uses the file after the backup completes. If you are not using a schedule for the backup, ensure that you close the application that uses the file before you start the backup.

Tivoli Storage Manager can back up the file even if it is open and gets changed during the backup. This is only useful if the file is usable even if it changes during backup. To back up these files, assign the files a management class with the serialization *dynamic* or *shared dynamic*.

Related concepts

"Display information about management classes and copy groups" on page 231

"Select a management class for files" on page 234

Wildcard characters

You can use the operating system wildcard characters in file specifications with Tivoli Storage Manager. These characters let you select groups of files that have similar names.

In a command, wildcard characters can only be used in the file name or extension. They cannot be used to specify destination files, file systems, or directories. When using wildcard characters in non-loop mode, as in `dsmc sel "/home/ledger.*"`, enclose the parameter containing the asterisk in quotation marks to ensure the system does not interpret the wildcard character and produce unexpected results. Wildcard character information is covered in the following table.

Important: Use an asterisk (*) instead of a question mark (?) as a wildcard character when trying to match a pattern on a multibyte code page, to avoid unexpected results.

This table shows some wildcard patterns and how to specify them.

* *(Asterisk)*	**Zero or more characters that match all files:**
*.cpp	With a cpp extension
hm*.*	Starting with hm, regardless of extension, but must have the **'.'** character
hm*	Starting with hm, whether an extension exists or not
h.*	With an h somewhere in the file name, regardless of extension, but must have **.**

? *(Question mark)*	**One character that matches all files with:**
?.cpp	The extension cpp with one, and only one, character in the file name
hm?.cpp	Three-character names beginning with hm and that have the cpp extension

* ? *(Asterisk and question mark)*	**Asterisk and question mark combinations matching:**
??hm.*	All four-character file names ending in hm., no matter what extension they have

In a path name for a file specification, you cannot specify a directory whose name contains an asterisk (*) or a question mark (?). Tivoli Storage Manager recognizes those characters only as wildcard characters.

Chapter 5. Restoring your data

Use Tivoli Storage Manager to restore backup versions of specific files, a group of files with similar names, or entire directories.

You can restore these backup versions if the original files are lost or damaged. Select the files you want to restore using file specification (file path, name, and extension), a directory list, or a subdirectory path to a directory and its subdirectories.

All client backup and restore procedures in this topic also apply to the Web client, except the following:
* Estimate
* Preferences editor

Note: When restoring a directory, its modification date and time is set to the date and time of the restore, not to the date and time the directory had when it was backed up. This is because Tivoli Storage Manager restores the directories first, then adds the files to the directories.

The following are the primary restore tasks:
* "Restoring an image"
* "Restoring data using the GUI" on page 193
* "Command line restore examples" on page 194
* "Restore data from a backup set" on page 179
* "Restoring data to a point in time" on page 187
* "Restore NAS file systems" on page 190
* "Authorizing another user to restore or retrieve your files" on page 198
* "Restoring or retrieving files from another client node" on page 199
* "Restore or retrieve files to another workstation" on page 200
* "Restoring a disk in case of disk loss" on page 200
* "Deleting file spaces" on page 201

Refer to *IBM Tivoli Storage Manager for Space Management for UNIX and Linux* for details about restoring migrated files and the **restoremigstate** option.

Related tasks

"Starting a Web client session" on page 111

Restoring an image

There are some items to consider before you begin restoring images on your system.

Before you restore an image (offline or online), you must have administrative authority on the system.

Here is a list of items to consider before you restore an image:
* Restoring the image of a volume restores the data to the same state that it was in when you performed your last image backup. Be absolutely sure that you

need to restore an image, because it replaces your entire current file system or raw volume with the image on the server.

- Ensure that the volume to which you are restoring the image is at least as large as the image that is being restored.
- The file system or volume you are restoring to must be the same type as the original.
- Ensure that the target volume of the restore is not in use. The client locks the volume before starting the restore. The client unlocks the volume after the restore completes. If the volume is in use when the client attempts to lock the file system, the restore fails.
- You cannot restore an image to where the Tivoli Storage Manager client program is installed.
- If you have run progressive incremental backups *and* image backups of your file system, you can perform an incremental image restore of the file system. The process restores individual files after the complete image is restored. The individual files restored are those backed up after the original image. Optionally, if files were deleted after the original backup, the incremental restore can delete those files from the base image.

  Deletion of files is performed correctly if the backup copy group of the Tivoli Storage Manager server has enough versions for existing and deleted files. Incremental backups and restores can be performed only on mounted file systems, not on raw logical volumes.
- If for some reason a restored image is corrupted, you can use the `fsck` tool to attempt to repair the image.

  You can use the verifyimage option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log.

  If bad sectors are present on the target volume, you can use the imagetofile option with the **restore image** command to specify that you want to restore the source image to a file. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume.

Related reference

"Imagetofile" on page 348
"Verifyimage" on page 451

Restoring an image using the GUI

You can use the GUI to restore an image of your file system or raw logical volume.

Follow these steps to restore an image of your file system or raw logical volume:

1. Click **Restore** from the main window. The Restore window appears.
2. Expand the directory tree.
3. Locate the object in the tree named **Image** and expand it. Click the selection box next to the image you want to restore. You can obtain detailed information about the object by highlighting the object and selecting **View** → **File Details...** from the main window or click the **View File details** button.
4. **(Optional)** To perform an incremental image restore, click the **Options** button to open the Restore Options window and select the **Image plus incremental directories and files** option. If you want to delete inactive files from your local file system, select the **Delete inactive files from local** check box. Click the **OK** button.

5. Click **Restore**. The Restore Destination window appears. The image can be restored to the volume with the mount point from which it was originally backed up. Alternatively, a different volume can be chosen for the restore location.

6. Click the **Restore** button to begin the restore. The **Task List** window appears showing the progress of the restore. The Restore Report window displays a detailed status report.

The following are some items to consider when you perform an image restore using the GUI:

- You can select **View → File Details** from the main window or click the **View File details** button to display the following statistics about file system images backed up by the client:
  - Image Size - This is the volume size which was backed up.
  - Stored Size - This is the actual image size stored on the server. The stored image on the Tivoli Storage Manager server is the same size as the volume capacity.
  - File system type
  - Backup date and time
  - Management class assigned to image backup
  - Whether the image backup is an active or inactive copy
- To modify specific restore options, click the **Options** button. Any options you change are effective during the current session *only*.
- In the Restore Options window, you can choose to restore the image only or the image and incremental directories files. If you choose **Image Only**, you restore the image from your last image backup only. This is the default.

  If you ran incremental-by-date image backup on a volume or image backups on a volume with incrementals, you can choose the **Image plus incremental directories and files** option. If you choose **Image plus incremental directories and files**, you can also select **Delete inactive files from local** to delete the inactive files that are restored to your local file system. If incremental-by-date image backup was the only type of incremental backup you performed on the file system, deletion of files will not occur.

  Important: Be absolutely sure that you need to perform an incremental restore because it replaces your entire file system with the image from the server and then restore the files that you backed up using the incremental image backup operation.

Restoring an image using the command line

Use the **restore image** command to restore an image using the Tivoli Storage Manager command line client.

Related reference

"Imagetofile" on page 348

"Verifyimage" on page 451

Restore data from a backup set

Your Tivoli Storage Manager administrator can generate a *backup set*, which is a collection of your files that reside on the server, onto portable media created on a device using a format that is compatible with the client device.

You can restore data from a backup set from the Tivoli Storage Manager server, or when the backup set is locally available as a file or on a tape device.

You can restore backup sets from the following locations:
- From the Tivoli Storage Manager server
- From portable media on a device attached to your client workstation
- From a backup set file on your client workstation

Backup sets can provide you with instant archive and rapid recovery capability as described in the following list.

Instant archive

    This capability allows an administrator to create an archive collection from backup versions already stored on the server.

Rapid recovery with local backup sets

    Typically, restores are performed from normal file backups that are stored on the Tivoli Storage Manager server outside of backup sets. This restore approach gives you the ability to restore the most recent backup version of every file. It is possible that a backup set does not contain the most recent backup version of your files.

    In some cases restoring data from a backup set can be a better option than restoring data from normal backup files on the Tivoli Storage Manager server. Restoring from a backup set can be a better option for the following reasons:
- A backup set restore can provide for a faster recovery because all of the required files for restore are contained together within a smaller number of storage volumes.
- A backup set provides a point-in-time collection of files. You can restore to a point in time rather than restoring what is currently available from a normal file-level restore from the server.

    Restoring a backup set from the Tivoli Storage Managerserver provides a larger set of restore options than restoring from a local backup set. However, restoring from a local backup set can be preferable in some cases:
- It is possible that you need to restore your data when a network connection to the Tivoli Storage Manager server is not available. This is possible in a disaster recovery situation.
- The local restore may be faster than restoring over a network connection to your Tivoli Storage Manager server.

A backup set can be restored from the Tivoli Storage Manager server while the backup set volumes are available to the server, or they can be moved to the client system for a local backup set restore. A backup set can be generated with or without a table of contents (TOC), and can contain file data or image data.

Your ability to restore data from backup sets is restricted by the location of the backup set and the type of data in the backup set. The command-line client can restore some data that the GUI cannot restore, but the GUI can allow you to browse and choose which objects to restore. Generally, backup sets from the server with a TOC allow more options when restoring. However, local backup sets provide options that are sometimes preferable to restoring from theTivoli Storage Manager server.

The restrictions for restoring data from backup sets using the GUI are summarized in the following table. Each interior cell represents one combination of data type and backup set location. For each situation, the cell indicates if you can use the GUI to restore only the entire backup set, to select objects within the backup set, or if you cannot use the GUI to restore the backup set.

Table 59. Backup set GUI restore restrictions

Data type in the backup set	Backup set location		
	Local (`location=file` or `location=tape`)	Tivoli Storage Manager Server (TOC available)	Tivoli Storage Manager Server (TOC not available)
file	Restore entire backup set only.	Restore entire backup set, or selected objects in the backup set.	Restore entire backup set only.
image	Cannot be restored.	Restore entire backup set, or selected objects in the backup set.	Cannot be restored.
system state or system object.	Restore entire backup set only.	Restore entire backup set, or selected objects in the backup set.	Restore entire backup set only.

The restrictions for restoring data from backup sets using the command-line client are summarized in the following table. Each interior cell represents one combination of data type and backup set location. For each situation, the cell lists the restore commands you can use. Except as noted, you can restore specific objects within a backup set, as well as the entire backup set.

Table 60. Backup set command-line restore restrictions

Data type in the backup set	Backup set location		
	Local (`location=file` or `location=tape`)	Tivoli Storage Manager Server (TOC available)	Tivoli Storage Manager Server (TOC not available)

file	Commands:
	`restore`
	`restore backupset`
		Commands:
		`restore`
		`restore backupset`
			Commands:
			`restore backupset`

image	Cannot be restored	Command: `restore image`	Cannot be restored

system state or system object	Command:
	`restore backupset`
		Commands:
		`restore backupset`
		`restore systemobject`
		`restore systemstate`
			Command:
			`restore backupset`

Restriction: When restoring system objects or system state using the **restore backupset** command, you cannot specify individual objects. You can only restore the entire system state or the entire system objects filespace

Restore backup sets: considerations and restrictions

This topic lists some considerations and restrictions that you must be aware of when restoring backup sets.

Backup set restore considerations

Consider the following when restoring backup sets:

- If the object you want to restore was generated from a Tivoli Storage Manager node whose name is different from your current node, specify the original node name with the **filespacename** parameter on any of the restore commands.
- If you are unable to restore a backup set from portable media, check with your Tivoli Storage Manager administrator to ensure that the portable media was created on a device using a compatible format.
- If you use the **restore backupset** command on the initial command line with the parameter -location=tape or -location=file, the client does not attempt to contact the Tivoli Storage Manager server.
- When restoring a group from a backup set:
  - The entire group, or all groups, in the virtual file space are restored. You cannot restore a single group by specifying the group name, if there are several groups in the same virtual file space. You cannot restore a part of a group by specifying a file path.
  - Specify a group by using the following values:
    - Specify the virtual file space name with the **filespacename** parameter.
    - Use the subdir option to include subdirectories.
- Limited support is provided for restoring backup sets from tape devices attached to the client system. A native device driver provided by the device manufacturer must always be used. The device driver provided by Tivoli to be used with the Tivoli Storage Manager server cannot be used on the client system for restoring local backup sets.
- If a backup set contains files from several owners, the backup set itself is owned by the root user ID, and non-root user IDs cannot see the backup set. In this case, non-root user IDs can restore their files by obtaining the backup set name from the Tivoli Storage Manager administrator. Non-root users can restore only their own files.
- To enable the client GUI to restore a backup set from a local device, without requiring a server connection, use the localbackupset option.

Backup set restore restrictions

Be aware of the following restrictions when restoring backup sets:

- A backup set data that was backed up with the API cannot be restored or used.

- You cannot restore image data from a backup set using the **restore backupset** command. You can restore image data from a backup set only with the **restore image** command.
- You cannot restore image data from a local backup set (`location=tape` or `location=file`). You can restore image data from a backup set only from the Tivoli Storage Manager server.

Related information

"Localbackupset" on page 363

"Restore" on page 525

"Restore Image" on page 539

"Restore Backupset" on page 529

Backup set restore

Tivoli Storage Manager considers a backup set as one object containing the whole file structure. You can restore the entire backup set or, in some cases, you can select portions. The backup set media is self-describing and contains all the information required to perform a successful restore.

If you are connected to a Tivoli Storage Manager Version 5.4 or later server, your Tivoli Storage Manager administrator can create backup sets that are stacked. Stacked backup sets can contain data from multiple client nodes, and they can contain different types of data for a particular client node. The types of data can be file data or image data.

Restriction: Image data and application data restore processing is only available when restoring from the Tivoli Storage Manager server. You cannot restore image data and application data from a client local backup set restore.

When a backup set is stacked, you can only restore data for your own node. Data for all other nodes is skipped. When restoring data from a stacked backup set on a local device, you can only restore file level data for your own client node. It is important that the nodename option is set to match the node name used to generate the backup set for one of the nodes in the stack.

Important: Due to the portability of local backup sets, you must take additional steps to secure your local backup sets on portable media. The backup set media should be physically secured because the backup set can be restored locally without authenticating with the server. Each user has access to all of the data on the stacked backup set, which means that the user has access to data that they do not own, by changing the node name or viewing the backup set in its raw format. Encryption or physical protection of the media are the only methods to ensure that the data is protected.

If you restore backup set data from the server, individual files, directories or entire backup set data can be restored in a single operation from the GUI or the command line. When you restore backup set data locally, the GUI can only display and restore an entire backup set. The command line can be used to restore individual files or directories stored in a backup set locally.

Restoring backup sets using the GUI

The client GUI can restore data from a backup set from the server, from a local file, or from a local tape device. You can use the GUI to restore individual files from a backup set from the Tivoli Storage Manager server with a TOC, but not from a local backup set nor from a backup set from the server without a TOC.

Important: Before you begin a restore operation, be aware that backup sets can contain data for multiple file spaces. If you specify a destination other than the original location, data from *all* file spaces are restored to the location you specify.

To restore a backup set from the GUI, perform the following steps:

1. Click **Restore** from the GUI main window. The Restore window appears.
2. Locate the **Backup Sets** directory tree object and expand it by clicking the plus sign (+) beside it.
   - To restore the backup set from a local device, expand the **Local** object and the Specify backup set location window is displayed. On the window, select **File name:** or **Tape name:** from the list and enter the tape or file name location. You can also click the **Browse** button to open a file selection window and select a backup set.
   - To restore data from backup set from the server, first expand the **Server** object and then either **Filelevel** or **Image**, depending on the type of restore requested.
3. Click the selection box next to the backup set or directory or file within the backup set that you want to restore.

   You can select files from within a backup set if that backup set is from the server and has a table of contents.
4. Click **Restore**. The Restore Destination window appears. Enter the appropriate information.
5. Click **Restore**. The Task List window displays the restore processing status.

Note:
- If the object you want to restore is part of a backup set generated on a node, and the node name is changed on the server, any backup set objects that were generated prior to the name change will not match the new node name. Ensure that the node name is the same as the node for which the backup set was generated.
- The client can be used to restore a backup set on an attached device with or without a server connection. If the server connection fails, a prompt appears to continue for purposes of local backup set restore. Also, thelocalbackupset option can be used to tell the client not to attempt the connection to the server.
- Certain local devices such as tape devices (tape devices do not apply to Mac OS X) require device drivers to be set up prior to performing a restore. See the device manual for assistance with this task. You also need to know the device address in order to perform the restore.
- The following features of a backup set restore from the server are not available when restoring locally:
  1. Image restore.
  2. The GUI display and restore of individual files and directories. The command line can be used to restore an individual directory or file from a local backup set.

Backup set restores using the client command-line interface

The client command line interface can restore data from a backup set from the server, from a local file, or from a local tape device. You can use the client command line interface to restore individual files from local backup sets and from backup sets without a TOC.

To restore a backup set from the client command line interface, use the **query backupset** command to display what backup set data is available, then use restore commands to restore the data.

You can use the commands below to restore data from backup sets. Use the appropriate command for the location of the backup set and the data in the backup set. (See the table in the parent topic)

> **restore**
>
> **restore backupset**
>
> **restore image**

Related reference

"Query Backupset" on page 508

"Query Image" on page 515

"Restore" on page 525

"Restore Backupset" on page 529

"Restore Image" on page 539

Restore an image to file

When you back up an image, Tivoli Storage Manager backs up the first sector of the volume, but when the data is restored, it skips the first sector to preserve the original logical volume control block of the destination volume.

When you restore an image to file, entire volume contents, including the first sector, are restored to the file.

AIX LVM volumes from original volume groups contain the Logical Volume Control Block (LVCB) on the first sector (512 bytes) of the volume. The LVCB contains volume specific meta-data that should be preserved by applications using the volume.

When you copy the file, containing the image, onto an LVM volume from the original volume group, you need to skip the LVCB from both the file and destination volume. The following **dd** command can be used for this purpose.

```
dd if=<filename> of=/dev/<vol> bs=512 skip=1 seek=1
```

The **dd** command sets the block size to 512 bytes, which makes copying very slow. It is better to use bs=1m or similar. Here is an alternative way to copy image data:

1. Save the original first sector to a file:

   ```
   dd if=/dev/<vol> of=firstblk.tmp bs=512 count=1
   ```

2. Copy the restored image:

   ```
   dd if=<filename> of=/dev/<vol> bs=1m
   ```

3. Restore the original first sector:

   ```
   dd if=firstblk.tmp of=/dev/<vol> bs=512 count=1
   ```

With the introduction of big and scalable volume group formats on AIX, it is possible that the first sector of the logical volume cannot contain LVCB and is available for the data. If you use big or scalable volume groups on your system, and need to restore the whole volume including the first sector, restore the volume to file and then copy it to a destination volume. The following **dd** command can be used for this purpose.

```
dd if=<filename> of=/dev/<vol> bs=1m
```

Related concepts

"Restoring an image using the command line" on page 179

Related tasks

"Restoring an image using the GUI" on page 178

Manage GPFS file system data with storage pools

GPFS 3.1 and 3.2 allow you to manage your data using storage pools. A storage pool is a collection of disks or RAIDs with similar properties that are managed together as a group.

The group under which the storage pools are managed together is the file system. The automated placement and management of files on the storage pool level is done by policies. A policy is a set of rules that describes the life cycle of user data, based on the attributes of the file.

When a file is created, the placement policy determines the initial location of the data of the file and assigns the file to a storage pool. All data written to that file is placed in the assigned storage pool. The management policy determines file management operation, such as migration and deletion. The files within a GPFS 3.1 or 3.2 file system are distributed over different storage pools, depending on the enabled placement and migration policies.

During restore, the files are placed on the correct storage pool. Tivoli Storage Manager is not aware of pool-to-pool migrations, so the files are placed on the storage pool from where the backup has taken place. The policy engine replaces the files based on migration policies.

GPFS 3.1 handles the placement of files after a restore as follows:
- The file is placed according to the storage pool ID that is stored in the extended attributes
- The file is placed based on the current placement policy
- The file is placed in the system pool

If a storage pool ID is stored in the extended attributes of the file, and that storage pool is available, the file is always placed in that storage pool. If the storage pool is not available, the file is placed according to the placement policy. If the placement policy does not match the file, the file is placed in the system pool.

GPFS 3.2 handles the placement of files after a restore as follows:
- The file is placed in the pool that can be selected by matching the saved file attributes to a RESTORE rule
- The file is placed in the pool that it was in when it was backed up
- The file is placed based on the current placement policy
- The file is placed in the system storage pool

The GPFS RESTORE rule allows you to match files against their saved attributes rather than the current file attributes. If the file attributes do not match, GPFS tries to restore the file in the sequence described above.

For more information about the GPFS RESTORE rule, read the GPFS documentation about policies and rules.

The following restrictions apply:
- Correct placement to storage pools only works with GPFS 3.1 PTF 6 or higher
- The restore of stub files does not work with multiple storage pools, or with files that have ACLs
- Filesets are not supported
- Unlink of filesets are not allowed
- The ctime option of GPFS should be set to *no* (default), to prevent unwanted Backup-Archive backups of files after GPFS file migration from pool to pool

See the *IBM Tivoli Storage Manager Administrator's Guide* for your operating system, for more information about storage pools.

Restoring data to a point in time

Use a *point-in-time* restore to restore files to the state that existed at a specific date and time.

A point-in-time restore can eliminate the effect of data corruption by restoring data from a time prior to known corruption, or recover a basic configuration to a prior condition.

You can perform a point-in-time restore of a file space, directory, or file.

You can also perform a point-in-time restore of image backups.

Perform incremental backups to support a point-in-time restore. During an incremental backup, the client notifies the server when files are deleted from a client file space or directory. Selective and incremental-by-date backups do not notify the server about deleted files. Run incremental backups at a frequency consistent with possible restore requirements.

If you request a point-in-time restore with a date and time that is prior to the oldest version maintained by the Tivoli Storage Manager server, the object is not restored to your system. Files which were deleted from your workstation prior to the point-in-time specified will not be restored.

Note:
1. Your administrator must define copy group settings that maintain enough inactive versions of a file to guarantee that you can restore that file to a specific date and time. If enough versions are not maintained, Tivoli Storage Manager might not be able to restore all objects to the point-in-time you specify.
2. If you delete a file or directory, the next time you run an incremental backup, the active backup version becomes inactive and the oldest versions that exceed the number specified by the *versions data deleted* attribute of the management class are deleted.

When performing a point-in-time restore, consider the following:

- Tivoli Storage Manager restores file versions from the most recent backup before the specified point-in-time date. Ensure the point-in-time that you specify is not the same as the date and time this backup was performed.
- If the date and time you specify for the object you are trying to restore is earlier than the oldest version that exists on the server, Tivoli Storage Manager cannot restore that object.
- Point-in-time restore restores files deleted from the client workstation after the point-in-time date but not files deleted before this date.
- Tivoli Storage Manager cannot restore a file created after the point-in-time date and time. When a point-in-time restore runs, files that were created on the client after the point-in-time date are not deleted.

To perform a point-in-time restore using the client GUI, use the following steps:

1. Click the **Restore** button in the main window. The Restore window appears.
2. Click the **Point-in-Time** button from the Restore window. The Point in Time Restore window appears.
3. Select the **Use a Point-in-Time Date** selection box. Select the date and time and click **OK**. The point in time that you specified appears in the Point in Time display field in the Restore window.
4. Display the objects you want to restore. You can search for an object by name, filter the directory tree, or work with the directories in the directory tree.
5. Click the selection boxes next to the objects you want to restore.
6. Click the **Restore** button. The Restore Destination window is displayed. Enter the appropriate information.
7. Click the **Restore** button to start the restore. The Restore Task List window displays the restore processing status.

Note: If there are no backup versions of a directory for the point-in-time you specify, files within that directory are not restorable from the GUI. However, you can restore these files from the command line.

You can start point-in-time restore from the command-line client using the pitdate and pittime options with the **query backup** and **restore** commands. For example, when you use the pitdate and pittime options with the **query backup** command, you establish the point-in-time for which file information is returned. When you use pitdate and pittime with the **restore** command, the date and time values you specify establish the point-in-time for which files are returned. If you specify pitdate without a pittime value, pittime defaults to 23:59:59. If you specify pittime without a pitdate value, it is ignored.

Related concepts

Chapter 9, "Storage management policies," on page 229

Related reference

"Backup Image" on page 473

Restore AIX encrypted files

When files are backed up in raw format from an AIX JFS2 Encrypted File System (EFS), you can only restore them to the same or another JFS2 EFS. They cannot be restored to any different file system, or on a different platform.

When EFS files are backed up in clear text, then you can restore them anywhere. If you restore them to a JFS2 EFS, they are automatically re-encrypted only if the directory to which they are restored has the AIX "EFS inheritance" option set.

After restoring a file that was backed up in raw format, you might find that the file cannot be decrypted. The encryption key originally used for the file might no longer be available in the keystore of the user. In this case, you must restore the keystore used at the time of backup. Restore it to a location different from /var/efs, and use the AIX command "efskeymgr -f " to add its keys to the kernel key ring.

For information on backing up EFS data, refer to "AIX JFS2 encrypted file system backup" on page 165.

Restore AIX workload partition file systems

All the files created by the local workload partition (WPAR) and backed up by the Tivoli Storage Manager client installed at the global WPAR can be restored by the Tivoli Storage Manager client installed at the global WPAR.

Here are some global partition and WPAR configuration examples:

```
Global partition:

   system name: shimla

   file system: /home /opt

WPAR #1 configuration:

   name: wpar1

   file system: /home; name in global WPAR: /wpars/wpar1/home

WPAR #2 configuration:

   name: wpar2

   file system: /data; name in global WPAR: /wpars/wpar2/data
```

There are two ways to restore WPAR data, depending on the method used to back up the WPAR data files:
- Restore all WPAR file systems as the file spaces within the global partition. The file space name must be used to identify the WPAR to which it belongs. All of the data is managed on one node using one schedule. Using the example configuration mentioned previously, here is a sample dsm.sys file with one server stanza for all file systems, both global and local:

```
SErvername   shimla

   TCPPort           1500

   TCPServeraddress  clinton.storage.sanjose.ibm.com

   nodename          shimla

   PasswordAccess    generate

   Domain            /wpars/wpar1/home /wpars/wpar2/data /home /opt
```
Use the following command to restore each file space:

```
dsmc restore /wpars/wpar1/home/*
dsmc restore /wpars/wpar2/data/*
dsmc restore /home/*
dsmc restore /opt/
```

- Restore each WPAR file system from a different node name, if it is backed up under a different node name. Each WPAR must have a separate node name and a scheduler running within the global partition. Also, three scheduler services must be set up, each using a different dsm.opt file corresponding to the server stanza name. This method allows each WPAR restore operation to be managed independent of the others. Using the example configuration mentioned previously, here is a sample dsm.sys file with three server stanzas: one for wpar1, one for wpar2, and one for global partition shimla:

```
SErvername  shimla_wpar1
  TCPPort       1500
  TCPServeraddress  clinton.storage.sanjose.ibm.com
  nodename      wpar1
  PasswordAccess  generate
  Domain      /wpars/wpar1/home

SErvername  shimla_wpar2
  TCPPort       1500
  TCPServeraddress  clinton.storage.sanjose.ibm.com
  nodename      wpar2
  PasswordAccess  generate
  Domain      /wpars/wpar2/data

SErvername  shimla
  TCPPort       1500
  TCPServeraddress  clinton.storage.sanjose.ibm.com
  nodename      shimla
  PasswordAccess  generate
  Domain      /home /opt
```

Table 61. Sample WPAR restore commands with dsm.opt file

In dsm.opt file	Sample restore command
servername shimla_wpar1	dsmc restore /wpars/wpar1/home/*
servername shimla_wpar2	dsmc restore /wpars/wpar2/data/*
servername shimla	dsmc restore /home/* dsmc restore /opt/*

Restore NAS file systems

You restore NAS file system images using the Web client or command line interface.

You can restore full or differential NAS file system images that were backed up previously. If you restore a differential image, Tivoli Storage Manager automatically restores the full backup image first, followed by the differential image. It is not necessary for a client node to mount a NAS file system to perform backup or restore operations on that file system.

Related tasks
"Configuring the Web client" on page 58

Restoring NAS file systems using the Web client

This section lists the steps to follow to restore NAS file systems using the Web client GUI.

1. Click the **Restore** button from the main window. The Restore window appears.
2. Expand the directory tree if necessary. To expand a node in the tree, click the plus sign (+) next to an object in the tree. Nodes shown are those that have been backed up and to which your administrator has authority. The root node called **Nodes** is not selectable. This node only appears if a NAS plug-in is present on the client workstation. NAS nodes display on the same level as the node of the client workstation. Only nodes to which the administrator has authority appear.
3. Expand the NAS node to reveal the Image object.
4. Expand the Image object to display volumes that you can restore. You cannot expand Volume objects.
5. Click the selection boxes next to the volumes under the Image object that you want to restore. If you want to restore a NAS image that was backed up on a particular date, click the **Point In Time** button. After you select a date, the last object that was backed up on or prior to that date appears, including any inactive objects. If you want to display all images (including active images and inactive images), before you select them, select **View** → **Display active/inactive files** from the menu bar.
6. Click **Restore**. The Restore Destination window appears. Enter the information in the Restore Destination window. If you choose to restore to a different destination, you can only restore one volume at a time to a different destination. You can restore NAS file system images to any volume on the NAS file server from which they were backed up. You cannot restore images to another NAS file server.
7. Click **Restore**. The NAS Restore **Task List** window displays the restore processing status and progress bar. If there is a number next to the progress bar, it indicates the size of the restore, if known. After the restore completes, the NAS Restore Report window displays processing details. If you must close the Web browser session, current NAS operations continue after you disconnect. You can use the **Dismiss** button on the NAS Restore **Task List** window to quit monitoring processes without ending the current operation.
8. (Optional) To monitor processing of an operation, select the **Actions** → **TSM Activities** from the main window.

Considerations:
- Workstation and remote (NAS) backups are mutually exclusive in a Restore window. After selecting an item for restore, the next item you select must be of the same type (either NAS or non NAS).
- Details will not appear in the right-frame of the Restore window for NAS nodes or images. To view information about a NAS image, highlight the NAS image and select **View** → **File Details** from the menu.
- To delete NAS file spaces, select **Utilities** → **Delete Filespaces**. You can delete both workstation and remote objects.

Options and commands to restore NAS file systems from the command line

This topic lists some examples of options and commands you can use to restore NAS file system images from the command line.

Table 62. NAS options and commands

Option or command	Definition	Page
query node	Displays all the nodes for which a particular administrative user ID has authority to perform operations. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the Web client.	"Query Node" on page 518
query backup	Use the **query backup** command with the class option to display information about file system images backed up for a NAS file server.	"Query Backup" on page 506
query filespace	Use the **query filespace** command with the *class* option to display a list of file spaces belonging to a NAS node.	"Query Filespace" on page 512
restore nas	Restores the image of a file system belonging to a Network Attached Storage (NAS) file server.	"Restore NAS" on page 542
monitor process	Displays current backup and restore processes for all NAS nodes for which an administrative user has authority. The administrative user can then select one process to monitor.	"Monitor Process" on page 500
cancel process	Displays current backup and restore processes for all NAS nodes for which an administrative user has authority. From the display, the administrative user can select one process to cancel.	"Cancel Process" on page 481
delete filespace	Use the **delete filespace** with the class option to display a list of file spaces belonging to a NAS node so that you can choose one to delete.	"Delete Filespace" on page 487

Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: `/vol/vol0`.

Note: When you initiate a NAS restore operation using the command line client or the Web client, the server starts a process to initiate, control, and monitor the operation. It might take several moments before you notice progress at the client command line interface because the server must perform a mount and other necessary tasks before data movement occurs. The Tivoli Storage Manager command line client might display an `Interrupted ...` message when the mount occurs. You can ignore this message.

Restore active or inactive backups

Your administrator determines how many backup versions Tivoli Storage Manager maintains for each file on your workstation.

Having multiple versions of a file permits you to restore older versions if the most recent backup is damaged. The most recent backup version is the *active* version. Any other backup version is an *inactive* version.

Every time Tivoli Storage Manager backs up your files, it marks the new backup version as the active backup, and the last active backup becomes an inactive backup. When the maximum number of inactive versions is reached, Tivoli Storage Manager deletes the oldest inactive version.

To restore a backup version that is inactive, you must display both active and inactive versions by clicking on the **View** menu → **Display active/inactive files** item. To display only the active versions (the default), click on the **View** menu → **Display active files only** item. If you try to restore more than one version at a time, only the active version is restored.

On the Tivoli Storage Manager command line, use the inactive option to display both active and inactive objects.

Related reference

"Inactive" on page 348

Restoring data using the GUI

This section lists the steps to follow to restore backup versions of individual files or subdirectories.

1. Click **Restore** from the main window. The Restore window appears.
2. Expand the directory tree. Select the selection boxes next to the files or directories you want to restore. To search or filter files, click the **Find** icon on the tool bar.
3. Enter your search criteria in the Find Files (Restore) window.
4. Click the **Search** button. The Matching Files (Restore) window appears.
5. Click the selection boxes next to the files you want to restore and close the Matching Files (Restore) window.
6. Enter your filter criteria in the Find Files (Restore) window.
7. Click the **Filter** button. The Restore window displays the filtered files.
8. Click the selection boxes next to the filtered files or directories you want to restore.
9. To modify specific restore options, click the **Options** button. Any options you change are effective during the current session *only*.
10. Click **Restore**. The Restore Destination window appears. Enter the information in the Restore Destination window.
11. Click **Restore**. The Restore **Task List** window displays the restore processing status.

Note: On Mac OS X, consider the following items when restoring data using the GUI:

1. When **TSM Tools for Administrators** is used to start the client, the client is running with a UID of zero. This means that if you create a folder to restore

your files to, that folder is owned by root. To access the files you must change the permissions of the folder. You can change the folder owner from a terminal window using the sudo **chown** command. See your operating system documentation for more information on how to accomplish this.

2. When restoring files with the replace option set to *no*, existing files will not be overwritten, but existing directories are overwritten. To leave existing directories intact during a restore operation, select the **Options** button ⇒ **All selected files and directories** dropdown menu ⇒ **Files only** option.

3. When folders are restored from a UFS or HFSX file system to a HFS file system and they differ only in case, the client restores the contents of both folders to one folder.

Command line restore examples

This topic lists some examples of **restore** commands to use for specific tasks.

The following table shows examples of how to use the **restore** command to restore objects from Tivoli Storage Manager server storage.

Table 63. Command-line restore examples

Task	Command	Considerations
Restore the most recent backup version of the /Users/monnett/ Documents/h1.doc file, even if the backup is inactive.	dsmc restore /Users/monnett/ Documents/h1.doc -latest	If the file you are restoring no longer resides on your workstation, and you have run an incremental backup since deleting the file, there is no active backup of the file on the server. In this case, use the latest option to restore the most recent backup version. Tivoli Storage Manager restores the latest backup version, whether it is active or inactive. See "Latest" on page 362 for more information.
Display a list of active and inactive backup versions of files from which you can select versions to restore.	dsmc restore "/Users/monnett/ Documents/*"-pick -inactive	If you try to restore both an active and inactive version of a file at the same time, only the active version is restored. See "Pick" on page 382 and "Inactive" on page 348 for more information.
Restore the /Users/monnett/ Documents/h1.doc file to its original directory.	dsmc restore /Users/monnett/ Documents/h1.doc	If you do not specify a destination, the files are restored to their original location.
Restore the /Users/monnett/ Documents/h1.doc file under a new name and directory.	dsmc restore /Users/monnett/ Documents/h1.doc /Users/gordon/Documents/h2.doc	None
Restore the files in the /Users directory and all of its subdirectories.	dsmc restore /Users/ -subdir=yes	When restoring a specific path and file, Tivoli Storage Manager recursively restores *all* subdirectories under that path, and any instances of the specified file that exist under *any* of those subdirectories. See "Subdir" on page 432 for more information about the *subdir* option.

Table 63. Command-line restore examples (continued)

Task	Command	Considerations

Restore all files in the /Users/gordon/Documents directory to their state as of 1:00 PM on August 17, 2003.	`dsmc restore -pitd=8/17/2003`
	`-pitt=13:00:00 /Users/gordon/`
	`Documents/`
		See "Pitdate" on page 383 and "Pittime" on page 384 for more information about the pitdate and pittime options.

Restore all files from the /Users/mike/Documents directory that end with `.bak` to the /Users/mike/projectn/ directory.	`dsmc restore "/Users/mike/`
	`Documents/*.bak"`
	`/Users/mike/projectn/`
		If the destination is a directory, specify the delimiter (/) as the last character of the destination. If you omit the delimiter and your specified source is a directory or a file spec with a wildcard, you receive an error. If the `projectn` directory does not exist, it is created.

Restore files specified in the `restorelist.txt` file to a different location.	`dsmc restore -filelist=/Users/`
	`user2/Documents/ restorelist.txt`
	`/Users/NewRestoreLocation/`
		See "Filelist" on page 333 for more information about restoring a list of files.

Related reference

"Restore" on page 525

Examples: Command line restores for large amounts of data

If you need to restore a large number of files, you can get faster performance by using the **restore** command instead of the GUI. In addition, you can improve performance by entering multiple **restore** commands at one time.

For example, to restore all the files in your /home file system, enter:

```
dsmc restore /home/ -subdir=yes -replace=all -tapeprompt=no
```

However, if you enter multiple commands for the directories in the /home file space, you can restore the files faster.

For example, you could enter these commands:

```
dsmc restore /home/monnett/ -subdir=yes -replace=all -tapeprompt=no
dsmc restore /home/gillis/ -subdir=yes -replace=all -tapeprompt=no
dsmc restore /home/stewart/ -subdir=yes -replace=all -tapeprompt=no
```

You can also use the quiet option with the **restore** commands to save processing time. However, you will not receive informational messages for individual files.

Note: If you already have the appropriate values set for the subdir, replace, tapeprompt, and quiet options in your client user-options file, you do not need to include those options in the commands.

When you enter multiple commands to restore your files, you must specify a unique part of the file space in each **restore** command. Be sure you do not use any overlapping file specifications in the commands.

To display a list of the directories in a file space, use the **query backup** command. For example:

```
dsmc query backup -dirsonly -subdir=no /Users/
```

As a general rule, you can enter from two to four **restore** commands at one time. The maximum number you can run at one time without degrading performance depends on factors such as how much memory you have and network utilization.

The speed at which you can restore the files also depends on how many tape drives are available on the server, and whether your administrator is using collocation to keep file spaces assigned to as few volumes as possible.

For example, if /Users/user1 and /Users/user2 are on the same tape, the restore for /Users/user2 must wait until the restore for /Users/user1 is complete. However, if /Users/user3 is on a different tape, and there are at least two tape drives available, the restore for /Users/user3 can begin at the same time as the restore for /Users/user1.

Set the system ulimit values to unlimited (-1) if you are restoring very large (2 GB) files with HSM or the backup-archive client. The Tivoli Storage Manager client can restore these large files with enough system resources. If the ulimits are set to lower values, there might be restore failures.

Standard query restore, no query restore, and restartable restore

This topic describes the standard (or classic) restore method, the no query restore method, and the restartable restore method.

Standard query restore process

The standard query restore process is also known as classic restore. This topic explains how standard query restore works.

Here is how standard query restore works:
- The client queries the server for a list of files backed up for the client file space you want to restore.
- The server sends a list of backed up files that match the restore criteria. If you want to restore both active and inactive files, the server sends information about all backed up files to the client.
- The list of files returned from the server is sorted in client memory to determine the file restore order and to minimize tape mounts required to perform the restore.
- The client tells the server to restore file data and directory objects.
- The directories and files you want to restore are sent from the server to the client.

No query restore process

The no query restore process is outlined below.
1. The client tells the server that a no query restore is going to be performed and provides the server with details about file spaces, directories, and files.
2. The server sorts the data using an internal sort table which minimizes tape mounts.
3. The data to be restored is sent to the client. File and directory objects stored on disk are sent immediately since sorting for such data is not required before restoring it.
4. You can use multiple sessions to restore the data. If the data resides on multiple tapes, there are multiple mount points available at the server. The combination of using the resourceutilization option and MAXNUMMP allows multiple sessions.

When you enter an unrestricted wildcard source file specification on the **restore** command and do not specify any of the options: inactive, latest, pick, fromdate,

todate, the client uses a *no query restore* method for restoring files and directories from the server. This method is called *no query restore* because instead of querying the server for each object to be restored, a single restore request is sent to the server. In this case, the server returns the files and directories to the client without further action by the client. The client merely accepts the data coming from the server and restores it to the destination named on the **restore** command.

Using the Tivoli Storage Manager GUI client, an example of an unrestricted wildcard command would be to select a folder from the restore tree window. An example of a restricted wildcard command would be to select individual files from a folder.

Using the command-line client, an example of an unrestricted wildcard command would be:
```
"/Users/user1/Documents/2004/*"
/home/mydocs/2004/*
```

An example of a restricted wildcard file specification would be:
```
/Users/user1/Documents/2004/sales.*
/home/mydocs/2004/sales.*
```

Restartable restore process

If the restore process stops because of a power outage or network failure, the server records the point at which this occurred.

This record is known to the client as a *restartable restore*. It is possible to have more than one restartable restore session. Use the **query restore** command or choose **restartable restores** from the Actions menu to find out if your client has any restartable restore sessions in the server database.

You must complete a restartable restore before attempting further backups of the file system. If you attempt to repeat the restore that was interrupted or try to back up the destination file space, the attempt fails because you did not complete the original restore. You can restart the restore at the point of interruption by entering the **restart restore** command, or you can delete the restartable restore using the **cancel restore** command.

From the Tivoli Storage Manager GUI **Restartable restores** dialog box you can select the interrupted restore and delete it, or you can choose to restart the restore. If you restart the interrupted restore, it restarts with the first transaction, which might consist of one or more files, not completely restored when the interruption occurred. Because of this, you might receive some replace prompts for files from the interrupted transaction which were already restored.

To perform restartable restores using the GUI, follow these steps:
1. Select **Actions –> Restartable restores** from the main panel.
2. Select the restartable restore session you want to complete.
3. Click the **Restart** button at the bottom of the panel.

Restoring Solaris Zettabyte (ZFS) file systems

Zettabyte File Systems (ZFS) use storage pools to manage physical storage.

You can back up and restore ZFS snapshots in one of the following ways:

1. Back up and restore each file of the snapshot by using the snapshotroot option. This option allows the administrator to replace the current snapshot path with the ZFS file system path.

   Using the snapshotroot option, the files can be directly restored in the file system. If you have many changes in the file system during a backup, a snapshot could be made first, and then the Tivoli Storage Manager Backup-Archive Client can make a backup of the static snapshot.

2. Another way is to back up and restore the complete snapshot. The administrator can then make a copy of the complete snapshot data in a single file, and this file can be backed up and restored by the Tivoli Storage Manager Backup-Archive Client to the Tivoli Storage Manager Server.

   The advantage of backing up and restoring the complete snapshot is that the full file system can be restored in a disaster recovery case.

Related tasks

Additional restore tasks

This section discusses some advanced considerations for restoring data.

Authorizing another user to restore or retrieve your files

You can authorize another user on the same workstation or a different workstation to restore backup versions or retrieve archive copies of your files.

This permits you to share files with other people or with other workstations that you use with a different node name. To authorize a user on another workstation to restore or retrieve your files, the other workstation must be running one of the UNIX clients and must be registered with your server.

Note: Mac OS X can *only* restore Mac OS X nodes.

To authorize another user to restore or retrieve your files:

1. Click **Utilities** → **Node Access List** from the main window. The Node Access List window appears.
2. Click the **Add** button. The Add Access Rule window appears.
3. In the Add Access Rule window, select an item in the Permit Access to field to specify the type of data that the other user can access. You can select either Backed up Objects or Archived Objects.
4. In the Grant Access to Node field, type the node name of the host workstation of the user that can access your data.
5. In the User field, type the name of the user on a node who can access your data.

6. In the Filespace and Directory field, select the file space and the directory that the user can access. You can select one file space and one directory at a time. If you want to give the user access to another file space or directory, you must create another access rule.

7. If you want to limit the user to specific files in the directory, type the name or pattern of the files on the server that the other user can access in the Filename field. You can make only one entry in the Filename field. It can either be a single file name or a pattern which matches one or more files. You can use a wildcard character as part of the pattern. Your entry must match files that have been stored on the server.

8. For the Java GUI: If you want to give access to all files that match the file name specification within the selected directory including its subdirectories, click **Include subdirectories**.

9. Click the **OK** button to save the access rule and close the Add Access Rule window.

10. The access rule that you created is displayed in the list box in the Node Access List window. When you have finished working with the Node Access List window, click the **OK** button. If you do not want to save your changes, click **Cancel** or close the window.

In the client command line interface, use the **set access** command to authorize another node to restore or retrieve your files. You can also use the **query access** command to see your current list, and **delete access** to delete nodes from the list.

Related reference

"Delete Access" on page 482

"Query Access" on page 503

"Set Access" on page 551

Restoring or retrieving files from another client node

After users grant you access to their files on the server, you can restore or retrieve those files to your local system.

You can display file spaces of another user on the server, restore the backup versions of another user, or retrieve the archive copies of another user to your local file system:

1. Click **Utilities** from the main window.

2. Click **Access Another Node**. The Access Another Node window appears.

3. Type the node name of the host workstation of the user in the Node name field. Type the user name in the User name field.

4. Click the **Set** button.

If you are using commands, use the fromnode and fromowner options to indicate the node name and the name of the user who owns the files.

For example, to restore files to one of your own file systems that were backed up from a workstation named Node1 and owned by a user named Ann, enter:

```
dsmc restore -fromn=node1 -fromo=ann "/home/proj/*" /home/gillis/
```

Use the **query filespace** command to get a list of file spaces. For example, to get a list of file spaces owned by Ann on Node1, enter:

```
dsmc query filespace -fromn=node1 -fromo=ann
```

Related reference

"Fromnode" on page 339
"Query Filespace" on page 512
"Restore" on page 525
"Retrieve" on page 544

Restore or retrieve files to another workstation

From a different workstation, you can restore or retrieve files you have already backed up from your own workstation. You must know the Tivoli Storage Manager password assigned to your node.

To restore or retrieve files to another workstation, use the virtualnodename option to specify the node name of the workstation from which you backed up the files. The virtualnodename option cannot be set to the hostname of the workstation. You can use the virtualnodename option when you start Tivoli Storage Manager or you can add the virtualnodename option to your client user options file dsm.opt. Use the virtualnodename option on the **dsmj** command if you are borrowing the workstation of another user and you do not want to update their client user-options file.

Tivoli Storage Manager prompts you for the password for your original node. After you enter the correct password, all file systems from your original workstation appear in the Restore or Retrieve window. You can restore or retrieve files as if you were working on your own workstation.

Important: When you use this method to access files, you have access to all files backed up and archived from your workstation. You are considered a virtual root user.

You can use the virtualnodename option in a command. For example, to restore your *projx* files, enter:

    dsmc restore -virtualnodename=nodeone "/home/monnett/projx/*"

If you do not want to restore or retrieve the files to the same directory name on the alternate workstation, enter a different destination.

The considerations for retrieving files are the same as restoring files.

Restoring a disk in case of disk loss

You can only recover your files if you can run the client. If the disk that contains the client is lost (from theft or hardware failure, for example), you must reinstall the client before you can recover your files. If you also lose the disk that contains the operating system and communication software, you must recover them before you can connect to the Tivoli Storage Manager server.

To protect yourself against these kinds of losses, you need to put together a set of installation media that you can use to restore your system to a state that lets you contact the server and begin recovering data. The installation media should contain:
1. A startable operating system that lets you perform basic functions.
2. A correctly configured communication program that lets you establish communications with the server.

3. A client with appropriate customized options files. You can use the client command line interface to complete this task.

The communication package you use determines what files you need. Consult your operating system and communication software manuals to set up your installation media.

If you also have the Tivoli Space Manager installed on your workstation, your installation media should include the space manager command line client. For information about restoring migrated files, see *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide*, SC32-0148.

Note: Your administrator can schedule restore operations which can be very useful when you need to restore a large number of files.

Deleting file spaces

If your Tivoli Storage Manager administrator gives you authority, you can delete entire file spaces from the server.

When you delete a file space, you delete all the files and images, both backup versions and archive copies, that are contained within the file space. For example, if you delete the file space for your /home/monnet file system, you are deleting every backup for every file in that file system and every file you archived from that file system. **Carefully consider whether you want to delete a file space**. You must be an authorized user to perform this task.

You can delete individual backup versions by using the **delete backup** command.

You can delete file spaces using the Tivoli Storage Manager client GUI or client command line interface. To delete NAS file spaces, use the Web client or client command line interface.

To delete a file space using the GUI, perform the following steps:
1. Select **Utilities→ Delete Filespaces** from the main window.
2. Click the selection boxes next to the file spaces you want to delete.
3. Click the **Delete** button. Tivoli Storage Manager prompts you for confirmation before deleting the file space.

You can also delete a file space using the **delete filespace** command. Use the *class* option with the **delete filespace** command to delete NAS file spaces.

Related reference
"Class" on page 279
"Delete Backup" on page 484
"Delete Filespace" on page 487

Enable SELinux to restore files on the Red Hat Enterprise Linux 4 or 5 client

If you are a non-root user, and you are trying to restore files on the Red Hat Enterprise Linux 4 or 5 client, you must first enable SELinux.

If you do not enable SELinux, you will have problems restoring files with modified extended attributes.

Chapter 6. Archive and retrieve your data (UNIX and Linux)

Archiving and retrieving files is similar to backing up and restoring files. Many of the windows and concepts are similar. In this section, the main archive and retrieve tasks are covered.

All client archive and retrieve procedures in this topic also apply to the Web client, except the following:
- Estimate
- Preferences editor

The following are the primary archive and retrieve tasks:
- "Archiving data using the GUI"
- "Archive data examples using the command line" on page 204
- "Deleting archive data" on page 207
- "Retrieving data using the GUI" on page 210
- "Retrieve data examples using the command line" on page 210

Related concepts

Chapter 4, "Backing up your data," on page 119

Related tasks

"Starting a Web client session" on page 111

Archive files

To archive files, you need to specifically select the files to archive. You can select the files by using a file specification or by selecting them from a directory tree.

Your administrator might have set up schedules to archive certain files on your workstation automatically. The following sections cover how to archive files without using a schedule.

Related tasks

"Set the client scheduler process to run as a background task and start automatically at startup" on page 218

Archiving data using the GUI

You can archive a file or a group of files using file names, or you can select files that match your search criteria using a directory tree.

Perform archives using the following procedure:

1. Click **Archive** from the main window. The Archive window appears.
2. Expand the directory tree by clicking the plus sign **(+)** or the folder icon next to an object in the tree. To search or filter files, click the **Search** icon from the tool bar.
3. Enter your search criteria in the Find Files (Archive) window.
4. Click the **Search** button. The Matching Files (Archive) window appears.
5. Click the selection boxes next to the files you want to archive and close the Matching Files (Archive) window.
6. Enter your filter criteria in the Find Files (Archive) window.

7. Click the **Filter** button. The Archive window displays the filtered files.

8. Click the selection boxes next to the filtered files or directories you want to archive.

9. Enter the description, accept the default description, or select an existing description for your archive package in the Description box. The maximum length of a description is 254 characters. When an existing archive description is used, the files or directories selected are added to the archive package. All archived packages with the same description are grouped for retrieves, queries, and deletions.

10. To modify specific archive options, click the **Options** button. Any options you change are effective during the current session *only*. The Java GUI does not have an **Estimate** button.

11. Click on **Archive**. The Archive **Task List** window displays the archive processing status.

Archive data examples using the command line

You request archive services when you want to preserve copies of files in their current state, either for later use or for historical or legal purposes. This topic lists some examples of archiving data using the command line.

You can archive a single file, a group of files, or all the files in a directory or subdirectory. After you archive a file, you can choose to delete the original file from your workstation.

The following table shows examples of using the **archive** command to archive objects.

Table 64. Command line archive examples

Task	Command	Considerations
Archive all files in the /home/proj1 directory with a file extension of .txt.	`dsmc archive "/home/proj1/*.txt"`	Use wildcards to archive more than one file at a time.
Archive all files in the /home/jones/proj/ directory and delete the files on your workstation.	`dsmc archive /home/jones/proj/ -deletefiles`	Retrieve the archived files to your workstation whenever you need them again. See "Deletefiles" on page 295 for more information about the deletefiles option.
Archive the /home/jones/h1.doc and /home/jones/test.doc files.	`dsmc archive /home/jones/h1.doc /home/jones/test.doc`	If you specify the removeoperandlimit option with the **archive** command, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits. This allows you to specify more than 20 files on a single command. See "Removeoperandlimit" on page 396 for more information about this option.
Archive a list of files in the /home/avi/filelist.txt file.	`dsmc archive -filelist=/home/avi/filelist.txt`	Use the filelist option to process a list of files. See "Filelist" on page 333 for more information.

Table 64. Command line archive examples (continued)

Task	Command	Considerations
Archive the /home/jones/ch1.doc file and assign a description to the archive.	`dsmc archive /home/jones/ch1.doc -description="Chapter 1, first version"`	If you do not specify a description with the **archive** command, the default is `Archive Date:x`, where x is the current system date. See "Description" on page 295 for more information about the description option.
Archive all of the files in the /home/jones/proj/ directory and its subdirectories.	`dsmc archive /home/jones/proj/ -subdir=yes`	See "Subdir" on page 432 for more information about the subdir option.
Use the v2archive option with the **archive** command to archive only files in the /home/relx/dir1 directory, but not the relx or dir1 directories.	`dsmc archive "/home/relx/dir1/" -v2archive`	Tivoli Storage Manager archives only files in the /home/relx/dir1 directory. Directories that exist in the path are not processed. See "V2archive" on page 449 for more information about the v2archive option.
Use the archmc option with the **archive** command to specify the available management class for your policy domain to which you want to bind your archived files.	`dsmc archive —archmc=ret2yrs /home/plan/proj1/budget.jan`	See "Archmc" on page 267 for more information about the archmc option. See Chapter 9, "Storage management policies," on page 229 for more information about management classes.
Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, archive the /usr/dir1/sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name /usr.	`dsmc archive /usr/dir1/sub1/ -subdir=yes -snapshotroot=/ snapshot/day1`	Tivoli Storage Manager considers the snapshotroot value as a file space name. See "Snapshotroot" on page 425 for more information.

Related reference

"Archive" on page 465

Associating a local snapshot with a server file space

To associate data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager, use the snapshotroot option.

Using the snapshotroot option with the **archive** command in conjunction with a vendor-acquired application that provides a snapshot of a logical volume you can associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server.

You cannot use the snapshotroot option to take a volume snapshot, but you can use the option to manage data created by a volume snapshot.

Related reference

"Snapshotroot" on page 425

Archiving data with client node proxy

Archives of multiple nodes that share storage can be consolidated to a common target node name on the Tivoli Storage Manager server.

This is useful when the workstation responsible for performing the archive can change over time, such as with a Xsan or cluster. The asnodename option also allows data to be restored from a different system than the one which performed the backup. Use the asnodename option with the appropriate command to back up, archive, restore, and retrieve data under the target node name on the Tivoli Storage Manager server. This support is only available with Tivoli Storage Manager Version 5.3 and higher.

Tivoli Storage Manager FastBack clients are also backed up using client node proxy.

To enable this option, follow these steps:
1. Install the backup-archive client on all nodes in a shared data environment.
2. Register each node with the Tivoli Storage Manager server, if it does not exist. Register the common target node name to be shared by each of the agent nodes used in your shared data environment.
3. Register each of the nodes in the shared data environment with the Tivoli Storage Manager server. This is the agent node name that is used for authentication purposes. Data is not stored using the node name when the asnodename option is used.
4. Grant proxy authority to all nodes in the shared environment to access the target node name on the Tivoli Storage Manager server, using the GRANT PROXYNODE command (Tivoli Storage Manager administrator).
5. Use the QUERY PROXYNODE administrative client command to display the client nodes of the authorized user, granted by the GRANT PROXYNODE command.

Follow these steps to set up encryption with the *encryptkey=save* option:
1. Specify *encryptkey=save* in the options file.
2. Back up at least one file with *asnode=ProxyNodeName* to create a local encryption key on each agent node in the multiple node environment.

Follow these steps to set up encryption with the *encryptkey=prompt* option:
1. Specify *encryptkey=prompt* in the options file.
2. Ensure that users of the agent nodes in the multiple node environment are using the same encryption key.
• If you change the encryption key, you must repeat the previous steps.
• Use the same encryption key for all files backed up in the shared node environment.

Follow these steps to enable multinode operation from the GUI:
1. Verify that the client node has proxy authority to a target node (or authorized to act as the target node) using the QUERY PROXYNODE administrative client command.
2. Select **Edit** → **Preferences** to open the preferences window.
3. Select the **General** tab and fill in the **As Node Name** field with the name of the proxy authorized target node.
4. Click **Apply** and then **OK** to close the preferences window.

Follow these steps to verify that your client node is now accessing the server as the target node:

1. Open the tree window and check that the target node name specified by the **As Node Name** field appears, or
2. Verify the target node name in the **Accessing As Node** field in the **Connection Information** window.

To return to single node operation, delete the **As Node Name** from the **Accessing As Node** field in the **General** → **Preferences** tab.

Tip:
- All agent nodes in the multiple node environment should be of the same platform type.
- Do not use target nodes as traditional nodes. Use them only for multiple node processing.

Restrictions enforced within a proxied session:
- You cannot perform a system object, system state or system services backup or restore.
- You cannot access another node (either from GUI drop down or use of the fromnode option).
- You cannot perform NAS backup or restore.

Related reference

"Asnodename" on page 268

Deleting archive data

If your administrator has given you authority, you can delete individual archive copies from the Tivoli Storage Manager server without deleting the entire file space.

To determine if you have this authority, select **File** → **Connection Information** from the Tivoli Storage Manager GUI or Web client main menu. Your authority status is provided in the **Delete Archive Files** field.

To delete archive copies from the Tivoli Storage Manager GUI or Web client:
1. Select **Delete Archive Data** from the **Utilities** menu. The Archive Delete window appears.
2. Expand the Directory tree by clicking the plus sign **(+)** or folder icon next to the object you want to expand. Objects on the tree are grouped by archive package description.
3. Click the selection boxes next to objects that you want to delete. If you want to estimate the amount of time it takes to process your files and directories, click the **Estimate** button.
4. Click the **Delete** button. Tivoli Storage Manager prompts you for confirmation before deleting archive selections. The Archive Delete Task List window displays the progress of the delete.

To delete archive copies using the Tivoli Storage Manager command line client, use the **delete archive** command.

Related reference

"Delete Archive" on page 483

Advanced archive tasks

Access permissions, symbolic links, and hard links are advanced functions to consider when you archive data.

Access permissions

When you archive a file, Tivoli Storage Manager saves standard UNIX access permissions assigned to the file.

Depending on your operating system, it also saves extended permissions. For example, for files on an AIX workstation, Tivoli Storage Manager saves access control lists.

If you are a user, and you archive a file to which you have read access, you own the archived copy of the file. You are the only user who can retrieve the archived file unless you grant access to another user.

Archive and retrieve symbolic links

Tivoli Storage Manager archives and retrieves symbolic links differently than it does regular files and directories.

The way that the client archives and retrieves symbolic links depends on options settings, whether the target directory is accessible, and the way you specify objects.

A *UNIX symbolic link* is a file that contains a pointer to another file or directory. The object the symbolic link points to is called the *target object*.

A symbolic link can be backed up as path information to a target directory, or it can be backed up as a directory. If the symbolic link is backed up as a directory, the files and folders in the target directory can also be backed up.

What you restore depends on how the symbolic link was backed up, the scope of the restore, the setting of the followsymbolic option, and whether the target directory is accessible at the time of restore.

For more information on how symbolic links are handled during archive, see the archsymlinkasfile option.

Note: Symbolic link processing as described here does not apply to Mac OS X. Symbolic links are always archived as files and are never followed.

The following table shows symbolic link archive and retrieve functions and the action taken:

Table 65. Symbolic link management table for archive and retrieve

Function	Action taken
Archive of a file link.	Archives the file to which the symbolic link points.
Archive of a directory link.	Archives the directory and its contents.
Archive of a file with subdir=yes.	Archives the file, directory path and all like-named files in the subtree.

Table 65. Symbolic link management table for archive and retrieve (continued)

Function	Action taken
Archive of a directory with subdir=yes.	Archives the directory, its contents, and contents of subdirectories.
Archive of a symbolic link that points to a file or directory that does not exist.	Archives the symbolic link.
Retrieve a symbolic link that points to file; the file and link exist.	Replaces the file if replace=y is set.
Retrieve a symbolic link that points to file; the symbolic link no longer exists.	Retrieves the file replacing the file name with the symbolic link name and places it in the directory where the symbolic link resided.
Retrieve a symbolic link that points to a directory; the symbolic link and directory no longer exist.	A directory is created in the directory where the symbolic link resides, and all files and subdirectories are restored to that directory. The symbolic link name is used as the new directory name.
Retrieve a symbolic link that points to a directory; the symbolic link and directory still exist.	Tivoli Storage Manager is not retrieved as long as the symbolic link exists.

Related reference
"Archsymlinkasfile" on page 268

Hard links

When you archive files that are hard-linked, Tivoli Storage Manager archives each instance of the linked file.

For example, if you archive two files that are hard-linked, Tivoli Storage Manager archives the file data twice.

When you retrieve hard-linked files, Tivoli Storage Manager reestablishes the links. For example, if you had a hard-linked pair of files, and only one of the hard-linked files is on your workstation, when you retrieve both files, they are hard-linked. The only exception to this procedure occurs if you archive two files that are hard-linked and then break the connection between them on your workstation. If you retrieve the two files from the server, Tivoli Storage Manager respects the current file system and not retrieve the hard link.

Tip: If you do not archive and retrieve all files that are hard-linked at the same time, problems can occur. To ensure that hard-linked files remain synchronized, archive all hard links at the same time and retrieve those same files together.

Retrieve archives

Retrieve a file when you want to return an archive copy from the server to your workstation.

Many of the advanced considerations for retrieving files are the same as for restoring files.

Important: When you retrieve a file without any specifications, and more than one version of the archive copy exists on the server, Tivoli Storage Manager retrieves

all of the copies. After the first copy is retrieved, the second copy is retrieved. If there is an existing copy on your client workstation, you are prompted to replace, skip, or cancel.

Related concepts

"Restore or retrieve files to another workstation" on page 200

Related tasks

"Authorizing another user to restore or retrieve your files" on page 198

"Restoring or retrieving files from another client node" on page 199

Retrieving data using the GUI

This topic guides you through the steps to retrieve an archived file using the GUI.

1. Click **Retrieve** from the client Java GUI main window. The Retrieve window appears.
2. Expand the directory tree by clicking the plus sign **(+)** or the folder icon next to an object you want to expand. To search or filter files, click the **Search** icon from the tool bar.
3. Enter your search criteria in the Find Files (Retrieve) window.
4. Click the **Search** button. The Matching Files (Retrieve) window appears.
5. Click the selection boxes next to the files you want to retrieve and close the Matching Files (Retrieve) window.
6. Enter your filter criteria in the Find Files (Retrieve) window.
7. Click the **Filter** button. The Retrieve window displays the filtered files.
8. Click the selection boxes next to the filtered files or directories you want to retrieve.
9. To modify specific retrieve options, click the **Options** button. Any options you change are effective during the current session *only*.
10. Click **Retrieve**. The Retrieve Destination window appears. Enter the appropriate information in the Retrieve Destination window.
11. Click **Retrieve**. The Retrieve **Task List** window displays the retrieve processing status. When folders are retrieved from a case sensitive file system (UFS, HFSX, or XSAN) to a case insensitive file system (HFS or HFS+) and they differ only in case, the client retrieves the contents of both folders to one folder.

Retrieve data examples using the command line

You can retrieve a single file, a group of files, or all the files in a directory or subdirectory.

When you retrieve a file, Tivoli Storage Manager sends you a copy of that file. The archived file remains in storage.

Use the **retrieve** command to retrieve files from storage to your workstation. The following table shows examples of using the **retrieve** command.

Table 66. Command line examples of retrieving archives

Task	Command	Considerations
Retrieve the /home/jones/h1.doc file to its original directory.	`dsmc retrieve /home/jones/h1.doc`	If you do not specify a destination, the files are retrieved to their original location.

Task	Command	Considerations
Retrieve the /home/jones/h1.doc file with a new name and directory.	`dsmc retrieve /home/jones/h1.doc /home/smith/h2.doc`	None.
Retrieve all files from the /home/jones directory that end with the characters .bak to the /home/smith directory.	`dsmc retrieve  "/home/jones/*.bak" /home/smith/`	None.
Retrieve the /home/jones/ch1.doc file and assign a description.	`dsmc retrieve /home/jones/ch1.doc -description="Chapter 1, first version"`	If you do not specify a description with the **retrieve** command, the default is `Retrieve Date:x`, where `x` is the current system date.
Use the pick option to display a list of archives from which you can select files to retrieve.	`dsmc retrieve "/home/jones/*" -pick`	None.
Retrieve a list of files that are specified in the retrievelist.txt file to their original directory.	`dsmc retrieve -filelist=/home/ dir2/retrievelist.txt`	None.

Related reference

"Retrieve" on page 544

"Description" on page 295

"Filelist" on page 333

"Pick" on page 382

Archive management classes

Tivoli Storage Manager checks the include options in your include-exclude options list to determine which management class to assign to your archived files.

If you do not assign a management class to a file with the include option, Tivoli Storage Manager assigns the file the default management class. Tivoli Storage Manager can only archive a file if the selected management class contains an archive copy group.

You can override the default management class by using the archmc option, or by selecting the management class from the **Options** menu in the GUI.

You can also add include-exclude statements in the Tivoli Storage Manager client Java GUI or Web client directory tree. Then you can use the **Utilities Preview Include-Exclude** function to preview the include-exclude list before sending data to the server.

Related concepts

"Assign a management class to files" on page 235

"Display information about management classes and copy groups" on page 231

Related reference

"Preview Archive" on page 501

"Preview Backup" on page 502

Chapter 7. Tivoli Storage Manager scheduler overview

The Tivoli Storage Manager central scheduler allows client operations to occur automatically at specified times.

In order to understand scheduling with Tivoli Storage Manager, several terms need to be defined:

schedule definition

A definition on the Tivoli Storage Manager server which specifies critical properties of the automated activity including the type of action, the time the action should take place, and how frequently the action takes place. There are numerous other properties which can be set (see the appropriate *Tivoli Storage Manager Administrator's Reference Guide* for a detailed description of the **define schedule**.)

schedule association

An assignment to a specific schedule definition for a client node. Multiple schedule associations allow single schedule definitions to be used by many client nodes. Because schedule definitions are included with specific policy domains, it is only possible for nodes defined to a certain policy domain to be associated with schedules defined in that domain.

scheduled event

A specific occurrence of when a schedule is executed for a node. The following conditions must be met before automatic scheduled events take place for a client:
- A schedule definition must exist for a specific policy domain.
- A schedule association must exist for the required node which belongs to that policy domain.
- The client scheduler process must be running on the client system.

When creating a schedule definition on the Tivoli Storage Manager server, schedule actions that you can take include incremental, selective, archive, restore, retrieve, image backup (does not apply to Mac OS X), image restore (does not apply to Mac OS X), command, and macro. The scheduled action that is most frequently used is incremental with the *objects* parameter left undefined. With this setting, the Tivoli Storage Manager client performs a domain incremental backup of all file systems defined by the client domain option. A schedule definition using the *command* action allows an operating system command or shell script to be executed. When automating tasks for *Tivoli Storage Manager for Data Protection* clients, you must use *command* action schedule definitions which invoke the command-line utilities for the "Tivoli Storage Manager for" application.

The schedule *startup window* indicates the acceptable time period for a scheduled event to start. The startup window is defined by these schedule definition parameters: *startdate*, *starttime*, *durunits*, and *duration*. The *startdate* and *starttime* options define the beginning of the startup window for the very first scheduled event. The beginning of the startup windows for subsequent scheduled events vary depending on the *period* and *perunit* values of the schedule definition. The *duration* of the schedule window defines the length of the startup window. The schedule action is required to start within the startup window. To illustrate, consider the results of the following schedule definition:

```
define schedule standard test1 action=incremental starttime=12:00:00 period=1
perunits=hour dur=30 duru=minutes
```

Event	Window start	Window end	Actual start (just an example, times vary)
1	12:00:00	12:30:00	12:05:33
2	13:00:00	13:30:00	13:15:02
3	14:00:00	14:30:00	14:02:00
and so on			

The variation in actual start times is a result of the randomization feature provided by the Tivoli Storage Manager central scheduler which helps to balance the load of scheduled sessions on the Tivoli Storage Manager server.

Examples: Blank spaces in file names in schedule definitions

When defining or updating a schedule *objects* parameter with file specifications that contain blank spaces, use quotation marks around each file specification that contains blanks, then single quotation marks around all of the specifications.

The following are some examples.

```
objects='"/home/proj1/Some file.doc"'
objects='"/home/proj1/Some file.doc" "/home/Another file.txt"
/home/noblanks.txt'
objects='"/home/My Directory With Blank Spaces/"'

objects='"/Users/user1/Documents/Some file.doc"'
objects='"/Users/user1/Documents/Some file.doc" "/Users/user5/Documents/
        Another file.txt" /Users/user3/Documents/noblanks.txt'
objects='"/Users/user1/My Directory With Blank Spaces/"'
```

This ensures that /home/proj1/Some file.doc is treated as a single file name, as opposed to two separate files (/home/proj1/Some, and file.doc).

You can also refer to the *objects* parameter information for the DEFINE SCHEDULE and UPDATE SCHEDULE commands in the appropriate *IBM Tivoli Storage Manager Administrator's Reference*.

Related concepts

"Specifying input strings that contain blank spaces or quotation marks" on page 110

Preferential start times for certain nodes

Occasionally, you might want to ensure that a particular node begins its scheduled activity as close as possible to the defined start time of the schedule. The need for this typically arises when prompted mode scheduling is in use.

Depending on the number of client nodes associated with the schedule and where the node is in the prompting sequence, the node might be prompted significantly later than the start time for the schedule.

In this case, you can perform the following steps:
1. Copy the schedule to a new schedule with a different name (or define a new schedule with the preferred attributes).

2. Set the new schedule priority attribute so that it has a higher priority than the original schedule.

3. Delete the association for the node from the original schedule, then associate the node to the new schedule.

Now the Tivoli Storage Manager server processes the new schedule first.

Scheduler processing options

This topic lists the processing options that impact the behavior of the client scheduler.

On the Tivoli Storage Manager client, you can define most of these options in the client options file. However, some of these options can be set globally on the Tivoli Storage Manager server for all Tivoli Storage Manager clients. The *Tivoli Storage Manager Administrator's Guide* provides detailed information on all topics described in the section.

This table shows which options are defined by the client and server, and which options are overridden by the server.

Option	Client defined	Server defined	Server global override
managedservices	X		
maxcmdretries	X		**SET MAXCMDRETRIES** command
maxschedsessions		X	
postschedulecmd, postnschedulecmd	X		
preschedulecmd, prenschedulecmd	X		
queryschedperiod	X		**SET QUERYSCHEDPERIOD** command
randomize		X	
retryperiod	X		**SET RETRYPERIOD** command
schedcmddisabled	X		
schedlogname	X		
schedlogretention	X		
schedmode	X		**SET SCHEDMODES** command
sessioninitiation	X	X	**UPDATE NODE** command
tcpclientaddress	X	X (also defined on server when sessioninit=*serveronly* as part of the node definition)	

Option	Client defined	Server defined	Server global override
tcpclientport	X	X (also defined on server when sessioninit=*serveronly* as part of the node definition)	

Client defined options are defined in the dsm.sys or dsm.opt file, depending on the option and platform. The Tivoli Storage Manager server can also define some options in a client options set, or as part of the options parameter of the schedule definition. The Tivoli Storage Manager server can also set some options globally for all clients. By default, the client setting for these options is honored. If the global override on the Tivoli Storage Manager server is set, the client setting for the option is ignored. Defining client options as part of the schedule definition is useful if you want to use specific options for a scheduled action that differ from the option settings normally used by the client node, or are different for each schedule the node executes.

The schedmode option controls the communication interaction between the Tivoli Storage Manager client and server. There are two variations on the schedule mode: *client polling* and *server prompted*.

Return codes from preschedulecmd and postschedulecmd scripts

This topic lists the return codes that you might see when you use the preschedulecmd and postschedulecmd options.

- If the command specified by the preschedulecmd option ends with a nonzero return code, Tivoli Storage Manager considers the command to have failed. In this case, neither the scheduled event nor any postschedulecmd or postnschedulecmd command runs. The administrative **query event** command with format=*detailed* option shows that the event failed with return code 12.
- If the command specified by the postschedulecmd option ends with a nonzero return code, Tivoli Storage Manager considers the command to have failed. The administrative **query event** command with format=*detailed* option shows that the event completed with return code 8, unless the scheduled operation completed with a higher return code, in which case the higher return code takes precedence. Therefore, if the scheduled operation completes with return code 0 or 4 and the postschedulecmd command fails, the administrative **query event** command shows that the event completed with return code 8. If the scheduled operation completes with return code 12, that return code takes precedence, and **query event** shows that the event failed with return code 12.

When interpreting the return code from a command, Tivoli Storage Manager considers 0 to mean success, and anything else to mean failure. While this behavior is widely accepted in the industry, it is not 100% guaranteed. For example, the developer of the widget command might exit with return code 3, if widget ran successfully. Therefore, it is possible that the preschedulecmd or postschedulecmd command might end with a nonzero return code and be successful. To prevent Tivoli Storage Manager from treating such commands as failed, you should wrap these commands in a script, and code the script so that it interprets the command return codes correctly. The script should exit with return code 0 if the command

was successful; otherwise it should exit with a nonzero return code. The logic for a script running widget might look like this:

```
run 'widget'
   if lastcc == 3
      exit 0
   else
      exit 1
```

Related reference

Client-acceptor scheduler services versus the traditional scheduler services

You can configure the Tivoli Storage Manager client to manage the scheduler process using the Tivoli Storage Manager client acceptor.

The client acceptor daemon provides a light-weight timer which automatically starts and stops the scheduler process as needed. Alternatively, the traditional method keeps the Tivoli Storage Manager scheduler process running continuously. Generally, using the client acceptor daemon to manage the scheduler is the preferred method.

The following information is a comparison of the client acceptor daemon-managed services and the traditional scheduler services methods.

client acceptor daemon-managed services
- Defined using the managedservices schedule option and started with client acceptor daemon services (dsmcad).
- The client acceptor daemon starts and stops the scheduler process as needed for each scheduled action.
- Requires fewer system resources when idle.
- Tivoli Storage Manager client options and Tivoli Storage Manager server override options are refreshed each time the client acceptor daemon services start a scheduled backup.
- Cannot be used with SESSIONINITiation=SERVEROnly backups.

Tivoli Storage Manager traditional scheduler services
- Started with command `dsmc sched` command.
- Remains active, even after scheduled backup is complete.
- Requires higher use of system resources when idle.
- Tivoli Storage Manager client options and Tivoli Storage Manager server override options are only processed once when `dsmc sched` is started.
- You must restart the scheduler process for updated Tivoli Storage Manager options to take effect.

Tip: Restart the traditional scheduler periodically to free system resources previously used by system calls.

Set the client scheduler process to run as a background task and start automatically at startup

You can configure the Tivoli Storage Manager client scheduler to run as a background system task which starts automatically when your system is started.

This is true for both client acceptor daemon-managed and traditional methods of running the Tivoli Storage Manager client scheduler.

When running a client acceptor daemon-managed schedule, only the client acceptor daemon process should be set to start automatically at startup time; not the scheduler process. For the traditional method, the scheduler process should be set up to start automatically at startup time.

You can configure the client acceptor daemon to run as a background system task which starts automatically when your system is started. To configure the client acceptor daemon to manage scheduled backups, you must set the managedservices option to manage the scheduler, or both the scheduler and Web client. The method for setting up the client acceptor daemon as a system task varies for each platform.

In order for the scheduler to start unattended, you must enable the client to store its password by setting the passwordaccess option to *generate*, and store the password by running a simple Tivoli Storage Manager client command such as `dsmc query session`. Note that for testing purposes, you can always start the scheduler in the foreground by running `dsmc sched` from a command prompt (without a 'managedservices' stanza set).

To start the scheduler automatically at startup time, use either of the following methods:

Client acceptor daemon-managed

1. In your dsm.sys file, set the managedservices option to *schedule* or *schedule webclient*.
2. Add the following entry into the system startup file (/etc/inittab for most platforms):

   ```
   tsm::once:/usr/bin/dsmcad > /dev/null 2>&1 # TSM Client Acceptor
   Daemon
   ```
3. In your dsm.sys file, set the passwordaccess option to *generate*.
4. Run a command like `dsmc query sess` to store the node password.

Traditional

1. In your dsm.sys file, either remove the managedservices option or set it to *webclient*.
2. Add the following entry into the system startup file (/etc/inittab for most platforms):

   ```
   tsmsched::once:/usr/bin/dsmc sched > /dev/null 2>&1 # TSM scheduler
   ```
3. In your dsm.sys file, set the passwordaccess option to *generate*.
4. Run a command like `dsmc query sess` to store the node password.
5. To start the client scheduler on your client node and connect to the server schedule, enter the following command:

   ```
   dsmc schedule
   ```

   If the current directory is not in your PATH environment variable, enter the following command:

```
./dsmc schedule
```

When you start the client scheduler, it runs continuously until you close the window, end the process, or log off your system.

To run the **schedule** command in the background and to keep the client scheduler running, even if you log off your system, enter the following:

```
nohup dsmc schedule 2> /dev/null &
```

Requirement: For Mac OS X, the Tivoli Storage Manager client acceptor daemon must be installed as a Startup Item. A System Administrator must use TSM Tools for Administrators to install and start the client acceptor daemon.

For Mac OS X, after the client acceptor daemon has been installed as a startup item, the SystemStarter utility can be used to start, restart, or temporarily stop the client acceptor daemon.

To start the client acceptor daemon use the following command:

```
sudo /sbin/SystemStarter start dsmcad
```

To restart the client acceptor daemon use the following command:

```
sudo /sbin/SystemStarter restart dsmcad
```

To stop the client acceptor daemon use the following command:

```
sudo /sbin/SystemStarter stop dsmcad
```

Examples: Display information about scheduled work

Schedules can be classic or enhanced, depending on how the interval to the next execution is defined.

Classic schedules allow the period to be as small as an hour. Enhanced schedules allow actions to be executed on specific days.

To view schedules that are defined for your client node, enter:

```
dsmc query schedule
```

Tivoli Storage Manager displays detailed information about all scheduled work for your client node. Table 67 on page 220 displays sample classic **query schedule** output.

Table 67. Sample classic query schedule output

```
 Schedule Name: DAILY_INC
   Description: Daily System-wide backup
Schedule Style: Classic
        Action: Incremental
       Options: QUIET
       Objects:
      Priority: 1
Next Execution: 30 minutes
      Duration: 4 Hours
        Period: 1 Day
   Day of Week: Any
         Month:
  Day of Month:
 Week of Month:
        Expire: Never

 Schedule Name: WEEKLY_INC
   Description: Weekly backup for project files
Schedule Style: Classic
        Action: Incremental
       Options: QUIET
       Objects: /proj
      Priority: 1
Next Execution: 60 minutes
      Duration: 8 Hours
        Period: 7 Days
   Day of Week: Friday
         Month:
  Day of Month:
 Week of Month:
        Expire: Never
```

The schedule name, **WEEKLY_INC**, starts a weekly incremental backup in the /proj file system.

The schedule name, **DAILY_INC**, starts a daily incremental backup. The next incremental backup starts in 30 minutes. Because no objects are listed, Tivoli Storage Manager runs the incremental backup on your default domain. The schedule has no expiration date.

To more accurately determine the status of scheduled events, the **query schedule** output for an enhanced schedule, on Tivoli Storage Manager Version 5.3 client and above, includes new fields. These fields are always displayed, even if it is a classic schedule or a Tivoli Storage Manager Version 5.3 client session with a pre-Version 5.3 server, but the new fields are blank. Note that for a down-level (prior to Tivoli Storage Manager Version 5.3) client, the server reports the period as indefinite and the day of week as an illegal day. Table 68 on page 221 displays sample enhanced **query schedule** output.

Table 68. Sample enhanced query schedule output

```
Schedule Name: QUARTERLY_FULL
   Description: Quarterly full backup
Schedule Style: Enhanced
        Action: Selective
       Options: subdir=yes
       Objects: /* /Volumes/fs2/*
      Priority: 5
Next Execution: 1744 Hours and 26 Minutes
      Duration: 1 Day
        Period:
   Day of Week: Friday
         Month: March, June, September, December
  Day of Month: Any
 Week of Month: Last
        Expire: Never
```

Display information about completed work

When you run the **schedule** command in the foreground, your screen displays output from the scheduled commands.

Output is also directed to the dsmsched.log file in the installation directory unless you change the directory and file name using the schedlogname option.

When you run the **schedule** command in the background, output from scheduled commands is directed to the dsmsched.log file in the current directory, or to the path and file name that you specified. The dsmsched.log cannot be a symbolic link.

Note: On Mac OS X, by default the log can be found in one of these locations:
```
~/Library/Logs/tivoli/tsm
/Library/Logs/tivoli/tsm
```

After scheduled work is performed, check the schedule log to verify that all work completed successfully.

When a scheduled command is processed the schedule log contains the following entry:

    Scheduled event *eventname* completed successfully

If the scheduled event does not complete successfully, you receive a message similar to the following:

    ANS1512E Scheduled event *eventname* failed. Return code = *code*.

The client indicates whether Tivoli Storage Manager successfully issued the scheduled command associated with the *eventname* (action=command). No attempt is made to determine the success or failure of the command. You can assess the status of the command by evaluating the return code from the scheduled command in the schedule log. The schedule log entry for the return code of the command is prefaced with the following text:

    Finished command. Return code is:

The schedule log continues to grow unless you prune it using the schedlogretention option or specify a maximum size using the schedlogmax option.

Related concepts

"Specify scheduling options"

Specify scheduling options

You can modify scheduling options in the client options file or the graphical user interface (GUI).

However, if your administrator specifies a value for these options, that value overrides the value in your client.

Related concepts

"Scheduling options" on page 254

Scheduler options for commands

The scheduler executes commands under a user ID of 0 (root); however, some commands might need to be executed under a user ID other than 0.

In this case, your Tivoli Storage Manager administrator can define schedules for commands that are executed under a user ID different from the scheduler user ID using the schedcmduser server option.

The schedcmduser option specifies the name of a valid user on the system where a scheduled command is executed. This option can only be defined by the Tivoli Storage Manager server administrator. If this option is specified, the command is executed with the authorization of the specified user. Otherwise, it is executed with the scheduler authorization.

▶▶──SCHEDCMDUser────*user_name*──▶◀

user_name
    Specifies the name of a valid user on the system where a scheduled command is executed.

Note: The schedcmduser option does *not* affect the user ID used for the pre-schedule and post-schedule commands. Pre-schedule and post-schedule always run as root (user ID 0).

Enable or disable scheduled commands

You can use the schedcmddisabled option to disable the scheduling of commands by the server.

Commands are scheduled by using the action=*command* option on the DEFINE SCHEDULE server command.

The schedcmddisabled option does not disable the preschedulecmd and postschedulecmd commands. However, you can specify preschedulecmd or postschedulecmd with a blank or a null string to disable the scheduling of these commands.

You can use the schedrestretrdisabled option to prevent the Tivoli Storage Manager Server administrator from executing restore or retrieve schedule operations.

You can use the srvprepostscheddisabled option to prevent the Tivoli Storage Manager Server administrator from executing pre-schedule and post-schedule commands when performing scheduled operations.

You can use the srvprepostsnapdisabled option to prevent the Tivoli Storage Manager Server administrator from executing pre-snapshot and post-snapshot commands when performing scheduled image snapshot backup operations.

Related reference

"Schedcmddisabled" on page 402

"Schedrestretrdisabled" on page 409

"Srvprepostscheddisabled" on page 428

"Srvprepostsnapdisabled" on page 429

Manage multiple schedule requirements on one system

In certain situations it is preferable to have more than one scheduled activity for each client system.

Normally, you can do this by associating a node with more than one schedule definition. This is the standard method of running multiple schedules on one system.

You must ensure that the schedule windows for each schedule do not overlap. A single client scheduler process is not capable of executing multiple scheduled actions simultaneously, so if there is overlap, the second schedule to start is missed if the first schedule does not complete before the end of the startup window of the second schedule.

Suppose that most of the file systems on your client system must be backed up daily, and that one file system containing critical data must be backed up hourly. In this case, you would need to define two schedules to handle this requirement. To avoid conflict between the hourly and daily backup schedule, the *starttime* of each schedule needs to be varied.

In certain cases, it is necessary to run more than one scheduler process on a system. Multiple processes require a separate options file for each process and must contain the following information:

- Define a unique node name for each process
- Specify unique schedule and error logs for each process
- When running in prompted mode, you must use the tcpclientport option to specify a unique port for each process.

The advantages of using multiple schedule processes:

- You can run more than one scheduled backup at the same time.
- You can specify different backup criteria for each schedule started, with the Tivoli Storage Manager client option file or Tivoli Storage Manager server override options.

The disadvantages of using multiple schedule processes:

- A unique file space for each node name on the Tivoli Storage Manager server is created.
- When restoring the data, you must use the same node name associated with the backup.

Multiple schedule processes can run on UNIX and Linux platforms with either the client acceptor daemon-managed method, or the traditional method of running the scheduler. In either case, there are certain setup requirements:

- Each process must run using a different node name.
- You must create multiple stanzas in the dsm.sys file for each scheduler process. In each stanza, you must define a unique node name, along with unique values for the options errorlogname and schedlogname. You might also choose to define customized domain, include, and exclude statements for each stanza.
- In your dsm.sys file, set the passwordaccess option to generate in each stanza. The password must be generated for each node name that is running a scheduler process, by running a command such as `dsmc query sess`.
- If running with the schedmode option set to *prompt*, you should set a unique tcpclientport value for each stanza.

You must start each `dsmc sched` command or instance with the -servername option to reference its unique stanza name in dsm.sys. For dsmcad, it is necessary to define the environment variable DSM_CONFIG for each instance of dsmcad to reference its unique option file.

The following is an example configuration of two schedule processes managed by the client acceptor daemon in the dsm.sys file. Note that you must use full paths for the log file names to avoid the files being written in the root directory):

```
servername tsm1_sched1
  nodename        aixsvt01_sched1
  tcpserv         firebat
  tcpclientport   1507
  passwordaccess  generate
  domain          /svt1
  schedmode       prompted
  schedlogname    /tsm/dsmsched1.log
  errorlogname    /tsm/dsmerror1.log
  managedservices schedule

servername tsm1_sched2
  nodename        aixsvt01_sched2
  tcpserv         firebat
  tcpclientport   1508
  passwordaccess  generate
  domain          /svt1
  schedmode       prompted
  schedlogname    /tsm/dsmsched2.log
  errorlogname    /tsm/dsmerror2.log
  managedservices schedule
```

Contents of /test/dsm.opt1:

```
servername tsm1_sched1
```

Contents of /test/dsm.opt2:

```
servername tsm1_sched2
```

Open two shell command windows:

- In shell command window 1, enter:
  ```
  export DSM_CONFIG=/test/dsm.opt1
  sudo dsmcad
  ```
- In shell command window 2, enter:
  ```
  export DSM_CONFIG=/test/dsm.opt2
  sudo dsmcad
  ```

Note: You should enter these commands into a shell script if you intend to have the dsmcad processes started directly from /etc/inittab so that the proper DSM_CONFIG variable can be set prior to launching dsmcad.

Chapter 8. Client return codes

The backup-archive command-line interface and the scheduler exit with return codes that accurately reflect the success or failure of the client operation.

Scripts, batch files, and other automation facilities can use the return code from the command-line interface. For operations that use the Tivoli Storage Manager scheduler, the return codes are shown in the output of the QUERY EVENT administrative command.

In general, the return code is related to the highest severity message during the client operation.
- If the highest severity message is informational (ANSnnnnI), then the return code is 0.
- If the highest severity message is a warning (ANSnnnnW), then the return code is 8.
- If the highest severity message is an error (ANSnnnnE or ANSnnnnS), then the return code is 12.

The exception to the above rules is warning or error messages that individual files could not be processed. For such a skipped file, the return code is 4. For cases where the return code is not 0, you can examine the dsmerror.log file (and, for scheduled events, the dsmsched.log file).

For a description of the return codes and their meanings, see the following table.

Table 69. Client return codes and their meanings

Code	Explanation
0	All operations completed successfully.
4	The operation completed successfully, but some files were not processed. There were no other errors or warnings. This return code is very common. Files are not processed for various reasons. The most common reasons are: • The file satisfies an entry in an exclude list. • The file was in use by another application and could not be accessed by the client. • The file changed during the operation to an extent prohibited by the copy serialization attribute. See "Copy serialization attribute" on page 233.
8	The operation completed with at least one warning message. For scheduled events, the status is `Completed`. Review dsmerror.log (and dsmsched.log for scheduled events) to determine what warning messages were issued and to assess their impact on the operation.
12	The operation completed with at least one error message (except for error messages for skipped files). For scheduled events, the status is `Failed`. Review the dsmerror.log file (and dsmsched.log file for scheduled events) to determine what error messages were issued and to assess their impact on the operation. As a general rule, this return code means that the error was severe enough to prevent the successful completion of the operation. For example, an error that prevents an entire file system or file specification from being processed yields return code 12.

Table 69. Client return codes and their meanings (continued)

Code	Explanation
other	For scheduled operations where the scheduled action is COMMAND, the return code is the return code from the command that was executed. If the return code is 0, the status of the scheduled operation is *Completed*. If the return code is nonzero, then the status is *Failed*. Some commands might issue a nonzero return code to indicate success. For these commands, you can avoid a `Failed` status by *wrapping* the command in a script that invokes the command, interprets the results, and exits with return code 0 if the command was successful (the script should exit with a nonzero return code if the command failed). Then ask your Tivoli Storage manager server administrator to modify the schedule definition to invoke your script instead of the command.

The return code for a client macro is the highest return code issued among the individual commands that comprise the macro. For example, suppose a macro consists of these commands:

```
selective "/home/devel/*" -subdir=yes
incremental "/home/devel/TestDriver/*" -subdir=yes
archive "/home/plan/proj1/*" -subdir=yes
```

If the first command completes with return code 0; the second command completed with return code 8; and the third command completed with return code 4, the return code for the macro is 8.

See the *IBM Tivoli Storage Manager Administrator's Reference* for your operating system for more information about QUERY EVENT.

Related concepts

"Scheduler options for commands" on page 222

Chapter 9. Storage management policies

Storage management policies are rules your administrator defines in order to manage your backups and archives on the server.

Your data is associated (or bound) to these policies; then when the data is backed up or archived, it is managed according to policy criteria. Policy criteria include a policy domain, a policy set, a management class, and a copy group.

Policies determine:
- Whether a file is eligible for backup or archive services.
- How many backup versions to keep.
- How long to keep inactive backup versions and archive copies.
- Where to place the copies in storage.
- For incremental backup, policies also determine:
  - How frequently a file can be backed up.
  - Whether a file must change before it is backed up again.

If you have the Tivoli HSM client installed, your administrator also defines rules that determine whether files are eligible for migration from your local file systems to storage.

This topic explains:
- Policy criteria (policy domains, policy sets, copy groups, and management classes).
- How to display policies.
- How Tivoli Storage Manager associates your data with policies.

Policy domains and policy sets

A *policy domain* is a group of clients with similar requirements for backing up and archiving data.

Policy domains contain one or more policy sets. An administrator uses policy domains to manage a group of client nodes in a logical way.

For example, a policy domain might include:
- A department, such as Accounting.
- A physical location, such as a particular building or floor.
- A local area network, such as all clients associated with a particular file server.

Tivoli Storage Manager includes a default policy domain named *Standard*. At first, your client node might be associated with the default policy domain. However, your administrator can define additional policy domains if there are groups of users with unique backup and archive requirements.

A *policy set* is a group of one or more management classes. Each policy domain can hold many policy sets. The administrator uses a policy set to implement different management classes based on business and user needs. Only one of these policy sets can be active at a time. This is called the *active policy set*. Each policy set contains a *default management class* and any number of additional management classes.

Management classes and copy groups

A *management class* is a collection of backup and archive copy groups that establishes and contains specific storage management requirements for backing up and archiving data.

An administrator can establish separate management classes to meet the backup and archive requirements for different kinds of data, such as:
- System data that is critical for the business.
- Application data that changes frequently.
- Report data that Management reviews monthly.
- Legal information that must be retained indefinitely, requiring a large amount of disk space.

Note: If you have the Tivoli HSM client installed, it can also contain specific requirements for migrating files to storage.

Most of the work you do with storage management policies is with management classes. Each file and directory that you back up, and each file that you archive, is associated with (or *bound* to) a management class, as follows:
- If your data is not associated with a management class, Tivoli Storage Manager uses the default management class in the active policy set.
- When backing up directories, you can specify a management class with an *include* statement or the dirmc option. If you do not specify a management class, Tivoli Storage Manager uses the management class in the active policy set specifying the longest "Retain Only" retention period. If there are multiple management classes that meet this criteria, Tivoli Storage Manager uses the last one found, in alphabetical order.
- For archiving directories, you can specify a management class with an *include.archive* statement or the archmc option. If you do not specify a management class, the server assigns the default management class to the archived directory. If the default management class has no archive copy group, the server assigns the management class that currently has the archive copy group with the shortest retention time.

You can use *include* statements in your include-exclude list to associate files with management classes. In your client options file, you can associate directories with a management class, using the dirmc option.

Within a management class, the specific backup and archive requirements are in *copy groups*. Copy groups define the specific storage management attributes that describe how the server manages backed up or archived data. Copy groups include both *backup copy groups* and *archive copy groups*. A management class can have one backup copy group, one archive copy group, both, or neither.

A *backup copy group* contains attributes that are used during the backup process to determine:
- How many days must elapse before a file is backed up again.
- How a file is processed during a backup if it is in use.

It also contains attributes to manage the backup versions of your files on the server. These attributes control:
- On which media type the server stores backup versions of your files and directories.
- How many backup versions the server keeps of your files and directories.

- How long the server keeps backup versions of your files and directories.
- How long the server keeps inactive backup versions.
- How long the last remaining inactive version of a file is kept.

An *archive copy group* contains attributes that control:
- Whether a file is archived if it is in use
- On which media type the server stores archived copies of your files
- How long the server keeps archived copies of your files

Related concepts

"Select a management class for files" on page 234

"Retention grace period" on page 238

Display information about management classes and copy groups

You can display policy information with the command-line interface or with a graphical user interface.

On a graphical user interface, click **View policy information** from the Utilities menu. The **Policy information** window displays the available management classes. On a command line, use the **query mgmtclass** command to view the available management classes. The detail option provides more information.

Table 70 shows the default values for the backup and archive copy groups in the standard management class.

Table 70. Default attribute values in the standard management class

Attribute	Backup default	Archive default
Copy group name	Standard	Standard
Copy type	Backup	Archive
Copy frequency	0 days	CMD (Command)
Versions data exists	Two versions	Does not apply
Versions data deleted	One version	Does not apply
Retain extra versions	30 days	Does not apply
Retain only version	60 days	Does not apply
Copy serialization	Shared static	Shared static
Copy mode	Modified	Absolute
Copy destination	Backuppool	Archivepool
Retain versions	Does not apply	365 days
Lan free	Destination	No
Deduplication enabled	No	No

Copy group name attribute

The *copy group name* attribute is the name of the copy group. The default value for both backup and archive is *standard*.

Copy type attribute

The *copy type* attribute is the type of the copy group. The value for backup is always *backup*, and the value for archive is always *archive*.

Copy frequency attribute

The *copy frequency* attribute is the minimum number of days that must elapse between successive incremental backups. Use this attribute during a full incremental backup.

Copy frequency works with the *mode* parameter. For example, if frequency is *zero (0)* and mode is *modified*, a file or directory is backed up *only if* it changed since the last incremental backup. If frequency is *zero (0)* and mode is *absolute*, a file is backed up every time you run an incremental backup against it. This attribute is not checked for selective backups.

For archive copy groups, copy frequency is always CMD (command). There is no restriction on how often you archive a file.

Copy frequency is ignored during a journal-based backup.

Versions data exists attribute

The *versions data exists* attribute specifies the maximum number of different backup versions retained for files and directories.

If you select a management class that permits more than one backup version, the most recent version is called the *active* version. All other versions are called *inactive* versions. If the maximum number of versions permitted is five, and you run a backup that creates a sixth version, the oldest version is deleted from server storage.

Versions data deleted attribute

The *versions data deleted* attribute specifies the maximum number of different backup versions retained for files and directories that you deleted.

This parameter is ignored until you delete the file or directory.

If you delete the file or directory, the next time you run an incremental backup, the active backup version is changed to inactive. Tivoli Storage Manager deletes the oldest versions in excess of the number specified by this parameter.

The expiration date for the remaining versions is based on the *retain extra versions* and *retain only version* parameters.

Retain extra versions attribute

The *retain extra versions* attribute specifies how many days all but the most recent backup version is retained.

The most recent version is the active version, and active versions are never erased. If *Nolimit* is specified, then extra versions are kept until the number of backup versions exceeds the *versions data exists* or *versions data deleted* parameter settings. In this case, the oldest extra version is deleted immediately.

Retain only version attribute

The *retain only version* attribute specifies the number of days the last remaining inactive version of a file or directory is retained.

If *Nolimit* is specified, the last version is retained indefinitely.

This parameter goes into effect during the next incremental backup after a file is deleted from the client system. Any subsequent updates to this parameter will not affect files that are already inactive. For example: If this parameter is set to 10 days when a file is inactivated during an incremental backup, the file is deleted from the server in 10 days.

Copy serialization attribute

The *copy serialization* attribute determines whether a file can be in use during a backup or archive, and what to do if it is.

The value for this attribute can be one of the following:
- **Static**. A file or directory must not be modified during a backup or archive. If the object is changed during a backup or archive attempt, it is not backed up or archived.
- **Shared static**. A file or directory must not be modified during backup or archive. Tivoli Storage Manager attempts to perform a backup or archive as many as four additional times, depending on the value specified on the changingretries option in your options file. If the object is changed during every backup or archive attempt, it is not backed up or archived.
- **Dynamic**. A file or directory is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.
- **Shared dynamic**. A file or directory is backed up or archived regardless of whether it changes during a backup or archive. Tivoli Storage Manager attempts to back up or archive as many as four additional times. The number of attempts depend on the value that was specified on the changingretries option in your options file, without the file changing during the attempt. The file is backed up or archived on the last try even if it has changed.

  If you select a management class that permits a file to be backed up or archived while it is in use, the backup version or archived copy that is stored on the server might be a fuzzy copy. A *fuzzy copy* is a backup version or archived copy that does not accurately reflect what is currently in the file. It might contain some, but not all, of the changes. If that is not acceptable, select a management class that creates a backup version or archive copy only if the file does not change during a backup or archive. When you use static serialization, applications cannot open a file for write access while the file is being backed up.

  If you restore or retrieve a file that contains a fuzzy copy, the file might not be usable. Do not use dynamic or shared dynamic serialization to back up files unless you are certain that a fuzzy copy that is restored is usable.

  Important: Be careful when you select a management class containing a copy group that specifies shared dynamic or serialization dynamic backup.

  Related reference
  "Snapshotproviderimage" on page 424

Copy mode attribute

The *copy mode* attribute determines whether a file or directory is considered for incremental backup regardless of whether it changed or not since the last backup.

Tivoli Storage Manager does not check the mode for selective backups.

The value for this parameter can be one of the following:

- **Modified.** The file is considered for incremental backup *only if* it has changed since the last backup. A file is considered changed if any of the following are true:
  - The date or time of the last modification is different.
  - The file size is different.
  - If only the file metadata changes (such as access permissions), but the file data does not change, Tivoli Storage Manager might back up only the metadata.
  - The file owner is different.
  - The file permissions are different.
- **Absolute.** The file is considered for incremental backup regardless of whether it changed since the last backup. For archive copy groups, the mode is always *absolute*, indicating that a file is archived regardless of whether it changed since the last archive request.

Copy destination attribute

The *copy destination* attribute names the destination where backups or archives are stored.

The destination can be either a storage pool of disk devices or a storage pool of devices that support removable media, such as tape.

Retain versions attribute

The *retain versions* attribute specifies the number of days an archived file remains in storage.

When the specified number of days elapse for an archived copy of a file, it is deleted from server storage.

Deduplicate data attribute

The *deduplicate data* attribute specifies whether redundant data is transferred to the Tivoli Storage Manager server during backup and archive processing.

Related concepts

"Client-side data deduplication" on page 74

Related reference

"Deduplication" on page 293

"Enablededupcache" on page 312

"Exclude options" on page 321

Select a management class for files

If the default management class meets the backup and archive requirements for all the files on your workstation, it is not necessary to take any action to associate your files with that management class. This is done automatically when you back up or archive your files.

When selecting a different management class for your files, consider these questions:

- Does the management class contain a backup copy group?

  If you attempt to back up a file associated with a management class that does not contain a backup copy group, the file is not backed up.

- Does the management class contain an archive copy group?

You cannot archive a file associated with a management class that does not contain an archive copy group.

- Does the backup copy group contain attributes that back up your files often enough?

  Mode and frequency work together to control how often a file is backed up when you use incremental backup. Tivoli Storage Manager does not check those attributes for selective backup.

- What serialization method does the copy group use?

  The serialization method determines how Tivoli Storage Manager functions when a file changes while it is being backed up.

- Does the backup copy group specify an adequate number of backup versions to keep, along with an adequate length of time to keep them?

- Does the archive copy group specify an adequate length of time to keep archived copies of files?

Related concepts

"Copy serialization attribute" on page 233

Assign a management class to files

A management class defines when your files are included in a backup, how long they are kept on the server, and how many versions of the file the server should keep.

The server administrator selects a default management class. You can specify your own management class to override the default management class.

To assign a management class other than the default to directories, use the dirmc option in your options file.

You can assign a management class for a file or file group by using an *include* statement in your options file. You can also assign a management class by using an *include* statement in include-exclude file specified by the inclexcl option. Management class names are not case-sensitive.

Using the command-line client, to associate all files in the costs directory with the management class **budget**, you would enter:

```
include /home/proj2/costs/* budget
```

To specify a management class named **managall** to use for all files to which you do not explicitly assign a management class, enter the following:

```
include /.../* managall
```

The example below demonstrates how to use a management class:

```
exclude /.../*.sno
include /home/winter/.../*.ice     mcweekly
include /home/winter/december/*.ice mcdaily
include /home/winter/january/*.ice  mcmonthly
include /home/winter/february/white.sno
```

Processing follows these steps:

1. The file white.sno in the february directory in the winter directory is backed up following bottom-up processing rules. Because you did not specify a management class on this statement, the file is assigned to the default management class.

2. Any file with an extension of `ice` in the **january** directory is assigned to the management class, **mcmonthly**.

3. Any file with an extension of `ice` in the **december** directory is assigned to the management class, **mcdaily**.

4. Any other files with an extension of `ice` in any directory under **winter** are assigned to the management class, **mcweekly**.

5. Any file with an extension of `sno` in any directory is excluded from backup. The exception to this rule is white.sno in the february directory in the winter directory.

To specify your own default management class `mgmt_class_name` for files that are not explicitly included, put the statement below at the top of your include list:

```
include /.../* mgmt_class_name
```

When you archive a file using the graphical user interface, you can select a different management class to override the management class assigned to the file.

Related reference

"Dirmc" on page 298

"Include options" on page 351

Override the management class for archived files

When you archive a file, you can override the assigned management class using the a graphical user interface (GUI), or by using the archmc option on the **archive** command.

Overriding the management class using the GUI is equivalent to using the archmc option on the **archive** command. To use the GUI, press the **Options** button on the archive tree to override the management class and select a different management class.

On the command line, to associate the file `budget.jan` with the management class **ret2yrs**, enter this command:

```
dsmc archive –archmc=ret2yrs /home/jones/budget.jan
```

Select a management class for directories

If the management class in your active policy set containing the longest "Retain only version" (RETONLY) setting meets your backup requirements for directories, it might not be necessary to take any action to associate directories with that management class. Tivoli Storage Manager does it automatically when it backs up your directories.

If there is more than one management class with the longest RETONLY setting, the Tivoli Storage Manager client selects the management class whose name is last in alphabetical order.

If the default management class does not meet your requirements, select a management class with an adequate retention period specified by the *retain only version* parameter. For example, if the management class happens to back up data directly to tape, but you want your directory backups to go to disk, you must choose a different management class. You should keep directories at least as long as you keep the files associated with those directories.

For backup directories, use the dirmc option to specify the management class to which directories are bound.

For archive directories, use the archmc option with the **archive** command.

You can use these methods to view the available management classes and their attributes:
- GUI or Web Client: Select **View Policy Information** from the **Utilities** menu.
- command-line client: Run **dsmc query mgmtclass -detail**.

Note: During expiration processing on a Tivoli Storage Manager server, if an archived directory is eligible for expiration, the server checks if any existing archived files require the archived directory to remain. If so, the archived directory is not expired and the Tivoli Storage Manager client updates the insert date on the archived directory to ensure that the directory is not expired before the files under it.

Bind management classes to files

Binding associates a file with a management class.

When you back up a file for the first time, Tivoli Storage Manager binds it to either the default management class or the management class specified in your include-exclude list.

If the backup copy group for the management class specifies keeping multiple backup versions of the file, and you request multiple backups, the server always has one active backup version (the current version) and one or more inactive backup versions of the file. All backup versions of a file are bound to the same management class and are managed based on the attributes in the backup copy group.

When you archive a file for the first time, Tivoli Storage Manager binds it to the default management class, to the management class specified in your include-exclude list, or to a management class you specify when modifying your archive options during an archive.

Archived files are never rebound to a different management class. If you change the management class for a file using an *include.archive* statement, the archmc option, or through a Tivoli Storage Manager GUI, any previous copies of the file that you archived remain bound to the management class specified when you archived them.

If a file is deleted on the client system then that inactive objects of the file are not rebound.

Refer to the *IBM Tivoli Storage Manager Administrator's Guide* for your operating system, for more information on how files and directories are associated with management classes.

Rebind backup versions of files

Rebinding associates a file or a logical volume image with a new management class.

Backups of files are bound again to a different management class in the following conditions. In each condition, the files (active and inactive) are not bound again until the next backup.

- You specify a different management class in an Include statement to change the management class for the file. The backups are managed based on the old management class until you run another backup.
- Your administrator deletes the management class from your active policy set. The default management class is used to manage the backup versions when you back up the file again.
- Your administrator assigns your client node to a different policy domain and the active policy set in that domain does not have a management class with the same name. The default management class for the new policy domain is used to manage the backup versions.

Refer to the *IBM Tivoli Storage Manager Administrator's Guide* for your operating system, for more information on how files and directories are associated with management classes.

Retention grace period

Tivoli Storage Manager also provides a *backup retention grace period* and an *archive retention grace period* to help protect your backup and archive data when it is unable to rebind a file to an appropriate management class.

The backup retention grace period is in the following cases:

- You change the management class for a file, but neither the default management class nor the new management class contain a backup copy group.
- The management class to which a file is bound no longer exists, and the default management class does not contain a backup copy group.

The backup retention grace period, defined in your policy domain, starts when you run an incremental backup. The default is 30 days. However, your administrator can lengthen or shorten this period.

When Tivoli Storage Manager manages a file using the backup retention grace period, it does not create any new backup versions of the file. All existing backup versions of the file expire 30 days (or the number of days specified in your policy domain) from the day they are marked inactive.

Archive copies are never rebound because each archive operation creates a different archive copy. Archive copies remain bound to the management class name specified when the user archived them. If the management class to which an archive copy is bound no longer exists or no longer contains an archive copy group, the server uses the default management class. If you later change or replace the default management class, the server uses the updated default management class to manage the archive copy. If the default management class does not contain an archive copy group, the server uses the archive retention grace period specified for the policy domain.

Event-based policy retention protection

All management classes with an archive copy group must specify a retention period, for example, the number of days that an archived object is stored on the server before being deleted.

Event-based policy provides the option of beginning the retention period either at the time the object is archived or at a later date when an activation event is sent to the server for that object.

Using the Tivoli Storage Manager copy group value RETINIT=CREATE starts the data retention period when the file is archived. Using the copy group value RETINIT=EVENT starts the data retention period when the server is notified that the event has occurred.

The following example demonstrates this concept:

The user has two files, `create.file` and `event.file`. The user has available two management classes; CREATE, with RETINIT= CREATE, and EVENT, with RETINIT=EVENT. Both management classes have a 60-day retention period. The user, on the same day, archives both files:

```
dsmc archive create.file -archmc=CREATE
dsmc archive event.file -archmc=EVENT
```

Ten days later, the user issues the **set event** *-type=hold* command for the `create.file` file, so the file cannot be deleted. On the same day the user issues the **set event** *-type=activate* for the `event.file` file. At this time, `create.file` has 50 days left on its retention period, and `event.file` has 60 days. If no other action is taken, `create.file` remains on the server forever, and `event.file` is expired 70 days after it was created (60 days after its event occurred). However, if 20 days after the initial archive, the user issues **set event** *-type=release* for the `create.file` file. Thirty days of its retention period have passed, so the file is expired in 30 days (the hold does not extend the retention period).

See the *IBM Tivoli Storage Manager Administrator's Guide* for your operating system, for more information about RETINIT.

Related reference

"Set Event" on page 553

Archive files on a data retention server

Up to this point, there is no difference between archiving files on a normal server or a data retention server.

The following example demonstrates the differences between the two servers, and what can be done at day 5:

If the files were archived on a non-data retention server, the user can issue the **delete archive** *create.file event.file* command and both files are deleted. If the files were archived on a data retention server, the same command fails both files. The data retention server forces the user to keep archives until the stated retention criteria are met.

Now here is the difference at day 15 (after the hold):

The **delete archive** *create.file event.file* command on the non-data retention server now deletes `event.file`, but returns a *cannot delete* error for `create.file` because it is in hold status. That same command to a data retention server still rejects the deletion of both files.

IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients: Installation and User's Guide

Chapter 10. Processing options

You can use defaults for processing options or you can tailor the processing options to meet your specific needs. This topic provides an overview of processing options and an options reference section that provides detailed information about each option.

Related concepts

"Using options with commands" on page 259

Related information

Reading syntax diagrams

Processing options overview

This section provides an overview of the types of options that you can use.

Tivoli Storage Manager uses *processing options* that you specify in your dsm.sys or client user-options file (dsm.opt) or on the command line to control communications, backup-archive processing, and other types of processing.
- Communication options
- Server and node options
- Backup and archive processing options
- Restore and retrieve processing options
- Scheduling options
- Format options
- Command processing options
- Authorization options
- Error processing options
- Transaction processing option
- Web client options
- Diagnostics options

Tivoli Storage Manager also includes a group of client command options that you can enter *only* on the command line with specific commands. You can override some of the options in your options file by entering them with appropriate backup-archive commands.

Related concepts

"Entering options with a command" on page 259

Related tasks

"Creating and modifying the client system-options file" on page 49

Communication options

You use communication options to specify how your client node communicates with a Tivoli Storage Manager server. This topic provides information about the types of communication options you can use.

For UNIX and Linux use one of the following communication protocols:
- TCP/IP
- Shared memory (AIX, HP-UX, Linux, and Solaris)

Use the commmethod option to specify the communication protocol.

Ask your Tivoli Storage Manager administrator for assistance in setting your communication options.

Related reference

"Commmethod" on page 281

TCP/IP options

To use the TCP/IP communication protocol, you must include the tcpserveraddress option in your client options file.

The other TCP/IP options have default values that you can modify if you want to change the default value. This topic provides information about the types of communication options you can use.

Table 71. TCP/IP options

Option	Description
httpport "Httpport" on page 343	Specifies a TCP/IP port address for the Tivoli Storage Manager Web client.
lanfreetcpport "Lanfreetcpport" on page 361	Specifies the TCP/IP port number where the Tivoli Storage Manager storage agent is listening.
lanfreetcpserveraddress "Lanfreetcpserveraddress" on page 361	Specifies the TCP/IP address for the Tivoli Storage Manager storage agent.
tcpbuffsize "Tcpbuffsize" on page 435	Specifies the size, in kilobytes, of the Tivoli Storage Manager internal TCP/IP communication buffer.
tcpnodelay "Tcpnodelay" on page 439	Specifies whether the server or client disables the delay of sending successive small packets on the network. This option is for all UNIX clients.
tcpadminport "Tcpadminport" on page 434	Specifies a separate TCP/IP port number on which the server is waiting for requests for administrative client sessions, allowing secure administrative sessions within a private network.
tcpcadaddress "Tcpcadaddress" on page 436	Specifies a TCP/IP address for dsmcad.
tcpport "Tcpport" on page 439	Specifies the TCP/IP port address for a Tivoli Storage Manager server.
tcpserveraddress "Tcpserveraddress" on page 440	Specifies the TCP/IP address for a Tivoli Storage Manager server.
tcpwindowsize "Tcpwindowsize" on page 441	Specifies the size, in kilobytes, of the TCP/IP sliding window for your client node.
webports "Webports" on page 455	Enables the use of the Web client outside a firewall by specifying the TCP/IP port number used by the client acceptor daemon and the Web Client Agent service (Web Client Agent service does not apply to Mac OS X) for communications with the Web GUI.

Related reference

"Nfstimeout" on page 373

Shared memory options

This topic provides information on the shared memory options that you can use.

Table 72. Shared memory communication options

Option	Description
lanfreeshmport "Lanfreeshmport" on page 360	Specifies the unique number that is used by the client and the storage agent to identify shared memory area used for communications.
lanfreeshmport "Shmport" on page 415	Specifies the unique number that is used by the client and the server to identify shared memory area used for communications.

Server and node options

Authorized User

This topic provides information about the server to contact for backup-archive services, and the client node to request backup-archive services.

Table 73. Server and node options

Option	Description
defaultserver "Defaultserver" on page 294	The name of the Tivoli Storage Manager server to contact for backup-archive services by default if more than one server is defined in the dsm.sys file.
	Also specifies the server to contact for space management services if you have the HSM client installed and do not specify a server with the migrateserver option. See *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide*, for more information.
asnodename "Asnodename" on page 268	Use the asnodename option to allow agent nodes to back up or restore data on behalf of another node (the target node). This enables concurrent operations from multiple nodes to store data to the same target node and file space in parallel.
nasnodename "Nasnodename" on page 371	Specifies the node name for the NAS file server when processing NAS file systems.
nodename "Nodename" on page 374	Use the nodename option in your dsm.sys file to identify your workstation to the server to establish communications.
servername "Servername" on page 412	In the dsm.sys file, this option specifies the name of a server and starts a server stanza. In the client user-options file (dsm.opt), this option specifies the Tivoli Storage Manager server to contact for services.
virtualnodename "Virtualnodename" on page 454	The virtualnodename option specifies the node name of your workstation when you want to restore or retrieve files to a different workstation.

Server options

Use the servername option in your dsm.sys file to begin a group of options (stanzas) used to connect to a Tivoli Storage Manager server.

You can set up multiple groups of stanzas in the dsm.sys file to connect to different servers. Each servername stanza must have listed below it all client option stanzas required to establish communication with a server. The stanza list can also contain other options for backup-archive operations.

If your client system-options file contains only one stanza - Your client node contacts the server you specify in that stanza for all services.

If your client system-options file contains more than one stanza - You can specify a default server with the defaultserver option. If you do not specify a default server, Tivoli Storage Manager contacts the server you specify in the first stanza of your dsm.sys file.

Place the defaultserver option at the beginning of your dsm.sys file before any server stanzas. See "Defaultserver" on page 294 for more information.

Use the servername option in the client user-options file (dsm.opt) or on the command line to specify a server to contact for backup-archive services. This overrides the default server specified in your (dsm.sys) file.

Note: You cannot override the migration server specified in the client system-options file.

Table 74 shows a sample dsm.sys file.

Table 74. Sample client system-options file

Sample dsm.sys file

```
DEFAULTServer               server2

SErvername     server1
  NODename                  node1
  COMMMethod                TCPip
  TCPPort                   1500
  TCPServeraddress          node.domain.company.com
  PASSWORDAccess            generate
  GRoups                    system adsm
  USERs                     ashton stewart kaitlin
  INCLExcl                  /adm/adsm/backup1.excl

SErvername     server2
  COMMMethod                SHAREdmem
  shmport                   1520
  PASSWORDAccess            prompt
  GRoups                    system adsm
  USERs                     danielle derek brant
  INCLExcl                  /adm/adsm/backup2.excl
```

Node options

This topic explains the node options that you can specify in your dsm.sys file to request backup-archive services.

Node options

This topic explains the node options that you can use to request backup-archive services.

Table 75. Node options

Option	Description
asnodename "Asnodename" on page 268	Use the asnodename option to allow agent nodes to back up or restore data on behalf of another node (the target node). This enables concurrent operations from multiple nodes to store data to the same target node and file space in parallel.
nodename "Nodename" on page 374	Use the nodename option in your client options file dsm.opt to identify your workstation to the server to establish communications.
virtualnodename "Virtualnodename" on page 454	The virtualnodename option specifies the node name of your workstation when you want to restore or retrieve files to a different workstation.

Backup and archive processing options

This topic describes the options that you can use to control some aspects of backup and archive processing.

Table 76. Backup and archive processing options

Option	Description
archmc "Archmc" on page 267	Use the archmc option with the **archive** command to specify the available management class for your policy domain to which you want to bind your archived files.
archsymlinkasfile "Archsymlinkasfile" on page 268	Specifies whether you want Tivoli Storage Manager to follow a symbolic link and archive the file or directory to which it points, or archive the symbolic link only.
asnodename "Asnodename" on page 268	Use the asnodename option to allow agent nodes to back up or restore data on behalf of another node (the target node). This enables concurrent operations from multiple nodes to store data to the same target node and file space in parallel.
automount "Automount" on page 276	Use this option with the domain option to specify all automounted file systems the Tivoli Storage Manager client tries to mount at the following points in time: • When Tivoli Storage Manager client starts • When the backup is started • When the Tivoli Storage Manager client has reached an automounted file system during backup
autofsrename "Autofsrename" on page 274	Specifies whether to rename an existing file space on a Unicode-enabled server so a Unicode-enabled file space can be created for the current operation.
changingretries "Changingretries" on page 278	Specifies the number of times a back up or archive action is retried on a file that is in use.
compressalways	The compressalways option specifies whether to continue compressing an object if it grows during compression. Use this option with the compression option.

Table 76. Backup and archive processing options (continued)

Option	Description
compression "Compression" on page 285	The compression option compresses files before you send them to the server. Compressing your files reduces data storage for backup versions and archive copies of your files. **Note:** The compression option also applies to migrated files if you install the Tivoli Storage Manager HSM client on your workstation.
createnewbase "Createnewbase" on page 287	The createnewbase option creates a base snapshot and uses it as a source to run a full incremental. This ensures that the backup of any files that might have been skipped during the snapshot difference incremental.
deduplication"Deduplication" on page 293	Specifies whether to eliminate redundant data on the client side when transferring data to theTivoli Storage Manager server, during backup or archive processing.
dedupcachepath"Dedupcachepath" on page 291	Specifies the location where the client-side data deduplication cache database is created, if the enablededupcache=yes option is set during backup or archive processing.
dedupcachesize"Dedupcachesize" on page 292	Determines the maximum size of the data deduplication cache file.
enablededupcache "Enablededupcache" on page 312	Specifies whether you want to enable client-side data deduplication cache, so that Tivoli Storage Manager gets the changed data from the cache.
ieobjtype"Ieobjtype" on page 345	Specifies an object type for a client-side data deduplication operation. This option is used with the include.dedup and exclude.dedup options.
deletefiles "Deletefiles" on page 295	Use the deletefiles option with the **archive** command to delete files from your workstation after you archive them. You can also use this option with the **restore image** command and the incremental option to delete files from the restored image if they were deleted after the image was created.
description "Description" on page 295	The description option assigns or specifies a description for files when performing archive, delete, retrieve, query archive, or query backupset operations.
detail "Detail" on page 296	Use the detail option to list management class, file space, backup, and archive information depending on the command with which it is used.
diffsnapshot "Diffsnapshot" on page 297	Use the diffsnapshot option to determine whether Tivoli Storage Manager creates a differential snapshot.
dirmc "Dirmc" on page 298	Specifies the management class to use for directories. If you do not specify this option, the client uses the management class in the active policy set of your policy domain with the longest retention period.
dirsonly "Dirsonly" on page 300	Backs up, restores, archives, retrieves, or queries directories *only*.

Table 76. Backup and archive processing options (continued)

Option	Description
diskcachelocation "Diskcachelocation" on page 302	Specifies the location where the disk cache database is created if the option memoryefficient=diskcachemethod option is set during an incremental backup.
domain "Domain" on page 303	Specifies the file systems to include in your default client domain for an incremental backup.
domain.image "Domain.image" on page 307	Specifies the mounted file systems and raw logical volumes that you want to include in your client domain for an image backup. This option is for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris *only*.
domain.nas "Domain.nas" on page 308	Specifies the volumes to include in your default domain for NAS image backups.
efsdecrypt "Efsdecrypt" on page 310	Specifies whether files encrypted by an AIX Encrypted File System (EFS) are read in encrypted or decrypted format.
enablearchiveretentionprotection "Enablearchiveretentionprotection" on page 311	Allows the client to connect to a data retention server.
enablelanfree "Enablelanfree" on page 314	Specifies whether to enable an available LAN-free path to a storage area network (SAN) attached storage device.
"Exclude options" on page 321 exclude exclude.backup exclude.file exclude.file.backup	Use these options to exclude a file or group of files from backup services and space management services (if the HSM client is installed). The exclude.backup option only excludes files from normal backup, but not from HSM.
encryptiontype "Encryptiontype" on page 315	Allows you to select AES 128 bit data encryption, providing a stronger form of data encryption than DES 56 bit data encryption.
encryptkey "Encryptkey" on page 316	Specifies whether to save the encryption key password locally when performing a backup-archive operation or whether to prompt for the encryption key password.
exclude.archive "Exclude options" on page 321	Excludes a file or a group of files that match the pattern from archive services *only*.
exclude.attribute.symlink "Exclude options" on page 321	Excludes a file or a group of files that are symbolic links or aliases (aliases apply to Mac OS X) from backup processing only.
exclude.compression "Exclude options" on page 321	Excludes files from compression processing if you set the compression option to *yes*. This option applies to backups and archives.
exclude.dir "Exclude options" on page 321	Excludes a directory, its files, and all its subdirectories and their files from backup processing.
exclude.encrypt "Exclude options" on page 321	Excludes specified files from encryption processing.
exclude.fs "Exclude options" on page 321	Excludes file spaces matching a pattern. This option is valid for all UNIX clients.

Table 76. Backup and archive processing options (continued)

Option	Description
exclude.fs.nas "Exclude options" on page 321	Excludes file systems on the NAS file server from an image backup when used with the **backup nas** command. This option is for AIX and Solaris clients *only*.
exclude.image "Exclude options" on page 321	Excludes mounted file systems and raw logical volumes that match the specified pattern from full image backup operations. This option is valid for AIX, HP-UX, Solaris, and all Linux clients *only*.
filelist "Filelist" on page 333	Specifies a list of files to be processed for the command. Tivoli Storage Manager opens the designated filelist and processes the files listed within according to the command.
filesonly "Filesonly" on page 336	Backs up, restores, retrieves, or queries files *only*.
groupname "Groupname" on page 341	Use this option with the **backup group** command to specify the fully qualified name of the group leader for a group.
guitreeviewafterbackup "Guitreeviewafterbackup" on page 342	Specifies whether the client GUI is returned to the Backup, Restore, Archive, or Retrieve window after an operation finishes.
ieobjtype"Ieobjtype" on page 345	Specifies an object type for a client-side data deduplication operation. This option is used with the include.dedup and exclude.dedup options.
imagegapsize "Imagegapsize" on page 347	Specifies the minimum size of empty regions on a volume that you want to skip during image backup. This option is valid for AIX JFS2 clients.
inclexcl "Inclexcl" on page 349	Specifies the path and file name of an include-exclude options file.
"Include options" on page 351 include include.backup include.file	Use these options to include files or assign management classes for backup processing.
include.archive "Include options" on page 351	Includes files or assigns management classes for archive processing.
include.attribute.symlink "Include options" on page 351	Includes a file or a group of files that are symbolic links or aliases (aliases apply to Mac OS X) within broad group of excluded files for backup processing only.
include.compression "Include options" on page 351	Includes files for compression processing if you set the compression option to *yes*. This option applies to backups and archives.
include.encrypt "Include options" on page 351	Includes the specified files for encryption processing. By default, Tivoli Storage Manager does not perform encryption processing.
include.fs "Include options" on page 351	Use the include.fs option to control how Tivoli Storage Manager processes your file space for incremental backup.

Table 76. Backup and archive processing options (continued)

Option	Description
include.fs.nas "Include options" on page 351	Use the include.fs.nas option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system image backup, using the toc option with the include.fs.nas option in your dsm.sys file. See "Toc" on page 444 for more information. This option is valid for AIX and Solaris clients *only*.
include.image "Include options" on page 351	Specifies a file system or logical volume to be included for image backup processing. This option also provides a way to specify an explicit management class assignment for a specified file system or logical volume. The backup image command ignores all other include options. This option is valid for AIX, HP-UX, Solaris, and all Linux clients.
incrbydate "Incrbydate" on page 357	Use with the **incremental** command to request an incremental backup by date.
incremental "Incremental" on page 358	Use with the **restore image** command to ensure that any changes that were made to the base image are also applied to the restored image. This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris *only*.
memoryefficientbackup "Memoryefficientbackup" on page 368	Specifies a memory-saving backup algorithm for incremental backups when used with the **incremental** command.
mode "Mode" on page 369	Use the mode option with these commands, as follows: **backup image** To specify whether to perform a selective or incremental image backup of client file systems. **backup nas** To specify whether to perform a full or differential image backup of NAS file systems. **backup group** To specify whether to perform a full or differential group backup containing a list of files from one or more file space origins.
monitor "Monitor" on page 371	Specifies whether you want to monitor an image backup of file systems belonging to a Network Attached Storage (NAS) file server.
noprompt "Noprompt" on page 376	suppresses the confirmation prompt that is presented by the **delete group**, **delete archive**, **expire**, and **set event** commands.
noprompt "Noprompt" on page 376	suppresses the confirmation prompt that is presented by the **delete group**, **delete archive**, **expire**, **restore image**, and **set event** commands.

Table 76. Backup and archive processing options (continued)

Option	Description
nojournal "Nojournal" on page 375	Use this option with the **incremental** command to specify that you want to perform the traditional full incremental backup, instead of the default journal-based backup.
optfile "Optfile" on page 378	Specifies the client user-options file you want to use when you start a Tivoli Storage Manager session.
postsnapshotcmd "Postsnapshotcmd" on page 386	During a snapshot-based backup, this option allows you to manually open an application after the snapshot is created. This option is valid only for AIX JFS2 or Linux LVM snapshot-based operations.
preservelastaccessdate "Preservelastaccessdate" on page 389	Use this option during a backup or archive operation to specify whether to reset the last access date of any specified files to their original value following a backup or archive operation. By default, the Tivoli Storage Manager client does not reset the last access date of any backed up or archived files to their original value before the backup or archive operation.
presnapshotcmd "Presnapshotcmd" on page 392	During a snapshot-based backup operation, this option allows you to manually quiesce an application before the snapshot is created. This option is valid only for AIX JFS2 or Linux LVM snapshot-based operations.
removeoperandlimit "Removeoperandlimit" on page 396	Specifies that Tivoli Storage Manager removes the 20-operand limit. If you specify the removeoperandlimit option with the **incremental**, **selective**, or **archive** commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits.
skipacl "Skipacl" on page 416	Specifies whether to skip ACL processing completely.
skipaclupdatecheck "Skipaclupdatecheck" on page 417	Specifies whether to perform checksum and size comparisons before and after backup and during incremental processing.
snapdiff "Snapdiff" on page 418	Specifies an incremental backup of the files reported as changed by NetApp, instead of scanning the volume looking for files that have changed. Use this option with a NAS full volume incremental backup.
snapshotcachesize "Snapshotcachesize" on page 422	Linux and AIX only: Use this option to specify an appropriate snapshot size so that all original data blocks can be stored during file modification and deletion. A snapshot size of 100 percent ensures a valid snapshot. This option is valid for AIX JFS2 file systems and Linux x86/x86_64 clients only and has a default value of 100 percent.
snapshotproviderfs "Snapshotproviderfs" on page 423	Use the snapshotproviderfs option to enable snapshot-based file backup and archive operations, and to specify a snapshot provider. You must be a root user to perform a snapshot-based file backup or archive operation. If you are not a root user, the operation fails with an error message.

Table 76. Backup and archive processing options (continued)

Option	Description
snapshotproviderimage "Snapshotproviderimage" on page 424	Use the snapshotproviderimage option to enable snapshot-based image backup, and to specify a snapshot provider. You must be a root user to perform a snapshot-based image backup operation. If you are not a root user, the operation fails with an error message.
snapshotroot "Snapshotroot" on page 425	Use the snapshotroot option with the **incremental**, **selective**, or **archive** commands in conjunction with a independent software vendor application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. This option is valid for all UNIX and Linux clients.
subdir "Subdir" on page 432	Specifies whether to include subdirectories of a named directory.
tapeprompt "Tapeprompt" on page 433	Specifies whether you want Tivoli Storage Manager to wait for a tape mount if it is required for a backup, archive, restore, or retrieve process, or to be prompted for a choice.
toc "Toc" on page 444	Use the toc option with the **backup nas** command or the include.fs.nas option to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup. If you save TOC information, you can use the **query toc** server command to determine the contents of a file system backup in conjunction with the **restore node** server command to restore individual files or directory trees. You can also use the Tivoli Storage Manager Web client to examine the entire file system tree and select files and directories to restore.
type "Type" on page 448	Use the type option with the **query node** command to specify the type of node to query.
v2archive "V2archive" on page 449	Use the v2archive option with the **archive** command to archive only files to the server. Tivoli Storage Manager will not process directories that exist in the path of the source file specification.
virtualfsname "Virtualfsname" on page 452 (does not apply to Mac OS X)	Use this option with the **backup group** command to specify the name of the container for the group on which you want to perform the operation.
virtualmountpoint "Virtualmountpoint" on page 452	Defines a virtual mount point for a file system if you want to consider files for backup that begin with a specific directory within that file system.

Restore and retrieve processing options

This topic describes the options that you can use for restoring and retrieving.

Table 77 lists the restore and retrieve processing options that are available.

Table 77. Restore and retrieve processing options

Option	Description
dirsonly "Dirsonly" on page 300	Qualifies the operation (backup, archive, restore, retrieve) to process directories alone.
disablenqr "Disablenqr" on page 300	Specifies whether the Tivoli Storage Manager Backup-Archive Client can use the no-query restore method for restoring files and directories from the server.
filelist "Filelist" on page 333	Specifies a file containing a list of files to be processed by the specified command.
filesonly "Filesonly" on page 336	Qualifies the operation (backup, archive, restore, retrieve) to process files alone.
followsymbolic "Followsymbolic" on page 336	Specifies whether you want to restore files to symbolic links or use a symbolic link as a virtual mount point.
fromdate "Fromdate" on page 338	Use the fromdate option with the fromtime option to specify a date and time from which you want to search for backups or archives during a restore, retrieve, or query operation.
fromnode "Fromnode" on page 339	Permits one node to perform commands for another node. A user on another node must use the **set access** command to permit you to query, restore, or retrieve files or images for the other node.
fromowner "Fromowner" on page 339	Displays file spaces for an alternate owner. Also specifies an alternate owner from which to restore or retrieve files.
fromtime "Fromtime" on page 340	Use the fromtime option with the fromdate option to specify a beginning time from which you want to search for backups or archives during a restore, retrieve or query operation.
guitreeviewafterbackup "Guitreeviewafterbackup" on page 342	Specifies whether the client GUI is returned to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.
ifnewer "Ifnewer" on page 346	Replaces an existing file with the latest backup version only if the backup version is newer than the existing file.
imagetofile "Imagetofile" on page 348	Use the imagetofile option with the **restore image** command to specify that you want to restore the source image to a file. You might need to restore the image to a file in the event of bad sectors present on the target volume, or if you want to do some manipulations with the image data. This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris *only*.
inactive "Inactive" on page 348	Displays a list of active and inactive files when used with the pick option.
latest "Latest" on page 362	Restores the most recent backup version of a file whether it is active or inactive.
localbackupset "Localbackupset" on page 363	Specifies whether the Tivoli Storage Manager GUI bypasses initial logon with the server to restore a local backup set on a standalone workstation.

Table 77. Restore and retrieve processing options (continued)

Option	Description
makesparsefile "Makesparsefile" on page 364 (does not apply to Mac OS X)	Use the makesparsefile option with the **restore** or **retrieve** commands to specify how sparse files are recreated.
monitor "Monitor" on page 371	Specifies whether you want to monitor an image restore of one or more file systems belonging to a Network Attached Storage (NAS) file server.
noprompt "Noprompt" on page 376	suppresses the confirmation prompt that is presented by the **delete group**, **delete archive**, **expire**, and **set event** commands.
noprompt "Noprompt" on page 376	suppresses the confirmation prompt that is presented by the **delete group**, **delete archive**, **expire**, **restore image**, and **set event** commands.
optfile "Optfile" on page 378	Specifies the client user-options file you want to use when you start a Tivoli Storage Manager session.
pick "Pick" on page 382	Creates a list of backup versions, images, or archive copies that match the file specification you enter. From the list, you can select the versions to process. Include the inactive option to view both active and inactive objects.
pitdate "Pitdate" on page 383	Use the pitdate option with the pittime option to establish a point in time for which you want to display or restore the latest version of your backups.
pittime "Pittime" on page 384	Use the pittime option with the pitdate option to establish a point in time for which you want to display or restore the latest version of your backups.
preservepath "Preservepath" on page 390	Specifies how much of the source path to reproduce as part of the target directory path when you restore or retrieve files to a new location.
replace "Replace" on page 397	Specifies whether to overwrite an existing file, or to prompt you for your selection when you restore or retrieve files.
showmembers "Showmembers" on page 416 (does not apply to Mac OS X)	Displays all members of a group.
subdir "Subdir" on page 432	Specifies whether you want to include subdirectories of a named directory.
tapeprompt "Tapeprompt" on page 433	Specifies whether you want Tivoli Storage Manager to wait for a tape required for a restore or retrieve to be mounted, or to prompt you for your choice.
todate "Todate" on page 445	Use the todate option with the totime option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation.
totime "Totime" on page 446	Use the totime option with the todate option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation.
type "Type" on page 448	Use the type option with the **query node** command to specify the type of node to query.

Table 77. Restore and retrieve processing options (continued)

Option	Description
verifyimage "Verifyimage" on page 451	Use the verifyimage option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log.

The following option is a backup-archive client option that applies only to HSM for Windows migrated files. Refer to *IBM Tivoli Storage Manager for Space Management for UNIX and Linux* for details about this option.

- Restoremigstate

Scheduling options

This topic discusses the options that you can use to regulate central scheduling. Tivoli Storage Manager uses scheduling options only when the Scheduler is running.

Table 78 lists the scheduling options that are available.

Table 78. Scheduling options

Option	Description
managedservices "Managedservices" on page 365	Specifies whether the Storage Manager client acceptor daemon manages the Web client, the scheduler, or both.
maxcmdretries "Maxcmdretries" on page 367	Specifies the maximum number of times the client scheduler attempts to process a scheduled command that fails.
postschedulecmd/postnschedulecmd "Postschedulecmd/Postnschedulecmd" on page 385	Specifies a command to process after running a schedule.
preschedulecmd/prenschedulecmd "Preschedulecmd/Prenschedulecmd" on page 387	Specifies a command to process before running a schedule.
queryschedperiod "Queryschedperiod" on page 394	Specifies the number of hours the client scheduler waits between attempts to contact the server for scheduled work.
retryperiod "Retryperiod" on page 401	Specifies the number of minutes the client scheduler waits between attempts to process a scheduled command that fails or between unsuccessful attempts to report results to the server.
schedcmddisabled "Schedcmddisabled" on page 402	Specifies whether to disable the scheduling of generic commands specified by your Tivoli Storage Manager administrator.

Table 78. Scheduling options (continued)

Option	Description
schedcmduser (server defined only) "Scheduler options for commands" on page 222	The scheduler executes commands under a uid of 0, however, there might be some users who have a different user ID. In this case, your Tivoli Storage Manager administrator can define schedules and allow these schedules to be executed under a uid other than 0, using this option. The Tivoli Storage Manager Client API does not support this option.
schedlogmax "Schedlogmax" on page 404	Specifies the maximum size of the scheduler log and Web client log, in megabytes.
schedlogname "Schedlogname" on page 405	Specifies the path and file name where you want to store schedule log information.
schedlogretention "Schedlogretention" on page 407	Specifies the number of days to keep log file entries in the schedule log and the Web client log, and whether to save pruned entries.
schedmode "Schedmode" on page 408	Specifies which schedule mode to use, *polling* or *prompted*.
schedrestretrdisabled "Schedrestretrdisabled" on page 409	Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing restore or retrieve schedule operations.
sessioninitiation "Sessioninitiation" on page 414	Use the sessioninitiation option to control whether the server or client initiates sessions through a firewall. The default is that the client can initiate sessions.
srvprepostscheddisabled "Srvprepostscheddisabled" on page 428	Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing pre-schedule and post-schedule commands when performing scheduled operations.
srvprepostsnapdisabled "Srvprepostsnapdisabled" on page 429	Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing pre-snapshot and post-snapshot commands when performing scheduled image snapshot backup operations.
tcpclientaddress "Tcpclientaddress" on page 437	Specifies a TCP/IP address if your client node has more than one address, and you want the server to contact an address other than the one that was used to make the first server contact. The server uses this address when it begins the server prompted scheduled operation. See schedmode *prompted* ("Schedmode" on page 408) for details.
tcpclientport "Tcpclientport" on page 438	Specifies a TCP/IP port number for the server to contact the client when the server begins the server prompted scheduled operation. See schedmode *prompted* ("Schedmode" on page 408) for details.

Format and language options

Format and language options allow you to select different formats for date, time, numbers, and for different languages if you have the appropriate client language pack installed for that language.

Format options allow you to select different formats for date, time, and numbers.

Table 79. Format and language options

Option	Description
dateformat "Dateformat" on page 289	Specifies the format for displaying dates.
numberformat "Numberformat" on page 376	Specifies the format for displaying numbers.
timeformat "Timeformat" on page 442	Specifies the format for displaying time.

Command processing options

This topic explains the options that you can use with the Tivoli Storage Manager commands.

Command processing options allow you to control some of the formatting of data on your terminal screen.

Table 80. Command processing options

Option	Description
quiet "Quiet" on page 396	Limits the number of messages that are displayed on your screen during processing. This option can be overridden by the server.
scrolllines "Scrolllines" on page 410	Specifies the number of lines of information that are displayed on your screen at one time. Use this option only when scrollprompt is set to *yes*.
scrollprompt "Scrollprompt" on page 411	Specifies whether you want Tivoli Storage Manager to stop and wait after displaying the number of lines of information you specified with the scrolllines option, or scroll through and stop at the end of the information list.
verbose "Verbose" on page 450	Specifies that processing information should be displayed on your screen. The alternative is quiet. This option can be overridden by the server.

Authorization options

Authorization options control access to a Tivoli Storage Manager server.

Table 81 lists the authorization options that are available.

Table 81. Authorization options

Option	Description
groups "Groups" on page 342	Specifies the groups on your workstation that you want to authorize to request Tivoli Storage Manager services from the server.

Table 81. Authorization options (continued)

Option	Description
password "Password" on page 379	Specifies a Tivoli Storage Manager password.
passwordaccess "Passwordaccess" on page 380	Specifies whether you want to use a generated password or be prompted for a password each time you start the client.
passworddir "Passworddir" on page 382	Specifies the directory in which you want to store the automatically generated password for your client node. The encryption key and password are encrypted and stored in the TSM.PWD file.
revokeremoteaccess "Revokeremoteaccess" on page 402	Restricts an administrator with client access privileges from accessing your workstation through the Web client.
users "Users" on page 448	Authorizes specific users on your workstation to request services from a server.

Error processing options

Error processing options specify the name of the error log file and how Tivoli Storage Manager treats the entries in the log file.

Table 82 lists the error processing options that are available.

Table 82. Error processing options

Option	Description
errorlogmax "Errorlogmax" on page 318	Specifies the maximum size of the error log, in megabytes.
errorlogname "Errorlogname" on page 318	Specifies the fully qualified path and file name of the file where you want to store information about errors that occur during processing.
errorlogretention "Errorlogretention" on page 319	Specifies how many days to maintain error log entries before pruning, and whether to save the pruned entries.

Transaction processing options

Transaction processing options control how Tivoli Storage Manager processes transactions between the client and server.

Table 83 lists the transaction processing options that are available.

Table 83. Transaction processing options

Option	Description
collocatebyfilespec "Collocatebyfilespec" on page 280	Specifies that you want the Tivoli Storage Manager client to use only one server session to send objects generated from one file specification. Setting the collocatebyfilespec option to *yes* eliminates interspersing of files from different file specifications, by limiting the client to one server session per file specification. Therefore, if you store the data to tape, files for each file specification are stored together on one tape (unless another tape is required for more capacity).

Table 83. Transaction processing options (continued)

Option	Description
commrestartduration "Commrestartduration" on page 283	Specifies the maximum number of minutes you want the client to try to reconnect to a Tivoli Storage Manager server after a communication error occurs.
commrestartinterval "Commrestartinterval" on page 283	Specifies the number of seconds you want the client to wait between attempts to reconnect to a Tivoli Storage Manager server after a communication error occurs.
diskbuffsize "Diskbuffsize" on page 301	Specifies the maximum disk I/O buffer size (in kilobytes) that the client can use when reading files.
largecommbuffers "Diskbuffsize" on page 301	This option has been replaced by the diskbuffsize option. At this time, largecommbuffers is still accepted by the Tivoli Storage Manager client in order to ease the transition to the new option. However, the value specified by largecommbuffers is ignored in favor of the diskbuffsize setting. **Important:** Discontinue the use of largecommbuffers because future releases of Tivoli Storage Manager might not accept this option.
nfstimeout "Nfstimeout" on page 373	Specifies the number of seconds the server waits for a status system call on an NFS file system before it times out.
resourceutilization "Resourceutilization" on page 398	Use the resourceutilization option in your dsm.sys file to regulate the level of resources the Tivoli Storage Manager server and client can use during processing.
txnbytelimit "Txnbytelimit" on page 446	Specifies the number of kilobytes the client program buffers before it sends a transaction to the server.

Web client options

This topic describes the options for the Tivoli Storage Manager Web Client.

Table 84 lists the Web client options that are available.

Table 84. Web client options

Option	Description
httpport "Httpport" on page 343	Specifies a TCP/IP port address for the Web client.
managedservices "Managedservices" on page 365	Specifies whether the Storage Manager client acceptor daemon manages the Web client, the scheduler, or both.
revokeremoteaccess "Revokeremoteaccess" on page 402	Restricts administrator access on a client workstation through the Web client.
webports "Webports" on page 455	Enables the use of the Web client outside a firewall by specifying the TCP/IP port number used by the client acceptor daemon and the Web Client Agent service for communications with the Web client.

Diagnostics options

Use the **query systeminfo** command to gather Tivoli Storage Manager system information and output this information to a file or the console.

The **query systeminfo** command is intended primarily as a diagnostic aid. You can submit the resulting information to technical support personnel for problem diagnosis.

Table 85 lists the diagnostics options that are available.

Table 85. Diagnostics options

Option	Description
console "Console" on page 286	Use the console option with the **query systeminfo** command to output system information to the console.
filename "Filename" on page 335	Use the filename option with the **query systeminfo** command to specify a file name in which to store the system information.

Related reference
"Query Systeminfo" on page 522

Using options with commands

You can override some of the options in your client options file (dsm.opt) file by entering them with appropriate Tivoli Storage Manager commands.

You can override some of the options in your dsm.sys file or client user-options file (dsm.opt) by entering them with appropriate Tivoli Storage Manager commands.

Tivoli Storage Manager processes options in the following order (precedence):
1. Options defined on the server with server-enforced client options. The server overrides client values.
2. Options entered locally on the command line.
3. Options defined on the server for a schedule using the options parameters.
4. Options entered locally in the options file.
5. Options received from the server with client option sets not set as forced by the server. The server *does not* override client values if not forced.
6. Default option values.

Tivoli Storage Manager also includes a group of client command options that you can enter *only* on the command line with specific commands. For a complete list of command-line options, a description, and where to go for more information, see Table 86 on page 260.

Entering options with a command

This topic lists the general rules for entering options with a command.
- Enter a command, a dash (–), the option name, an equal sign (=), and the option value or parameter. There should be no spaces on either side of the = sign.
- For example, enter this command:

  For example, enter this command:

      dsmc archive -description="year end 1999" /home/
- For options that do not include parameters, enter a command, a dash (–) and the option name. For example,

      dsmc incremental -quiet

  Note: Use a leading dash (-) to indicate that the following text is the name of an option. If an object name begins with a dash, you must surround it in either

single quotation marks (') or quotation marks ("). Most operating system command-line processors strip the quotation marks before submitting the command-line arguments to the Tivoli Storage Manager client application. In such cases, using escape characters or doubling the quotation marks allows the client to receive the quoted object name. In loop mode, surround such objects in either single quotation marks (') or quotation marks (").

- Enter either the option name, or an abbreviation for the option name. For example, to enter the latest option, enter either `-lat` or `-latest`. The capital letters in the syntax of each option indicate the minimum abbreviation for that option name.
- Enter options before or after command parameters. For example, you can enter the option before or after a file specification:

  ```
  dsmc selective -subdir=yes "/home/devel/proj1/*"
  dsmc selective "/home/devel/proj1/*" -subdir=yes
  ```

- When entering several options on a command, separate them with a blank space.
- Enclose the value in quotation marks (" ") if the option value that you enter contains a blank space. For example:

  ```
  dsmc archive -description="Project A" "/home/devel/proj1/*"
  ```

- Most options that you enter on the command line override the value set in the preferences file. However, when you use the domain option with the **incremental** command, it adds to the domain specified in your client options file rather than overriding the current value.
- The maximum number of characters for a file name is 255. The maximum combined length of the file name and path name is 1024 characters. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.
- For Mac OS X, the maximum length of a file name is limited to 504 bytes (not characters). The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name contains can vary.

Table 86 lists client command options that you can enter only on the command line with specific commands.

Table 86. Client command options

Command option	Description	Commands
archmc "Archmc" on page 267	Use the archmc option with the **archive** command to specify the available management class for your policy domain to which you want to bind your archived files.	**archive**
class "Class" on page 279	Specifies whether to display a list of NAS objects or client objects when using the following commands:	**query backup** **delete filespace** **query filespace**
console "Console" on page 286	Use the console option with the **query systeminfo** command to output system information to the console.	**query systeminfo**
deletefiles "Deletefiles" on page 295	Deletes the local copy of files from your workstation after they are archived on the server.	**archive**
deletefiles "Deletefiles" on page 295	Deletes the local copy of files from your workstation after they are archived on the server. Can also be used with the **restore image** command and the incremental option to delete files from the restored image that are deleted from the file space after the image is created.	**archive** **restore image**

Table 86. Client command options (continued)

Command option	Description	Commands
description "Description" on page 295	Assigns or specifies a description for files when performing archive, delete, retrieve, or query archive operations.	**archive** **delete archive** **query archive** **query backupset** **retrieve**
detail "Detail" on page 296	Displays management class, file space, backup, and archive information depending on the command with which it is used.	**delete filespace** **query archive** **query backup** **query filespace** **query mgmtclass**
dirsonly "Dirsonly" on page 300	Backs up, restores, archives, retrieves, or queries directories *only*.	**archive** **incremental** **query archive** **query backup** **restore** **restore backupset** **retrieve** **selective**
dynamicimage "Dynamicimage" on page 309	Performs a dynamic image backup.	**backup image**
filelist "Filelist" on page 333	Specifies a list of files to be processed for the command. Tivoli Storage Manager opens the designated filelist and processes the files listed within according to the command.	**archive** **backup group** **delete archive** **delete backup** **expire** **incremental** **query archive** **query backup** **restore** **retrieve** **selective**
filename "Filename" on page 335	Use the filename option with the **query systeminfo** command to specify a file name in which to store the system information.	**query systeminfo**
filesonly "Filesonly" on page 336	Backs up, restores, retrieves, or queries files *only*.	**archive** **incremental** **query archive** **query backup** **restore** **restore backupset** **retrieve** **selective**

Table 86. Client command options (continued)

Command option	Description	Commands
fromdate "Fromdate" on page 338	Use the fromdate option with the fromtime option to specify a date and time from which you want to search for backups or archives during a restore, retrieve, or query operation.	**delete backup** **query archive** **query backup** **restore** **restore group** **retrieve**
fromnode "Fromnode" on page 339	Permits one node to perform commands for another node. A user on another node must use the **set access** command to permit you to query, restore, or retrieve files or images for the other node.	**query archive** **query backup** **query filespace** **query group** **query image** **query mgmtclass** **restore** **restore group** **restore image** **retrieve**
fromowner "Fromowner" on page 339	Displays file spaces for an alternate owner. Also specifies an alternate owner from which to restore or retrieve files.	**query archive** **query backup** **query group** **query image** **restore** **restore group** **restore image** **retrieve**
fromtime "Fromtime" on page 340	Specifies a beginning time on the specified date. Use with the fromdate option. This option is ignored if the fromdate option is absent.	**query archive** **query backup** **restore** **restore group** **retrieve**
groupname "Groupname" on page 341	Specifies the fully qualified name for a group.	**backup group**
ifnewer "Ifnewer" on page 346	Replaces existing files with the latest backup version only if the backup version is newer than the existing version.	**restore** **restore backupset** **restore group** **retrieve**
imagetofile "Imagetofile" on page 348	Use the imagetofile option with the **restore image** command to specify that you want to restore the source image to a file. You might need to restore the image to a file in the event of bad sectors present on the target volume, or if you want to do some manipulations with the image data. This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris *only*.	**restore image**

Table 86. Client command options (continued)

Command option	Description	Commands
inactive "Inactive" on page 348	Displays a list of active and inactive files when used with the pick option.	**delete group** **query backup** **query group** **query image** **query nas** **restore** **restore group** **restore image** **restore nas**
incrbydate "Incrbydate" on page 357	Requests an incremental backup by date.	**incremental**
incremental "Incremental" on page 358	Applies changes to the base image using information from incremental backups made after the original image backup. This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris *only*.	**restore image**
latest "Latest" on page 362	Restores the most recent backup version of a file whether it is active or inactive.	**restore** **restore group**

mode "Mode" on page 369	Use the mode option with these commands, as follows:
	
	backup image
	    To specify whether to perform a selective or incremental image backup of client file systems.
	
	backup nas
	    To specify whether to perform a full or differential image backup of NAS file systems.
	
	backup group
	    To specify whether to perform a full or differential group backup containing a list of files from one or more file space origins.
		backup group
		backup nas
		backup image
		restore nas

monitor "Monitor" on page 371	Specifies whether you want to monitor an image backup or restore of one or more file systems belonging to a Network Attached Storage (NAS) file server. Specifies whether you want to monitor a restore of one or more file systems belonging to a Network Attached Storage (NAS) file server.	**backup nas** **restore nas**
nojournal "Nojournal" on page 375	Use this option with the with the **incremental** command to specify that you want to perform the traditional full incremental backup, instead of the default journal-based backup.	**incremental**
noprompt "Noprompt" on page 376	suppresses the confirmation prompt that is presented by the **delete group**, **delete archive**, **expire**, and **set event** commands.	**delete archive** **delete backup** **delete group** **expire**
noprompt "Noprompt" on page 376	suppresses the confirmation prompt that is presented by the **delete group**, **delete archive**, **expire**, **restore image**, and **set event** commands.	**delete archive** **delete backup** **delete group** **expire** **restore image**

Table 86. Client command options (continued)

Command option	Description	Commands
optfile "Optfile" on page 378	Specifies the client user-options file you want to use when you start a Tivoli Storage Manager session.	**dsmc**
pick "Pick" on page 382	Creates a list of backup versions, images, or archive copies that match the file specification you enter. From the list, you can select the versions to process. Include the inactive option to view both active and inactive objects.	**delete archive** **delete group** **expire** **query nas** **restore** **restore group** **restore image** **restore nas** **retrieve**
pitdate "Pitdate" on page 383	Use the pitdate option with the pittime option to establish a point in time for which you want to display or restore the latest version of your backups.	**query backup** **query group** **query image** **query nas** **restore** **restore group** **restore image** **restore nas**
pittime "Pittime" on page 384	Use the pittime option with the pitdate option to establish a point in time for which you want to display or restore the latest version of your backups.	**query backup** **query image** **query nas** **restore** **restore image** **restore nas**
preservepath "Preservepath" on page 390	Specifies how much of the source path to reproduce as part of the target directory path when you restore or retrieve files to a new location.	**restore** **restore backupset** **restore group** **retrieve**
removeoperandlimit "Removeoperandlimit" on page 396	Specifies that Tivoli Storage Manager removes the 20-operand limit. If you specify the removeoperandlimit option with the **incremental**, **selective**, or **archive** commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits.	**incremental** **selective** **archive**
showmembers "Showmembers" on page 416	Displays all members of a group.	**query group** **restore group**
todate "Todate" on page 445	Use the todate option with the totime option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation.	**query archive** **query backup** **restore** **restore group** **retrieve**

Table 86. Client command options (continued)

Command option	Description	Commands
totime "Totime" on page 446	Use the totime option with the todate option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation.	**query archive** **query backup** **restore** **restore group** **retrieve**
type "Type" on page 448	Use the type option with the **query node** command to specify the type of node to query.	**query node**
v2archive "V2archive" on page 449	Use the v2archive option with the **archive** command to archive only files to the server. Tivoli Storage Manager will not process directories that exist in the path of the source file specification.	**archive**
verifyimage "Verifyimage" on page 451	Use the verifyimage option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log. This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris *only*.	**restore image**
virtualfsname "Virtualfsname" on page 452	Specifies the name of the virtual file space for the group on which you want to perform the operation.	**backup group**

Initial command-line-only options

There is a subset of client options that are valid on the initial command line only. Many of these options establish the runtime environment, such as the commmethod and optfile options. Options in this category are not valid in interactive, macro, or scheduler modes. They generate an error and cause processing to stop.

Table 87 lists the options that are valid only on the initial command line.

Table 87. Options that are valid on the initial command line only

Options valid on the initial command line

commmethod	preschedulecmd/prenschedulecmd (can be included in the schedule definition)
deduplication	
diskbuffsize	queryschedperiod
editor	resourceutilization
enablededupcache	retryperiod
enablelanfree	schedlogmax
errorlogmax	schedlogname
errorlogname	schedlogretention
errorlogretention	schedmode
lanfreecommmethod	servername
lanfreeshmport	sessioninitiation
lanfreetcpport	tcpbuffsize
maxcmdretries	tcpcadaddress
nfstimeout	tcpclientaddress
nodename	tcpclientport
optfile	tcpwindowsize
password	txnbytelimit
postschedulecmd/postnschedulecmd (can be included in the schedule definition)	virtualnodename

Client options that can be set by the Tivoli Storage Manager server

Some client options can be set by the Tivoli Storage Manager server.

Table 88 lists the options that can be set by the server.

Table 88. Options that can be set by the Tivoli Storage Manager server

Options that can be set by the Tivoli Storage Manager server

- "Archsymlinkasfile" on page 268
- "Changingretries" on page 278
- "Collocatebyfilespec" on page 280
- "Compressalways" on page 284
- "Compression" on page 285
- "Dirmc" on page 298
- "Disablenqr" on page 300
- "Diskcachelocation" on page 302
- "Domain" on page 303
- "Domain.image" on page 307
- "Domain.nas" on page 308
- "Encryptiontype" on page 315
- "Encryptkey" on page 316
- "Exclude options" on page 321
- "Inclexcl" on page 349
- "Include options" on page 351
- "Managedservices" on page 365
- MAXCANDPROCS
- "Maxcmdretries" on page 367
- MAXMIGRATORS
- "Memoryefficientbackup" on page 368
- "Nfstimeout" on page 373
- "Postschedulecmd/Postnschedulecmd" on page 385
- "Postsnapshotcmd" on page 386
- "Preschedulecmd/Prenschedulecmd" on page 387
- "Preservelastaccessdate" on page 389
- "Presnapshotcmd" on page 392

- "Queryschedperiod" on page 394
- "Quiet" on page 396
- "Resourceutilization" on page 398
- "Retryperiod" on page 401
- "Schedmode" on page 408
- "Scrolllines" on page 410
- "Scrollprompt" on page 411
- "Snapshotcachesize" on page 422
- "Snapshotproviderfs" on page 423
- "Snapshotproviderimage" on page 424
- "Subdir" on page 432
- "Tapeprompt" on page 433
- "Txnbytelimit" on page 446
- "Verbose" on page 450
- VSSALTSTAGINGDIR [2]

Note:

1. See Tivoli Storage Manager for Space Management for UNIX and Linux, client command line interface help or user's guide.

2. See *Tivoli Storage Manager for Mail, Data Protection for Microsoft Exchange Server Installation and User's Guide*.

Related information

> For information about creating client options sets on the Tivoli Storage Manager server, see the Tivoli Storage Manager Administrator's Guide for your server platform.

Client options reference

The following sections contain detailed information about each of the Tivoli Storage Manager processing options.

Information for each option includes the following information:
- a description
- a syntax diagram
- detailed descriptions of the parameters
- examples of using the option in the client options file (if applicable)
- examples of using the option on the command line (if applicable)

Options with a command-line example of **Does not apply** cannot be used with command line or scheduled commands.

Note:
1. Do not enclose an option value with single or quotation marks, unless the value is a file specification that contains spaces or wildcard characters. For example, the following option is not valid:

   ```
   passwordaccess  "generate"
   ```
2. All options in the dsm.sys file, except for the defaultserver option, must be placed within a server stanza. A server stanza is a collection of options statements in dsm.sys that begins with a SERVERName option and ends either at the next SERVERName option or the end of the file.

Archmc

Use the archmc option with the **archive** command to specify the available management class for your policy domain to which you want to bind your archived files and directories.

When you archive a file, you can override the assigned management class using the archmc option on the **archive** command or by using the Web client. Overriding the management class using the Web client is equivalent to using the archmc option on the **archive** command.

If you do not use the archmc option, the server binds archived directories to the default management class. If the default management class has no archive copy group, the server binds archived directories to the management class with the shortest retention period.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──ARCHMc =──managementclass──────────────────────────────────────►◄
```

Parameters

managementclass
> Specifies an available management class in the active policy set of your policy domain. This management class overrides the default management class and any include statements for the files and directories you are archiving.

Examples

Command line:
```
dsmc archive –archmc=ret2yrs /Users/van/Documents/budget.jan
dsmc archive –archmc=ret2yrs /home/plan/proj1/budget.jan
```

Archsymlinkasfile

The archsymlinkasfile option specifies whether Tivoli Storage Manager follows a symbolic link and archives the file or directory to which it points, or archives the symbolic link only. Use this option with the **archive** command.

Supported Clients

This option is valid for all UNIX clients except Mac OS X. The server can also define this option.

Options File

Place this option in the client user options file (dsm.opt).

Syntax

```
                         ┌─Yes─┐
►►──ARCHSYMLinkasfile────┼─────┼──────────────────────────────────►◄
                         └─No──┘
```

Parameters

Yes
> Specifies that Tivoli Storage Manager follows a symbolic link and archives the associated file or directory. This is the default.

No Specifies that Tivoli Storage Manager archives the symbolic link and not the associated file or directory.

Examples

Options file:
```
archsymlinkasfile no
```

Command line:
```
-archsyml=no
```

Asnodename

Use the asnodename option to allow agent nodes to back up or restore data on behalf of another node (the target node). This enables concurrent operations from multiple nodes to store data to the same target node and file space in parallel.

Your client node must be granted access to the target node by the Tivoli Storage Manager server administrative client **grant proxynode** command, and you must be a root user to use the asnodename option.

When the Tivoli Storage Manager administrator grants a node proxy authority, and you use the asnodename option to become that node, you can query and restore all files as if you had root authority.

An agent node is a client node that has been granted authority to perform client operations on behalf of a target node.

A target node is a client node that grants authority to one or more agent nodes to perform client operations on its behalf.

For example, you can use the following command to back up shared data for file space stored under the node name MyCluster:

```
/cluster1/mydata
dsmc incremental /Users -asnodenodename=MyCluster
```

You can also use the asnodename option to restore data under another node name on the Tivoli Storage Manager server. You can only restore the data that you own.

The asnodename option differs from the nodename option as follows:
- When using the nodename option, you must enter the password for the node name you specify.
- When using the asnodename option, you must enter the password for your client agent node to access the data stored for the client target node.

Restrictions: You cannot use the asnodename option with -fromnode and you cannot perform NAS backup using asnodename. Also, asnodename can be used for clustered systems, although no specific cluster software is supported.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **General** tab of the Preferences editor.

Syntax

►►—ASNODEname— —*targetnode*———————————————————————————►◄

Parameters

targetnode
> Specifies the node name on the Tivoli Storage Manager server under which you want to back up or restore data.

Examples

Options file:
```
asnodename mycluster
```

Command line:
```
-asnodename=mycluster
```

This option is not valid in interactive mode, but it can be defined in the options portion of a schedule definition.

Auditlogging

Use the auditlogging option to generate an audit log that contains an entry for each file that is processed during an incremental, selective, archive, restore, or retrieve operation.

The audit log can be configured to capture either a basic level of information or a more inclusive (full) level of information.

The basic level of the audit logging feature captures the information that is in the schedule log and it records information that a file has been backed up, archived, updated, restored, retrieved, expired, deleted, skipped or failed during an incremental backup, selective backup, archive, restore or retrieve operation. In addition, the basic level of audit logging captures the input command for commands run through the Backup-Archive command line or scheduler clients.

The full level of audit logging records an action for each file that is processed by the Backup-Archive client. In addition to all of the events recorded by the basic level of audit logging, the full level of audit logging records information for a file that has been excluded or not sent during a progressive incremental backup operation because the file had not changed.

The following is an example of the messages that are issued when the audit log is configured to capture the basic level of information:
```
04/21/07 15:25:05 ANS1650I Command:
  sel /home/spike/test/*
04/21/07 15:25:05 ANS1651I Backed Up:
  /home/spike/test/file.txt
04/21/07 15:25:05 ANS1652I Archived:
  /home/spike/test/file.txt
04/21/07 15:25:05 ANS1653I Updated:
  /home/spike/test/file.txt
04/21/07 15:25:05 ANS1654E Failed:
  /home/spike/test/file.txt
04/21/07 15:25:05 ANS1655I Restored:
  /home/spike/test/file.txt
04/21/07 15:25:05 ANS1656I Retrieved:
  /home/spike/test/file.txt
04/21/07 15:25:05 ANS1657I Expired:
  /home/spike/test/file.txt
04/21/07 15:25:05 ANS1658I Deleted:
  /home/spike/test/file.txt
04/21/07 15:25:05 ANS1659I Skipped:
  /home/spike/test/file.txt
```

The following messages can be issued when the audit log is configured to capture the full level of information (in addition to all messages issued for the basic level of audit logging):
```
04/21/07 15:25:05 ANS1660I Excluded:
  /home/spike/test/file.txt
04/21/07 15:25:05 ANS1661I Unchanged:
  /home/spike/test/file.txt
```

The audit log is not a substitute or a replacement for the standard error log (dsmerror.log) or for the schedule log (dsmsched.log). If an error occurs that prevents a file from being processed, a message indicating that an error has occurred is written to the audit log, but the message will not indicate the nature of the error. For problem diagnostics the standard error log must still be used.

The audit log entries only contain a time stamp and object name. There is no information to distinguish between files and directories or any information about the size of an object.

The Mac OS X Backup-Archive Client creates the audit log as a Unicode (UTF-16) file.

By default, the name of the audit log is dsmaudit.log and it is contained in the same directory as the error log, dsmerror.log. The name and location of the audit log can be configured using the auditlogname option. There are no parameters to control the size of the audit log or to prune the audit log. The auditlogname option cannot be set as an option in a Tivoli Storage Manager Server client options set.

The **auditlogging** command is supported with backup commands that interact with file-level objects such as **backup groups**.

The **auditlogging** command is not supported with backup commands which interact with image-level objects such as **backup image** or **restore image**. The **auditlogging** command is supported with backup commands that interact with file-level objects such as **backup groups**.

If you have enabled audit logging for an operation and there is a failure trying to write to the audit log (for example, the disk on which the audit log resides is out of space), the audit logging is disabled for the rest of the operation and the return code for the operation is set to 12, regardless of the outcome of the operation.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza.

Syntax

```
►►──AUDITLOGGing──┬──────┬───────────────────────────────────►◄
                  ├─off──┤
                  ├─basic─┤
                  └─full─┘
```
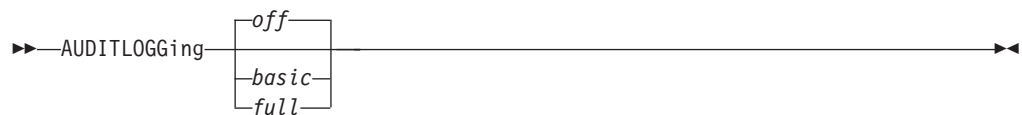
Parameters

off   Specifies that the audit logging facility is not engaged. This is the default.

basic
    Specifies that the audit log captures a basic level of information.

full
    Specifies that the audit log captures a more extensive level of information.

Examples

Run an incremental backup with audit logging enabled.

Command line:

    dsmc i -auditlogging=basic

Back up a list of files using the maximum level of auditing, which enables
a separate application, such as a Perl script, to verify the results.

Auditlogname

The auditlogname option specifies the path and file name where you want to store
audit log information. This option applies when audit logging is enabled.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza.

Syntax

▶▶──AUDITLOGName──*filespec*───◀◀

Parameters

filespec

Specifies the path and file name where you want Tivoli Storage Manager to
store audit log information.

If you specify a file name only, the file is stored in your current directory. The
default is the installation directory with a file name of `dsmaudit.log`. The
`dsmaudit.log` file *cannot* be a symbolic link.

Examples

Run an incremental backup with audit logging enabled.

Sample output

The following is a sample execution and output file:

    > dsmc inc /SMSVT/mfs1 -auditlogging=full
      -auditlogname=/home/cliv3/audit.log
    IBM Tivoli Storage Manager
    Command Line Backup/Archive Client Interface
      Client Version 5, Release 5, Level 0.0
      Client date/time: 07/03/07   12:05:13
    (c) Copyright by IBM Corporation and other(s)
        1990, 2007. All Rights Reserved.

    Node Name: NAXOS_CLUSTER
    Session established with server
      ODINHSMSERV: AIX-RS/6000
      Server Version 5, Release 4, Level 0.0
      Server date/time: 07/03/07   12:05:18
      Last access: 07/03/07   12:01:57

```
            Incremental backup of volume '/SMSVT/mfs1'
            Directory-->              4,096 /SMSVT
              /mfs1/ [Sent]
            Normal File-->           32,768 /SMSVT
              /mfs1/test0 [Sent]
            Normal File-->           32,768 /SMSVT
              /mfs1/test1 [Sent]
            Normal File-->           32,768 /SMSVT
              /mfs1/test2 [Sent]
            Normal File-->           32,768 /SMSVT
              /mfs1/test3 [Sent]
            Normal File-->           32,768 /SMSVT
              /mfs1/test4 [Sent]
            Normal File-->           32,768 /SMSVT
              /mfs1/test5 [Sent]
            Normal File-->           32,768 /SMSVT
              /mfs1/test6 [Sent]
            Normal File-->           32,768 /SMSVT
              /mfs1/test7 [Sent]
            Normal File-->           32,768 /SMSVT
              /mfs1/test8 [Sent]
            Normal File-->           32,768 /SMSVT
              /mfs1/test9 [Sent]
            Successful incremental backup of '/SMSVT/mfs1'

            Total number of objects inspected:      11
            Total number of objects backed up:      11
            Total number of objects updated:         0
            Total number of objects rebound:         0
            Total number of objects deleted:         0
            Total number of objects expired:         0
            Total number of objects failed:          0
            Total number of bytes transferred:    320.31 KB
            Data transfer time:                     0.01 sec
            Network data transfer rate:        17,141.84 KB/sec
            Aggregate data transfer rate:         297.43 KB/sec
            Objects compressed by:                    0%
            Elapsed processing time:             00:00:01
```

The following are the audit log contents:

```
07/03/07   12:05:14 ANS1650I Command:
  inc /SMSVT/mfs1
07/03/07   12:05:15 ANS1651I Backed Up:
  /SMSVT/mfs1/
07/03/07   12:05:15 ANS1651I Backed Up:
  /SMSVT/mfs1/test0
07/03/07   12:05:15 ANS1651I Backed Up:
  /SMSVT/mfs1/test1
07/03/07   12:05:15 ANS1651I Backed Up:
  /SMSVT/mfs1/test2
07/03/07   12:05:15 ANS1651I Backed Up:
  /SMSVT/mfs1/test3
07/03/07   12:05:15 ANS1651I Backed Up:
  /SMSVT/mfs1/test4
07/03/07   12:05:15 ANS1651I Backed Up:
  /SMSVT/mfs1/test5
07/03/07   12:05:15 ANS1651I Backed Up:
  /SMSVT/mfs1/test6
07/03/07   12:05:15 ANS1651I Backed Up:
  /SMSVT/mfs1/test7
07/03/07   12:05:15 ANS1651I Backed Up:
  /SMSVT/mfs1/test8
07/03/07   12:05:15 ANS1651I Backed Up:
  /SMSVT/mfs1/test9
```

Related information

For more information about the audit logging facility refer to "Auditlogging" on page 270.

Autofsrename

The autofsrename option renames an existing file space that is not Unicode-enabled on the Tivoli Storage Manager server so that a Unicode-enabled file space with the original name can be created for the current operation.

Note: You must be an authorized user to use this option.

When you specify `autofsrename yes` in your client options file, and the server value of autofsrename is set to `client`, Tivoli Storage Manager generates a unique name by appending `_OLD` to the file space name you specify in the current operation. For example, Tivoli Storage Manager renames the file space `Jaguar` to `Jaguar_OLD`. If the new file space name is too long, the suffix replaces the last characters of the file space name. For example, the `mylongfilesystemname` file space name is renamed to:

    mylongfilesystem_OLD

If the new file space name already exists on the server, Tivoli Storage Manager renames the new file space `Jaguar_OLDx`, where `x` is a unique number.

Tivoli Storage Manager creates new Unicode-enabled file spaces that contain only the data specified in the current operation. For example, assume that `Jaguar` is the name of your startup disk and you archive all of the `.log` files in the `/Users/user5/Documents` directory. Before the archive takes place, the server renames the file space to `Jaguar_OLD`. The archive places the data specified in the current operation into the Unicode-enabled file space named `Jaguar`. The new Unicode-enabled file space now contains only the `/Users/user5/logs` directory and the `*.log` files specified in the operation. Tivoli Storage Manager stores all subsequent full and partial incremental, selective backup, and archive data in the new Unicode-enabled file spaces.

For example, assume that `Jaguar` is the name of your startup disk and you archive all of the `.log` files in the `/Users/user5/Documents` directory. Before the archive takes place, the server renames the file space to `Jaguar_OLD`. The archive places the data specified in the current operation into the Unicode-enabled file space named `Jaguar`. The new Unicode-enabled file space now contains only the `/Users/user5/logs` directory and the `*.log` files specified in the operation. Tivoli Storage Manager stores all subsequent full and partial incremental, selective backup, and archive data in the new Unicode-enabled file spaces.

Renamed file spaces remain on the server as stabilized file spaces. *These file spaces contain all the original data, which you can restore as long as they remain on the server.*

Note: When an existing file space is renamed during Unicode conversion, any access rules defined for the file space remain applicable to the original file space. New access rules must be defined to apply to the new Unicode file space.

After installation, perform a full incremental backup and rename all existing file spaces that are not Unicode-enabled and back up the files and directories within them under the new Unicode-enabled file spaces. This operation requires increased processing time and storage on the server.

File spaces that are not Unicode-enabled can be viewed in the character set of the locale from which Tivoli Storage Manager backed up the files. A workstation running in a different locale might be unable to view or restore from these file spaces. Unicode-enabled file spaces that are backed up in one locale are visible in all other locales, provided that the workstation has the proper fonts installed.

The server can define the autofsrename option and override the autofsrename setting on the client.

Supported Clients

This option is valid for Mac OS X only. The server can define the autofsrename option and override the autofsrename setting on the client. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **General** tab, **Rename non-Unicode filespaces during backup/archive** drop-down list box of the Preferences editor.

Syntax

```
                    ┌─Prompt─┐
►►──AUTOFsrename──────┼────────┼──────────────────────────────────────►◄
                    ├─Yes────┤
                    └─No─────┘
```

Parameters

Yes
　　Specifies that Tivoli Storage Manager automatically renames all file spaces that are not Unicode-enabled in the current backup or archive operation.

No　Specifies that Tivoli Storage Manager does not rename file spaces that are not Unicode-enabled in the current backup or archive operation.

Prompt
　　Specifies that Tivoli Storage Manager prompts you whether to rename the file spaces that are not Unicode-enabled in the current operation. This is the default.

Considerations:
- This option applies *only* when the server sets the autofsrename option to `client`.
- When the client scheduler is running, the default behavior is to not prompt you. The next interactive session prompts you to rename the file space.
- The client prompts you *only* one time per file space. If you specify `no` at the prompt, the client cannot rename the file spaces later. However, the Tivoli Storage Manager administrator can rename the file spaces on the server.
- When backing up files to a file space that is not Unicode-enabled, the Unicode-enabled client skips the files and directories with names containing characters from a code page that is different from the current locale.
- If files and directories with names containing characters from a code page other than the current locale were previously backed up with a client that was not Unicode-enabled, they might be expired. The Unicode-enabled client expires

these files if you do not migrate the file space to a Unicode-enabled file space. You can back up and archive these files to a Unicode-enabled file space.

Examples

Options file:
```
autofsrename yes
```

Related information

Automount

Authorized User: The automount option adds an automounted file system into the domain by mounting it. Use this option with the domain option.

Use this option to specify all automounted file systems the Tivoli Storage Manager client tries to mount at the following points in time:
- When Tivoli Storage Manager client starts
- When the backup is started
- When the Tivoli Storage Manager client has reached an automounted file system during backup

Mount the file system before Tivoli Storage Manager does a backup of that file system. If the file system is always mounted before the backup is done, it is unnecessary to explicitly specify an automounted file system in the automount option. However, add this file system in the automount option to ensure that the file system has been mounted at all the points in time mentioned previously. The automounted file systems are remounted if they have gone offline in the meantime during a backup.

Supported Clients

This option is valid for all UNIX platforms except Mac OS X. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client user options file (dsm.opt).

Syntax

```
►►──AUTOMount──┬──filespacename──┬──────────────────────────────►◄
               └─◄──────────────┘
```

Parameters

filespacename
 Specifies one or more fully qualified automounted file systems that are mounted and added into the domain.

Examples

Options file:
```
automount /home/Fred /home/Sam
```

Command line:
> Does not apply.

Related information

See "Domain" on page 303 for more information about working with automounted file systems and the domain option.

Backmc

The backmc option specifies the Tivoli Storage Manager management class to apply to the **backup fastback** command for retention purposes.

Use the backmc option with the **backup fastback** command.

If you back up an object more than once and specify a different management class for each backup, all backup versions of the object are rebound to the last management class specified.

Supported Clients

This option is valid for Linux x86/x86_64 clients.

Options File

None. You can specify this option only on the command line or on the scheduler.

Syntax

```
►►──BACKMc=──management_class_name──────────────────────────────────►◄
```

Parameters

management_class_name
    Specifies the management class name.

Examples

Command line:
```
dsmc backup fastback -fbpolicyname=policy1 -fbserver=server1
-backmc=ret2yrs
```

Backupsetname

The backupsetname option specifies the name of a backup set from the Tivoli Storage Manager server.

You must be an authorized user to use the backupsetname option.

You can use backupsetname option with the following commands:
- **query backup**
- **query filespace**
- **query image**
- **restore image**

Note: The following commands take **backupsetname** as a positional parameter.
The **backupsetname** positional parameter behaves differently from the
backupsetname option. See the command explanations for a discussion of how the
backupsetname positional parameter affects each of these commands:

> **query backupset**
>
> **restore**
>
> **restore backupset**

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager
client API does not support this option.

Options File

None. You can specify this option only on the command line.

Syntax

►►──BACKUPSETName──*backupsetname*──►◄

Parameters

backupsetname
> Specifies the name of a backup set from the Tivoli Storage Manager server. You
> cannot use wildcards.

Examples

Command line:

```
dsmc query backup /Volumes/bkSets/file.1
  -backupsetname=YEAR_END_ACCOUNTING.12345678

dsmc query backup /usr/projects -subdir=yes
  -backupsetname=YEAR_END_ACCOUNTING.12345678

dsmc restore image /home/proj
  -backupsetname=ACCOUNTING_2007.12345678

dsmc query image -backupsetname=WEEKLY_BSET.21435678
```

Related information

"Restore data from a backup set" on page 179

Changingretries

The changingretries option specifies how many additional times you want the
client to attempt to back up or archive a file that is in use. Use this option with the
archive, **incremental**, and **selective** commands.

You must be an authorized user to use the changingretries option.

This option is applied only when *serialization*, an attribute in a management class
copy group, is *shared static* or *shared dynamic*.

With *shared static* serialization, if a file is open during an operation, the operation repeats the number of times that you specify. If the file is open during each attempt, the operation does not complete.

With *shared dynamic* serialization, if a file is open during an operation, the operation repeats the number of times that you specify. The backup or archive occurs during the last attempt whether the file is open or not.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Backup** tab, **Number of retries if file is in use** field of the Preferences editor.

Syntax

```
►►──CHAngingretries─ numberretries────────────────────────────────►◄
```

Parameters

numberretries
>   Specifies the number of times a backup or archive operation is attempted if the file is in use. The range of values is zero through 4; the default is 4.

Examples

Options file:
>   changingretries 3

Command line:
>   -cha=3

Class

The class option specifies whether to display a list of NAS or client objects when using the **delete filespace**, **query backup**, and **query filespace** commands.

For example, to display a list of the file spaces belonging to a NAS node, enter the following command:
```
query filespace -class=nas
```

Supported Clients

This option is valid for AIX, Linux, and Solaris clients *only*. The Tivoli Storage Manager client API does not support this option.

Options File

None. You can specify this option only on the command line.

Syntax

```
►►──CLASS =──┬─client─┬──────────────────────────────────►◄
             └─nas────┘
```

Parameters

client

Specifies that you want to display a list of file spaces for a client node. This is the default.

nas

Specifies that you want to display a list of file spaces for a NAS node.

Examples

None. You can specify this option only on the command line.

Command line:
```
q backup -nasnodename=nodename -class=nas
```

Collocatebyfilespec

Use the collocatebyfilespec option to specify whether the Tivoli Storage Manager client uses only one server session to send objects generated from one file specification.

Setting the collocatebyfilespec option to *yes* attempts to eliminate interspersing of files from different file specifications, by limiting the client to one server session per file specification. Therefore, if you store the data to tape, files for each file specification are stored together on one tape (unless another tape is required for more capacity).

Considerations:
• Use the collocatebyfilespec option only if the storage pool is going directly to tape. If you use this option going to a disk storage pool, you could affect some load balancing, and therefore, performance.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option.

Options File

Place this option in the client user-options file (dsm.opt).

Syntax

```
►►──COLlocatebyfilespec──┬─No──┬──────────────────────────►◄
                         └─Yes─┘
```

Parameters

Yes
> Specifies that you want the Tivoli Storage Manager client to use only one server session to send objects generated from one file specification. Therefore, if you store the data to tape, files for each file specification are stored together on one tape, unless another tape is required for more capacity. Restore performance can increase as a result.

No Specifies that the Tivoli Storage Manager client can (depending on the execution dynamics and on the setting of the resourceutilization option of 3 or higher), use more than one server session to send the files from one file specification. This is the default.

> Backup performance might increase as a result. If the files are backed up to tape, files are stored on multiple tapes. Generally, the files specified in the file specification are still contiguous.

Examples

Options file:
> collocatebyfilespec yes

Command line:
> -collocatebyfilespec=yes

This option is valid only on the initial command line. It is not valid in interactive mode.

Commmethod

The commmethod option specifies the communication method you use to provide connectivity for client-server communication.

You must be an authorized user to use the commmethod option.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Communication** tab of the Preferences editor.

Syntax

```
                      ┌─TCPip─────┐
►►──COMMMethod────────┼───────────┼──────────────────────────────►◄
                      ├─SHAREdmem─┤
                      └─V6TCPIP───┘
```

Parameters

TCPip
> The Transmission Control Protocol/Internet Protocol (TCP/IP) communication method. This is the default.

V6Tcpip

Indicates that either TCP/IP Version 4 or Version 6 should be used, depending on the system configuration and the results of a domain name service lookup. A valid DNS environment must be available.

SHAREdmem

Use the shared memory communication method when the client and server are running on the same system. This provides better performance than the TCP/IP protocol. This communication method is valid for AIX, HP-UX, Linux, and Solaris clients *only*.

When specifying this communication method on AIX, the client can be logged in as root or non-root, as long as the server is running as root. If the server is not running as root, the user ID running the client must match the user ID running the server.

Important: When using commmethod *sharedmem* on Linux, you might receive error message: `ANR8294W Shared Memory Session unable to initialize` on the server or storage agent console. By default, Linux is not set up with sufficient system resources to create the message queues. You must increase the kernel parameter, `MSGMNI`, to 128 (the default is 16). You can modify this parameter by performing the following command:

`echo 128 > /proc/sys/kernel/msgmni`

To enable this parameter to remain persistent after rebooting the system, you can instead add the following line to the file `/etc/sysctl.conf`, then reboot the system:

`kernel.msgmni=128`

To view the current ipc settings, run this command:

`ipcs -l`

Now look at the `max queues system wide` value. The default is 16.

Examples

Options file:

    `commmethod tcpip`

    Use only TCP/IP Version 4

    `commmethod V6Tcpip`

    Use both TCP/IP Version 4 and Version 6, depending on how the system is configured, and the results of a domain name service lookup.

    Note: The dsmc schedule command cannot be used when both SCHEDMODe *prompt* and commmethod *V6Tcpip* are specified.

Command line:

    `-commm=tcpip`

    `-commm=V6Tcpip`

This option is valid only on the initial command line. It is not valid in interactive mode.

Commrestartduration

The commrestartduration option specifies the maximum number of minutes you want the client to try to reconnect to a Tivoli Storage Manager server after a communication error occurs.

You must be an authorized user to use the commrestartduration option.

Note: A scheduled event continues if the client reconnects with the server before the commrestartduration value elapses, even if the startup window of the event has elapsed.

You can use the commrestartduration option and the commrestartinterval in busy or unstable network environments to decrease connection failures.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Communication** tab, **Common Options** section of the Preferences editor.

Syntax

```
►►──COMMRESTARTDuration── minutes────────────────────────────────►◄
```

Parameters

minutes
   The maximum number of minutes you want the client to attempt to reconnect with a server after a communication failure occurs. The range of values is zero through 9999; the default is 60.

Examples

Options file:
   ```
   commrestartduration 90
   ```

Command line:
   Does not apply.

Commrestartinterval

The commrestartinterval option specifies the number of seconds you want the client to wait between attempts to reconnect to a Tivoli Storage Manager server after a communication error occurs.

You must be an authorized user to use the commrestartinterval option.

Note: Use this option only when commrestartduration is a value greater than zero.

You can use the commrestartduration option and the commrestartinterval in busy or unstable network environments to decrease connection failures.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Communication** tab, **Common Options** section of the Preferences editor.

Syntax

►►──COMMRESTARTInterval─ *seconds*──►◄

Parameters

seconds
> The number of seconds you want the client to wait between attempts to reconnect with a server after a communication failure occurs. The range of values is zero through 65535; the default is 15.

Examples

Options file:
> commrestartinterval 30

Command line:
> Does not apply.

Compressalways

The compressalways option specifies whether to continue compressing an object if it grows during compression.

Use this option with the compression option, and with the **archive**, **incremental**, and **selective** commands.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Backup** tab, **Continue Compressing if Object Grows** check box of the Preferences editor.

Syntax

```
                       ┌─Yes─┐
►►──COMPRESSAlways──────┼─────┼──────────────────────────────────────────────►◄
                       └─No──┘
```

Parameters

Yes
File compression continues even if the file grows as a result of compression. This is the default.

No   Backup-archive client objects are resent uncompressed if they grow during compression. API behavior depends on the application. Application backups might fail.

Examples

Options file:
        compressalways yes

Command line:
        -compressa=no

This option is valid only on the initial command line. It is not valid in interactive mode.

Compression

The compression option compresses files *before* you send them to the server.

Compressing your files reduces data storage for backup versions and archive copies of your files. It can, however, affect Tivoli Storage Manager throughput. A fast processor on a slow network connection benefits from compression, but a slow processor on a fast network connection does not.

You must be an authorized user to use the compression option.

Use the compression option with the **archive**, **incremental**, and **selective** commands.

The **backup image** command uses the compression option value specified in the dsm.sys file. This option is valid on the initial command line and in interactive mode. The server can also define this option which overrides the client value.

Tivoli Storage Manager backs up a sparse file as a regular file if client compression is off. Set the compression option to *yes* to enable file compression when backing up sparse files to minimize network transaction time and maximize server storage space.

If you set the compressalways option to *yes*, compression continues even if the file size increases. To stop compression if the file size grows, and resend the file uncompressed, set the compressalways option to *no*.

If you set the compression option to *yes*, you can control compression processing in the following ways:

- Use the exclude.compression option in your client system-options file (dsm.sys) to exclude specific files or groups of files from compression processing.
- Use the include.compression option in your client system-options file (dsm.sys) to include files within a broad group of excluded files for compression processing.

This option controls compression *only if* your administrator specifies that your client node can compress files before sending them to the server.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Backup** tab, **Compress objects** check box of the Preferences editor.

Syntax

```
►►──COMPRESSIon──┬──No──┬──────────────────────────────────►◄
                 └──Yes─┘
```

Parameters

No Files are not compressed before they are sent to the server. This is the default.

Yes
    Files are compressed before they are sent to the server.

Examples

Options file:
        compression yes

Command line:
        -compressi=no

This option is valid only on the initial command line. It is not valid in interactive mode.

Related information

Console

Use the console option with the **query systeminfo** command to output information to the console.
- DSMOPTFILE - The contents of the dsm.opt file.
- DSMSYSFILE - The contents of the dsm.sys file.
- ENV - Environment variables.
- ERRORLOG - The Tivoli Storage Manager error log file.
- FILE - Attributes for the file name that you specify.
- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- OPTIONS - Compiled options.

- OSINFO - Name and version of the client operating system (includes ULIMIT information for UNIX and Linux).
- POLICY - Policy set dump.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually dsmsched.log).
- CLUSTER - AIX cluster information.

Note: The **query systeminfo** command is intended primarily as an aid for IBM support to assist in diagnosing problems, although users who are familiar with the concepts addressed by this information might also find it useful. If you use the console option, no special formatting of the output is performed to accommodate screen height or width. Therefore, the console output might be difficult to read due to length and line-wrapping. In this case, use the filename option with the **query systeminfo** command to allow the output to be written to a file that can subsequently be submitted to IBM support.

Supported Clients

This option is valid for all clients.

Syntax

```
►►──CONsole──────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Command line:
```
query systeminfo dsmoptfile errorlog -console
```

Related information

"Filename" on page 335

Createnewbase

The createnewbase option creates a base snapshot and uses it as a source to run a full incremental backup.

Some files might not be backed up when the snapshot difference incremental backup command is run. If the files are skipped, you can run a snapshot difference incremental backup with the createnewbase option to backup these files. See "Snapdiff" on page 418 for a list of reasons why a file might not be backed up when the snapshot difference command is run.

NetApp Data ONTAP versions 8.0 and versions lower than 7.3.3 do not support file names that are not within the 7 bit ASCII character set. If you upgraded to NetApp Data ONTAP version 7.3.3, that supports file names with non 7 bit ASCII characters, use the createnewbase=migrate option to create a new base snapshot and run a full incremental backup. The createnewbase=migrate command backs up any files with non 7 bit ASCII names, that might have been skipped previously.

Supported Clients

This option is valid for the following clients:

- AIX 64 bit clients
- Linux x86/86_64 clients

Enter the createnewbase option on the command line. Specify this option with the snapdiff option.

Syntax

```
►►──Createnewbase──┬──No──┬──────────────────────────────────────►◄
                   │      │
                   ├─Yes──┤
                   ├─IGNore─┤
                   └─MIGRate─┘
```
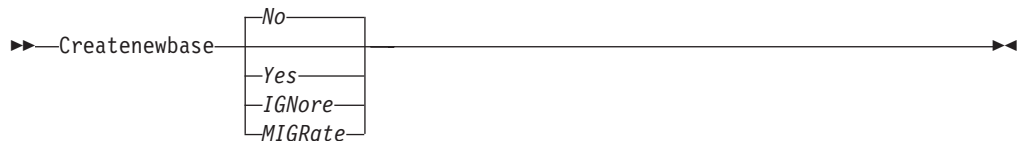
Parameters

NO
> Specifies that a snapshot difference incremental is run. If Tivoli Storage Manager detects that the NetApp Data ONTAP file server has been migrated from a version that does not support unicode-enabled file names to a file server that does, a warning message is recorded to the error log and the Tivoli Storage Manager server activity log. The warning message indicates that you must run a full incremental backup and logs a return code of 8 even if the operation completed successfully.
>
> This parameter is the default value.

YES
> Specifies that a full incremental is run by creating a new base snapshot and using it to run a scan-based incremental backup. Use this option to backup any file changes that might not have been detected by the snapshot difference API.
>
> If the operation finished successfully, the command ends with a return code of 0.
>
> It is recommended that you do not set **createnewbase=yes** for any schedule that runs a daily snapshot difference backup. Instead, create a separate, monthly schedule that has the **createnewbase=yes** option.

IGNore
> Specifies that a snapshot difference incremental backup is run when Tivoli Storage Manager detects that the NetApp Data ONTAP file server was upgraded to support Unicode file names.
>
> The **IGNore** parameter is different from the **No** parameter because Ignore suppresses the warning message. Instead, an informational message is recorded in the error log and the Tivoli Storage Manager activity log that tells you to run a full incremental backup.
>
> If the command finishes successfully, it returns a code of 0.
>
> Use the **IGNore** parameter if you have upgraded the NetApp Data ONTAP file server to support Unicode but you have not yet run a full incremental backup. This option is used only when the Tivoli Storage Manager client has detected that the file server was migrated and a full incremental has not yet been run. The option is ignored for all other times.

MIGRate

Specifies that if the NetApp Data ONTAP file server was upgraded to a version that supports Unicode file names, a base snapshot is taken and a scan-based incremental backup runs. The **MIGRate** parameter is different from the **YES** parameter because the **MIGRate** parameter creates a base snapshot only when Tivoli Storage Manager detects that the NetApp Data ONTAP file server version was updated. The **YES** parameter creates a base snapshot every time the command is run.

After the incremental backup finishes, no additional migration-related messages are recorded to the error log or the Tivoli Storage Manager server activity log. When the operation finishes, the command ends with a return code of 0.

Use the **MIGRate** parameter if you have upgraded the NetApp Data ONTAP file server to support Unicode but you have not yet run a full incremental backup. The **MIGRate** parameter is ignored if the NetApp Data ONTAP file server has not been upgraded.

Examples

Command line:

```
dsmc incremental -snapdiff -createnewbase=yes /net/home1
```

Related information

"Snapdiff" on page 418

"Configuring the file server connection for Tivoli Storage Manager" on page 91

Dateformat

The dateformat option specifies the format you want to use to display dates.

The AIX, Solaris, and HP-UX clients support locales other than English that describe every user interface that varies with location or language.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time you start the client. Consult the documentation on your local system for details about setting up your locale definition.

Note:

1. The dateformat option does not affect the Web client. The Web client uses the date format for the locale that the browser is running in. If the browser is not running in a locale that Tivoli Storage Manager supports, the Web client uses the date format for US English.
2. When you change the date format and use the schedlogretention option to prune the schedule log, Tivoli Storage Manager removes all entries in the schedule log with a different date format when pruning the log. When you change the date format and use the errorlogretention option to prune the error log, Tivoli Storage Manager removes all entries in the error log with a different date when pruning the log. When changing the date format, copy the schedule log and error log if you want to preserve log entries that contain a different date format.

You can use the dateformat option with the following commands.

- **delete archive**
- **delete backup**
- **expire**
- **query archive**
- **query backup**
- **query filespace**
- **query image**
- **restore**
- **restore image**
- **restore nas**
- **retrieve**
- **set event**

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Regional Settings** tab, **Date Format** drop-down list of the Preferences editor.

Syntax

►►──DATEformat── ──*format_number*───────────────────────────────────►◄

Parameters

format_number

Displays the date using one of the following formats. Select the number that corresponds to the date format you want to use:

0   Use the locale-specified date format (does not apply to Mac OS X).

For AIX, HP-UX, and Solaris: This is the default if the locale-specified date format consists of digits and separator characters.

1   MM/DD/YYYY

For AIX, HP-UX, and Solaris: This is the default if the locale-specified date format consists of anything but digits and separator characters.

This is the default for the following available translations:
- US English
- Chinese (Traditional)
- Korean

2   DD-MM-YYYY

This is the default for the following available translations:
- Brazilian Portuguese
- Italian

3   YYYY-MM-DD

This is the default for the following available translations:
- Japanese
- Chinese (Simplified)
- Polish

4   DD.MM.YYYY

This is the default for the following available translations:

- German
- French
- Spanish
- Czech
- Russian

5 YYYY.MM.DD

This is the default for the following available translations:
- Hungarian

6 YYYY/MM/DD

7 DD/MM/YYYY

For AIX, HP-UX, and Solaris: To set a particular date format, edit the source file for your locale and modify the **d_fmt** line to support your needs. Whatever date format you select applies both to output and to input; however, the input year can be either 2 or 4 digits.

"%m/%d/%y"
    Displays the date in the form MM/DD/YY

"%d.%m.%Y"
    Displays the date in the form DD.MM.YYYY

When you include the dateformat option with a command, it must precede the fromdate, pitdate, and todate options.

Examples

Options file:
```
dateformat 3
```

Command line:
```
-date=3
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

Dedupcachepath

Use the dedupcachepath option to specify the location where the client-side data deduplication cache database is created.

This option is ignored if the enablededupcache=no option is set during backup or archive processing.

Supported Clients

This option is valid for all clients. This option is also valid for the Tivoli Storage Manager client API.

Options File

Place this option in the system-options file (dsm.sys). You can set this option on the **Deduplication → Deduplication Cache Location** field of the Preferences editor. The option can be set in the client option set on the Tivoli Storage Manager server.

Syntax

►►──DEDUPCACHEPath──*path*──◄◄

Parameters

path

> Specifies the location where the client-side data deduplication cache database is created if the enablededupcache option is set to *yes*. The default location is to create the data deduplication cache file in the Tivoli Storage Manager client or API installation directory.

Examples

Options file:
> dedupcachepath /volumes/temp

Command line:
> Does not apply.

Related reference

"Enablededupcache" on page 312

Dedupcachesize

Use the dedupcachesize option to determine the maximum size of the data deduplication cache file. When the cache file reaches its maximum size, the contents of the cache are deleted and new entries are added.

Supported Clients

This option is valid for all clients. This option is also valid for the Tivoli Storage Manager client API.

Options File

Place this option in the system-options file (dsm.sys). You can set this option on the **Deduplication → Deduplication Cache → Maximum Size** field of the Preferences editor. The option can be set in the client option set on the Tivoli Storage Manager server.

Syntax

►►──DEDUPCACHESize──*dedupcachesize*──◄◄

Parameters

dedupcachesize

> Specifies the maximum size, in megabytes, of the data deduplication cache file. The range of values is 1 - 2048; the default is 256.

Examples

Options file:
> dedupcachesize 1024

Command line:
> Does not apply.

Related reference

"Deduplication"

Deduplication

Use the deduplication option to specify whether to enable redundant client-side data elimination when data is transferred to the Tivoli Storage Manager server during backup and archive processing.

Data deduplication is disabled if the enablelanfree or subfilebackup options are set. Tivoli Storage Manager client encrypted files are excluded from client-side data deduplication. Files from encrypted file systems are also excluded.

To support client-side data deduplication, the following criteria must be met:
- Client-side data deduplication for the node is enabled on the Tivoli Storage Manager server.
- The storage pool destination for the data must be a storage pool that is enabled for data deduplication. The storage pool must have a device type of "file".
- A file can be excluded from client-side data deduplication processing (by default all files are included).
- The server can limit the maximum transaction size for data deduplication by setting the CLIENTDEDUPTXNLIMIT option on the server. Refer to the *Administrator's Guide* for details.
- The file size must be larger than 2 K.

Supported Clients

This option is valid for all clients. This option is also valid on the Tivoli Storage Manager client API.

Options File

Place this option in the system-options file (dsm.sys) within a server stanza. You can set this option by selecting the **Deduplication ▸ Enable Deduplication** check box of the Preferences editor. The option can be set in the client option set on the Tivoli Storage Manager server.

Syntax

```
►►──DEDUPLication──┬──No──┬─────────────────────────────────────────►◄
                   └─Yes─┘
```

Parameters

No Specifies that you do not want to enable client-side data deduplication for backup and archive processing. *No* is the default.

Yes
> Specifies that you want to enable client-side data deduplication for backup and archive processing.

Examples

Options file:
>   deduplication yes

Command line:
>   -deduplication=yes

This option is valid only on the initial command line. It is not valid in interactive mode.

Related reference

"Include options" on page 351

"Exclude options" on page 321

Defaultserver

Authorized user: Use the defaultserver option to specify the name of the Tivoli Storage Manager server to contact for backup-archive services if more than one server is defined in the dsm.sys file.

By default, Tivoli Storage Manager contacts the server defined by the first stanza in the dsm.sys file. This option is only used if the servername option is not specified in the client user-options file (dsm.opt).

If you have the HSM client installed on your workstation, and you do not specify a migration server with the migrateserver option, use this option to specify the server to which you want to migrate files. See *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide* for more information.

Supported Clients

This option is valid for all UNIX clients.

Options File

Place this option *at the beginning* of the dsm.sys file *before* any server stanzas.

Syntax

```
►►──DEFAULTServer── ──servername─────────────────────────────────────────►◄
```

Parameters

servername
> Specifies the name of the default server to which you back up or archive files. The server to which files are migrated from your local file systems can also be specified with this option.

Examples

Options file:
>   defaults server_a

Command line:
>   Does not apply.

Deletefiles

Use the deletefiles option with the **archive** command to delete files from your workstation after you archive them.

You can also use this option with the **restore image** command and the incremental option to delete files from the restored image if they were deleted after the image was created. Deletion of files is performed correctly if the backup copy group of the Tivoli Storage Manager server has enough versions for existing and deleted files.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──DELetefiles─────────────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc archive "/Users/dgordon/Documents/*.c" –deletefiles

dsmc archive "/home/foo/*.c" –deletefiles
dsmc restore image /local/data -incremental -deletefiles
```

Description

The description option assigns or specifies a description for files when performing archive, delete archive, retrieve, query archive, or query backupset.

For example, if you want to archive a file named budget.jan and assign to it the description **2002 Budget for Proj 1**, you would enter:

```
dsmc archive –des="2003 Budget for Proj 1" /home/plan/
proj1/budget.jan
```

Note:
1. The maximum length of a description is 254 characters.
2. Enclose the value in quotation marks (" ") if the option value that you enter contains a blank space.

Use the description option with the following commands:
* **archive**
* **delete archive**
* **query archive**
* **query backupset**
* **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──DEScription =── ──description────────────────────────────────────────────────►◄
```

Parameters

description

Assigns a description to the file you are archiving. If you do not specify a description with the **archive** command, the default is `Archive Date:x`, where x is the current system date. Note that the date is always 10 characters long. If your date format uses a two digit year, there are two blank spaces at the end of the date. For example, a default description using a four-digit year might be `"Archive Date: 2002/05/03"`, and the same default with a two-digit year might be `"Archive Date: 02/05/03 "` (note the two spaces at the end). When retrieving files using the two-digit year description, you can enter the -description option string in either of the following ways:

```
-description="ArchiveDate: 02/05/03  "
 or
-description="ArchiveDate: 02/05/03*"
```

If you use the **archive** command to archive more than one file, the description you enter applies to each file. For example, to archive a group of files and assign the same description, *Project X*, to each file, you would enter:

```
dsmc archive –description="Project X" "/Users/van/Documents/*.x"
   dsmc archive –description="Project X" "/home/allproj/.*"
```

You can then use the description to retrieve all of the files.

Examples

Command line:

```
        dsmc archive "/Users/van/Documents/*.prj" -des="2003 Budget for Proj 1"

        dsmc archive "/home/foo/*.prj" –des="2003 Budget for Proj 1"
        dsmc query backupset –loc=server –descr="My Laptop"
```

Detail

Use the detail option to display management class, file space, backup, and archive information depending on the command with which it is used.

Use the detail option with the **query mgmtclass** command to display detailed information about each management class in your active policy set. If you do not use the detail option, only the management class name and a brief description are displayed on the screen. If you specify the detail option, information about attributes in each copy group contained in each management class is displayed on the screen. A management class can contain a backup copy group, an archive copy group, both, or neither.

A Unicode-enabled file space might not display correctly if the server cannot display the Unicode name. In this case, use the file space identifier (fsID) of the file

space to identify these file spaces on the server. Use the detail option with the **delete filespace** and **query filespace** commands to determine the fsID of a file space. The fsID also appears in the file information dialog in the backup-archive client and Web client GUIs.

Use the detail option with the **query backup** and **query archive** commands to display these attributes of the file that you specify:

- Last modification date
- Last access date
- Compression
- Encryption type
- Client-side data deduplication

Use the detail option with the following commands:
- **delete filespace**
- **query archive**
- **query backup**
- **query filespace**
- **query inclexcl**
- **query mgmtclass**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──DETail──────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc query mgmtclass -detail
dsmc query filespace -detail
dsmc query backup file1 -detail
```

Diffsnapshot

The diffsnapshot option controls whether Tivoli Storage Manager should create the differential snapshot when running a snapshot difference incremental backup .

If the differential snapshot is not created by Tivoli Storage Manager, the latest snapshot found on the volume is used as the differential snapshot and as the source for the backup operation.

The default value is to create the differential snapshot. This option is ignored the first time the snapdiff option is used. The first time the snapdiff option is used on a volume, a snapshot must be created and used as the source for a full incremental

backup. Snapshots that are created by Tivoli Storage Manager are deleted by Tivoli Storage Manager after the next snapshot difference incremental backup is complete.

Snapshots can be created with the Network Appliance FilerView tool. Use the *latest* parameter if you want Tivoli Storage Manager to use the most recent snapshot that was created with this or any other method. Snapshots created by methods outside of Tivoli Storage Manager are never deleted by Tivoli Storage Manager.

Supported Clients

This option is valid for the following clients:
- AIX 64 bit clients
- Linux x86/86_64 clients

Syntax

```
►►──DIFFSNAPSHOT─────┬──create──┬────────────────────────────────────►◄
                     └──latest──┘
```

Parameters

create
  Specifies that you want to create a new, persistent, snapshot to use as the source snapshot. This is the default.

latest
  Specifies that you want to use the latest snapshot that is found on the file server as the source snapshot.

Examples

Command line:
  Perform a snapshot difference incremental backup of an NFS mounted file-system /vol/vol1 hosted on the file server homestore.example.com, where /net/home1 is the mount point of /vol/vol1.

  `incremental -snapdiff -diffsnapshot=latest /net/home1`
  The -diffsnapshot option value of *latest* means that the operation uses the latest snapshot (the active snapshot).

Related information

"Snapdiff" on page 418

"Createnewbase" on page 287

"Configuring the file server connection for Tivoli Storage Manager" on page 91

Dirmc

The dirmc option specifies the management class you want to use for directories.

If you do not specify this option to associate a management class with directories, the client program uses the management class in the active policy set of your policy domain with the longest retention period. Select a management class for individual directories that retains directories at least as long as it retains the files associated with them.

You must be an authorized user to use the dirmc option.

If you specify a management class with this option, all directories specified in a backup operation are bound to that management class.

The dirmc option specifies the management class of directories that you back up and it does not affect archived directories. Use the archmc option with the **archive** command to specify the available management class for your policy domain to which you want to bind your archived directories and files. If you do not use the archmc option, the server binds archived directories to the default management class. If the default management class has no archive copy group, the server binds archived directories to the management class with the shortest retention period.

Important: Only extended attributes and ACLs are stored in storage pools. The directory information, other than extended attributes and ACLs, remains in the database. On Windows systems, directories occupy storage pool space.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Backup** tab, **Directory Management Class** section in the Preferences editor.

Syntax

▶▶──DIRMc── ──*mgmtclassname*──▶◀

Parameters

mgmtclassname
    Specifies the name of the management class that you want to associate with directories. The client uses the management class name that you specify for all of the directories that you back up. If you do not specify this option, the client associates the management class with the longest retention period with directories.

Examples

Options file:
        dirm managdir

Command line
        Does not apply.

Related information

If you want to back up specific files to a management class see "Assign a management class to files" on page 235 for more information.

Dirsonly

The dirsonly option processes directories *only*. The client does not process files.

Use the dirsonly option with the following commands:
- **archive**
- **incremental**
- **query archive**
- **query backup**
- **restore**
- **restore backupset**
- **retrieve**
- **selective**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►──DIrsonly──►◄

Parameters

There are no parameters for this option.

Examples

Command line:
     dsmc query backup -dirsonly "/Users/*"

Command line:
     dsmc query backup -dirsonly "*"

Disablenqr

The disablenqr option specifies whether the Tivoli Storage Manager Backup-Archive Client can use the "no query restore" method for restoring files and directories from the server.

If you set the disablenqr option to *no* (the default), the client can use the "no query restore" process.

If you set the disablenqr option to *yes*, the client can use only the standard restore process (also known as "classic restore").

Note: There is no option or value to specify that the client can use only "no query restore" method.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the dsm.opt file.

Syntax

```
►►──DISABLENQR──┬──No──┬──────────────────────────────────►◄
                └─Yes──┘
```

Parameters

No Specifies that Tivoli Storage Manager can use the "no query restore" method. This is the default.

Yes
    Specifies that the client uses only the "standard restore" method. The "no query restore" method is not allowed.

Examples

Options file:
```
        disablenqr yes
```

Command line
```
        -disablenqr=yes
```

Diskbuffsize

The diskbuffsize option specifies the maximum disk I/O buffer size (in kilobytes) that the client can use when reading files. The diskbuffsize option replaces the largecommbuffers option.

You must be an authorized user to use the diskbuffsize option.

Optimal backup, archive, or HSM migration client performance can usually be achieved if the value for this option is equal to or smaller than the amount of file read ahead provided by the client file system. A larger buffer requires more memory and it might not improve performance.

Important: Use the default setting, unless otherwise directed by IBM support personnel.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza.

Syntax

```
►►──DISKBuffsize── ──size────────────────────────────────────────────►◄
```

Parameters

size

> Specifies the maximum disk I/O buffer size (in kilobytes) that the client uses when reading files. The range of values is 16 through 1023; the default is 32. For AIX: If the enablelanfree option is set to *no*, the default setting for diskbuffsize is 256.

Examples

Options file:
```
diskbuffsize 64
```

Command line:
> Does not apply.

Diskcachelocation

The diskcachelocation option specifies the location where the disk cache database is created if the option memoryefficientbackup=*diskcachemethod* is set during an incremental backup.

You can specify the diskcachelocation option in your option file, or with the include.fs option. If the diskcachelocation option appears in the option file, its value is used for all file systems not represented by an include.fs option containing the diskcachelocation option.

You must be an authorized user to use the diskcachelocation option.

The disk cache is a temporary file which is deleted after the **incremental** command is run. Use this option to select one of the following:

1. A location that has more free disk space if, when you are using memoryefficientbackup=*diskcachemethod*, you get the message that the disk cache file cannot be created because you do not have enough disk space.

2. A location on a different physical volume to reduce contention for the disk access mechanism, and therefore improve performance.

Important: For performance reasons, do not use a remote drive for diskcachelocation.

The actual amount of disk space required for the disk cache file created by disk cache incremental backups depends on the number of files and directories included in the backup and on the average length of the files and directories to be backed up. For UNIX and Linux, estimate 1 byte per character in the path name. For Mac OS X, estimate 4 bytes per character in the path name. For example, if there are 1 000 000 files and directories to be backed up and the average path length is 200 characters, then the database occupies approximately 200 MB for UNIX and Linux, and 800 MB for Mac OS X clients. Another way to estimate for planning purposes is to multiply the number of files and directories by the length of the longest path to establish a maximum database size.

A second disk cache file is created for the list of migrated files when backing up an HSM managed file system. The combined disk cache files, created by disk cache incremental backups and HSM managed file system backups, can require above 400 MB of disk space for each million files being backed up. The disk cache file can become very large. Large file support must be enabled on the file system that is being used for the disk cache file.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

Place this option in the dsm.sys file within a server stanza.

Syntax

►►──DISKCACHELocation── ──*path*──◄◄

Parameters

path
> Specifies the location where the disk cache database is created if memoryefficientbackup is set to *diskcachemethod*. The default location is to create the disk cache file in the root of the file space being processed.

Examples

Options file:

```
diskcachelocation /home
diskcachelocation /Volumes/hfs2
```

Command line:
> Does not apply.

Domain

The domain option specifies what you want to include for incremental backup.

Domain objects are backed up only if you invoke the **incremental** command without a file specification.

Tivoli Storage Manager uses the domain value in the following situations to determine which file systems to process during an incremental backup:
- When you run an incremental backup using the **incremental** command and you do not specify which file systems to process
- When your Tivoli Storage Manager administrator defines a schedule to run an incremental backup for you, but does not specify which file systems to process
- When you select the **Backup Domain** action from the Tivoli Storage Manager backup-archive client GUI

There are several places where the domain option can be defined:

- In an options file
- On the command line, when entered with a Tivoli Storage Manager command
- In a client option set, defined on the Tivoli Storage Manager server with the **define clientopt** command
- As an option on a scheduled command, defined on the Tivoli Storage Manager server with the **define schedule** command

If any of these sources contain a domain definition, the client backs up that domain. If more than one source specifies a domain, the client backs up all specified domains. The same domain object can be defined more than once, but the effect is the same as defining it only once. If you do not specify a domain, Tivoli Storage Manager backs up the default domain, as described in the **all-local** parameter.

You can exclude objects from the domain by specifying the exclusion operator (-) before the object. If any domain definition excludes an object, that object is excluded from the domain, even if another definition includes the object. You cannot use the domain exclusion operator (-) in front of any domain keyword that begins with all-.

If you invoke the incremental command with a file specification, Tivoli Storage Manager ignores any domain definitions and backs up only the file specification.

You can include a virtual mount point in your client domain.

Important: If you are running GPFS for AIX or GPFS for Linux x86/x86_64 in a multinode cluster, and all nodes share a mounted GPFS file system, Tivoli Storage Manager processes this file system as a local file system. Tivoli Storage Manager backs up the file system on each node during an incremental backup. To avoid this situation, you can do *one* of the following tasks:

- Explicitly configure the *domain* statement in the client user options file (dsm.opt) to list the file systems you want that node to back up.
- Set the exclude.fs option in the client system-options file to exclude the GPFS file system from backup services.

Automounted file systems

When performing a backup with the domain option set to *all-local*, files handled by automounter and loopback file systems are not backed up.

If you back up a file system with the domain option set to *all-local*, any subdirectories that are mount points for an automounted file system (autofs) are excluded from backup. Any files that exist on the server for the automounted subdirectory are expired.

When performing a backup with the domain option set to *all-lofs*, all explicit loopback file systems (LOFS) are backed up and all automounted file systems are excluded. For loop devices and local file systems handled by automounter, set the domain option to *all-auto-lofs*.

Use the automount option with the domain parameters, all-auto-nfs and all-auto-lofs, to specify one or more automounted file systems to be mounted and added into the domain. If you specify the automount option, automounted file systems are remounted if they have gone offline during the execution of the **incremental** command.

Virtual mount points cannot be used with automounted file systems.

For HP-UX, the domain option is enhanced with the keywords all-auto-lofs and all-auto-nfs to support automounted file systems. To use this enhancement, you must use the automounter, AutoFS. To activate AutoFS, set the autofs parameter to 1 in the /etc/rc.config.d/nfsconf file. Changing this parameter requires a reboot. If you encounter problems with NFS automounted file systems, install patches PHCO_24777 and PHNE_26388 (or later patches). For further information, refer to the HP-UX documentation.

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the options file, dsm.opt or dsm.sys. In the dsm.sys file, you must place this option within a server stanza. You can set this option on the **Backup** tab, **Domain for Backup** section of the Preferences editor.

Syntax

```
>>──DOMain──┬──all-local──────┬──────────────────────────><
            │                 │
            ├──domain─────────┤
            ├── -domain───────┤
            ├──all-lofs───────┤
            ├──all-nfs────────┤
            ├──all-auto-nfs───┤
            └──all-auto-lofs──┘
```
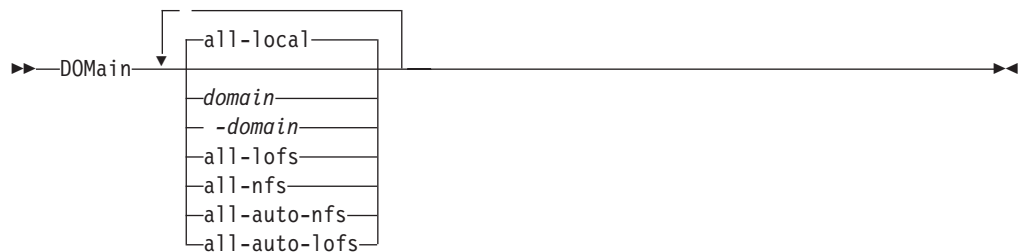
Parameters

all-local
> Backs up all local file systems except LOFS file systems and LOFS through automounter. This parameter is the default. The /tmp directory is not included.

domain
> Defines the file systems to include in your default client domain.
>
> When you use *domain* with the **incremental** command, it processes these file systems in addition to those file systems you specify in your default client domain.

-domain
> Defines the file systems to exclude in your default client domain.

all-lofs
> Backs up all loopback file systems, except those file systems handled by automounter. This parameter is not supported on Mac OS X.

all-nfs
> Backs up all network file systems, except those file systems handled by automounter. This parameter is not supported on Mac OS X.

all-auto-nfs
Backs up all network file systems (but not local file systems) which are handled by automounter. This parameter is not supported on Mac OS X.

all-auto-lofs
Backs up all loop devices and local file systems which are handled through automounter. This parameter is not supported on Mac OS X.

object
Specifies the domain objects to include in the domain.

An object name must be enclosed in quotation marks if the name includes any spaces.

-object
Specifies the domain objects to exclude from the domain.

An object name must be enclosed in quotation marks if the name includes any spaces.

Examples

Options file:

An options file can contain more than one domain statement. However, each of the domain statements are an example of a single statement in an options file.

```
domain all-local
domain all-local -/Volumes/volume2
domain all-local '-/Volumes/Macintosh HD'

domain /tst /datasave /joe
"domain all-local"
domain ALL-LOCAL -/home
domain ALL-NFS -/mount/nfs1
```

A single domain statement can list one or more objects for the domain. You can use more than one domain statement. The following two examples from two options files yield the same domain result:

Example 1
```
...
domain fs1
domain all-local
domain -fs3
...
```

Example 2
```
...
domain all-local fs1 -fs3
...
```

Command line:
```
-domain="/ /Volumes/volume2"
-domain="all-local -/Volumes/volume2"

-domain="/fs1 /fs2"
-domain=/tmp
-domain="ALL-LOCAL -/home"
```

Domain definition interaction

Domain can be defined in several sources, and the result is a summation of all domain definitions. As an example of the interaction of domain definitions, consider how domain definitions from several sources yield different backup

results. In the table, FS*n* is a file system. This table shows only commands that are entered on the command line. For scheduled commands, the command line column is not relevant, and options from the scheduled command must be considered.

Table 89. Interaction of domain definitions from several sources

Options file	Command line	Client option set	Objects backed up using the incremental command
domain FS1	incremental -domain=FS2	domain FS3	FS1 FS2 FS3
domain FS1	incremental	domain FS3	FS1 FS3
	incremental -domain=FS2		FS2
	incremental -domain=FS2	domain FS3	FS2 FS3
	incremental	domain FS3	FS3
	incremental		all-local
domain all-local	incremental	domain FS3	all-local + FS3
domain all-local domain -FS1	incremental		all-local, but not FS1
domain -FS1	incremental		none
domain FS1 FS3	incremental	domain -FS3	FS1
domain all-local	incremental	domain -FS3	all-local, but not FS3
	incremental FS1 -domain=all-local		FS1
	incremental FS1	domain all-local	FS1
domain -FS1	incremental FS1		FS1

Related information

For information about defining a virtual mount point, see "Virtualmountpoint" on page 452.

For information about specifying one or more automounted file systems to be mounted and added into the domain, see "Automount" on page 276.

Domain.image

The domain.image option specifies what you want to include in your client domain for an image backup.

If you do not specify a file system with the **backup image** command, the file systems you specify with the domain.image option are backed up.

When you specify a file system with the **backup image** command, the domain.image option is ignored.

If you do not use the domain.image option to specify file systems in your client options file, and you do not specify a file system with the **backup image** command, a message is issued and no backup occurs.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Backup** tab → **Domain for Backup** box of the Preferences editor.

Syntax

```
►►──DOMAIN.Image──┬──────────┬──────────────────────────────►◄
                  │ ◄──────┐ │
                  └─domain─┘
```

Parameters

domain
    Defines the file systems or raw logical volumes to include in your default client image domain.

Examples

Options file:
        domain.image /fs1 /fs2

Command line:
        Does not apply.

Domain.nas

The domain.nas option specifies the volumes to include in your NAS image backups.

You can specify *all-nas* to include all the mounted file systems on the NAS file server, except those you exclude with the exclude.fs.nas option.

Tivoli Storage Manager uses your domain for NAS image backups when you run a **backup nas** command and you do not specify which volumes to process.

When you use this option in your client system options file (dsm.sys), the domain.nas option defines your default domain for NAS image backups. When you perform a NAS file system image backup using the **backup nas** command, Tivoli Storage Manager adds the volumes that you specify on the command line to the volumes defined in your dsm.sys file. For example, if you enter domain.nas nas1/vol/vol0 nas1/vol/vol1 in your dsm.sys file and you enter dsmc backup nas

-nasnodename=nas1 /vol/vol2 on the command line, Tivoli Storage Manager backs up the `vol/vol0`, `vol/vol1`, and `vol/vol2` volumes on node `nas1`.

If you set the domain.nas option to *all-nas* in the dsm.opt file, Tivoli Storage Manager backs up all mounted volumes on the NAS file server. When performing a backup, if you use a file specification and set the domain.nas option to *all-nas* in the dsm.sys file, *all-nas* takes precedence.

Supported Clients

This option is valid for AIX and Solaris clients *only*. The server can also define this option.

Options File

Place this option in the dsm.sys file within a server stanza.

Syntax

```
►►──DOMAIN.Nas──┬──┬─all-nas─┬──┬────────────────────►◄
                │  └─domain──┘  │
                └──────◄────────┘
```

Parameters

domain
Defines the volumes you want to process. You cannot exclude volumes by specifying the dash (-) operator.

all-nas
Processes all mounted volumes on the NAS file server, except those you exclude with the exclude.fs.nas option. This is the default. If there is no domain.nas statement in the dsm.opt file and no volumes specified on the command line, Tivoli Storage Manager backs up all mounted volumes on the NAS server.

Examples

Options file:

```
domain.nas nas1/vol/vol0 nas1/vol/vol1
domain.nas all-nas
```

Command line:
Does not apply.

Dynamicimage

Use the dynamicimage option with the **backup image** command or the include.image option to specify that you want to perform a dynamic image backup.

Supported Clients

This option is valid for AIX, Solaris, HP-UX and all Linux clients. The Tivoli Storage Manager API does not support this option.

Options File

Place the include.image statement containing the dynamicimage value in the server stanza in your system-options file, dsm.sys. You can also set this option using the Preferences editor.

Syntax

►►──DYNAMICImage── ──*value*───►◄

Parameters

value
> Specifies one of the following values:

yes
> Use this option only if the volume cannot be unmounted and remounted as read-only. Tivoli Storage Manager backs up the volume as is without remounting it as read-only. Corruption of the backup can occur if applications write to the volume while the backup is in progress. In this case, run fsck after a restore and manually mount the file system in order to regain access to the volume. This option is valid for AIX, Solaris, HP-UX and all Linux clients.
>
> **Note:** This option is not allowed for AIX JFS2 file systems.

no Use this option if you do not want to perform a dynamic image backup. This is the default. The default behavior depends on the platform and file system type. For platforms and file systems that support snapshot, namely AIX JFS2 file systems and LINUX LVM file systems, the default is snapshot-based image backup. For all other UNIX platforms and file systems, the default is static image backup.

Examples

Options file:
> ```
> include.image /kalafs1 dynamicimage=yes
> ```

Command line on backup image:
> ```
> dynamicimage=yes
> ```

Efsdecrypt

The efsdecrypt option allows you to control whether or not files encrypted by an AIX Encrypted File System (EFS) are read in encrypted or decrypted format.

The efsdecrypt option default is *no*, which is to back up the encrypted or raw data. If you specify *yes*, the files are backed up as clear text, which means that they are backed up as normal files, as if the files existed in unencrypted form on the file system.

Important: Whenever you run a backup that includes any files encrypted on an EFS, you must ensure that you use the correct specification of the efsdecrypt option. If the efsdecrypt option value changes between two incremental backups, all encrypted files on EFS file systems are backed up again, even if they have not changed since the last backup. For example, if you are running an incremental backup of encrypted files that were previously backed up as "raw," then ensure

that efsdecrypt is specified as *no*. If you change efsdecrypt to *yes*, all the files are backed up again in clear text even if they are unchanged, so ensure that you use this option carefully.

Note: This is a global option that is applied to the complete backup. Two separate invocations of the client are required to back up some encrypted files as raw data and others as clear text.

Supported Clients

This option is valid for AIX clients.

Options File

Place this option in the dsm.sys file or the client user-options file (dsm.opt). In the dsm.sys file, you must place this option within a server stanza.

Syntax

```
►►──EFSDecrypt──┬──No──┬────────────────────────────────►◄
                └──Yes──┘
```

Parameters

No Encrypted files are read in encrypted or raw data format, and Tivoli Storage Manager encryption and compression is forced off. This is the default.

Yes
    Encrypted files are read in decrypted or clear text format.

Examples

Options file:
        EFSDecrypt yes

Command line:
        -EFSDecrypt=no

Enablearchiveretentionprotection

The enablearchiveretentionprotection option allows the client to connect to a Tivoli Storage Manager data retention server. This ensures that archive objects will not be deleted from the server until policy-based retention requirements for that object have been satisfied.

This option is ignored if the client connects to a server that is not retention protection enabled. If the option is *no* (the default) and an attempt is made to connect to a data retention server, the connection is refused.

The data retention server is specially configured for this task, so normal backup or restore processing is rejected by the server. When the client is connected to a data retention server, the following commands will not be available. If you attempt to use these commands, a message is displayed indicating that they are not valid with this server.
- **incremental**
- **backup** (all subcommands)

- **selective**
- **restore** (all subcommands except **restore backupset** -location=file or -location=tape)

  Note: restore backupset -location=file or -location=tape do not connect to any server (except the virtual one) and thus will not be blocked under any circumstances.
- **restart restore**
- **delete backup**
- **delete group**
- **expire**
- All queries *except:*
  – **query access**
  – **query archive**
  – **query filespace**
  – **query inclexcl**
  – **query managementclass**
  – **query node**
  – **query options**
  – **query schedule**
  – **query session**
  – **query systeminfo**
  – **query tracestatus**

Supported Clients

This option is valid for all clients.

Options File

This option is valid only in the dsm.sys file *within* a server stanza and is not valid in a client option set from the server. It is not valid on any command line.

Syntax

```
►►──ENABLEARCHIVERETENTIONProtection──┬─No──┬──────────────────────────────►◄
                                      └─Yes─┘
```
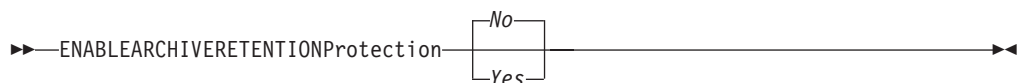
Parameters

No The data retention server connection is refused. This is the default.

Yes
    The client connects to a data retention server.

Enablededupcache

Use the enablededupcache option to specify whether you want to use a cache during client-side data deduplication. Using a local cache can reduce network traffic between the Tivoli Storage Manager server and the client.

When you perform a backup or archive operation with the data deduplication cache enabled, the specification of data extents that are backed up or archived are saved to the cache database. The next time you run a backup or archive, the client

queries the data deduplication cache and identifies the extents of data that have been previously saved to the server. Data extents that are identical to data extents on the server are not resent to the server.

If the server and the cache are not synchronized, the cache is removed and a new one is created.

Only one process can access the distributed data deduplication cache at a time. Concurrent backup instances on a workstation, that use the same server and storage pool, must either use unique node names or unique cache specifications. In this way, all the instances can use a local cache and optimize the client-side data deduplication.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API also supports this option.

Options File

Place this option in the system-options file (dsm.sys) within a server stanza. You can set this option on the **Deduplication → Enable Deduplication Cache** check box of the Preferences editor. The option can be set in the client option set on the Tivoli Storage Manager server.

Syntax

```
►►──ENABLEDEDUPCache──┬──Yes*──┬──────────────────────────►◄
                      └──No────┘
```

Parameters

Yes
    Specifies that you want to enable data deduplication cache. If data deduplication is not enabled, this setting is not valid. *Yes* is the default for the Tivoli Storage Manager client. *No* is the default for the Tivoli Storage Manager API.

No Specifies that you do not want to enable data deduplication cache.

Examples

Options file:
        enablededupcache no

Command line:
        -enablededupcache=no

This option is valid only on the initial command line. It is not valid in interactive mode.

Enablelanfree

The enablelanfree option specifies whether to enable an available LAN-free path to a storage area network (SAN) attached storage device.

A LAN-free path allows backup, restore, archive, and retrieve processing between the Tivoli Storage Manager client and the SAN-attached storage device.

You must be an authorized user to use the enablelanfree option.

To support LAN-free data movement you must install and configure the Tivoli Storage Manager Managed System for SAN Storage Agent on the client workstation. For more information, refer to the Tivoli Storage Manager Managed System for SAN Storage Agent publications.

Note:
1. If you place the enablelanfree option in the client option file (dsm.opt), but zero (0) bytes were transferred through the SAN during an operation, ensure that you bind the data to a LAN-free enabled management class.
2. To restore backup sets in a SAN environment, see "Restore Backupset" on page 529 for more information.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris clients.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **General** tab → **Enable Lanfree** check box of the Preferences editor.

Syntax

```
                          ┌─No──┐
►►──ENABLELanfree─────────┼─────┼──────────────────────────────────────►◄
                          └─Yes─┘
```

Parameters

Yes
> Specifies that you want to enable an available LAN-free path to a SAN-attached storage device.

No Specifies that you do not want to enable a LAN-free path to a SAN-attached storage device. This is the default.

Examples

Options file:
>   enablelanfree yes

Command line:
>   -enablelanfree=yes

This option is valid only on the initial command line. It is not valid in interactive mode.

Related information

To specify a communication protocol between the Tivoli Storage Manager client and Storage Agent, see "Lanfreecommmethod" on page 358.

Encryptiontype

The encryptiontype option allows you to choose the algorithm for data encryption.

The encryptiontype option allows you to use AES 128-bit data encryption, providing a stronger form of data encryption than DES 56-bit data encryption. The encryption type only affects backup and archive operations. The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received. During restore and retrieve operations, the encrypted data is decrypted with the proper encryption algorithm, regardless of the setting for this option.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can also set this option on the **Authorization** tab of the Preferences editor. The server can override this.

Syntax

```
                       ┌─AES128─┐
►►──ENCRYPTIONType──────┼────────┼───────────────────────────────►◄
                       └─DES56──┘
```
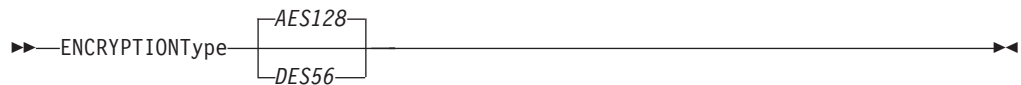
Parameters

AES128
>   AES 128-bit data encryption. AES 128-bit data encryption provides a stronger form of data encryption than DES 56-bit data encryption. This is the default.

DES56
>   DES 56-bit data encryption.

Examples

Options file:
>   encryptiontype des56

Command line:
     Does not apply.

Encryptkey

The Tivoli Storage Manager client supports the option to encrypt files being backed up or archived to the Tivoli Storage Manager server. This option is enabled with the include.encrypt option.

All files matching the pattern on the include.encrypt specification are encrypted before the data is sent to the server. There are three options for managing the key used to encrypt the files (prompt, save, and generate). All three options can be used with either the backup-archive client or the Tivoli Storage Manager API.

The encryption key password is case-sensitive and can be up to 63 characters in length and include the following characters:

A-Z   Any letter, A through Z, uppercase or lowercase. You cannot specify national language characters.
0-9   Any number, 0 through 9
+   Plus
.   Period
_   Underscore
-   Hyphen
&   Ampersand

Note:
1. The Tivoli Storage Manager API has an alternate way of specifying encryptkey=*generate*; the previous enableclientencryptkey=*yes* option can also be specified to request generate encryption processing.
2. The enableclientencryptkey=*yes* API option is still supported, so it is possible when using the API to specify two conflicting options. For example, enableclientencryptkey=*yes* and encryptkey=*prompt* or encryptkey=*save*.
3. When conflicting values are specified, the Tivoli Storage Manager API returns an error message.

Attention: When using the prompt option, your encryption key is not saved in the Tivoli Storage Manager password file on UNIX. If you forget the key, your data cannot be recovered.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Authorization** tab, **Encryption Key Password** section of the Preferences editor.

Syntax

```
>>─ENCRYPTKey──┬─save─────┬──────────────────────────────────────────><
               ├─prompt───┤
               └─generate─┘
```

Parameters

save

The encryption key password is saved in the Tivoli Storage Manager client's password file. A prompt is issued for an initial encryption key password, and after the initial prompt, the saved encryption key password in the password file is used for the backups and archives of files matching the include.encrypt specification. The password can be up to 63 bytes in length. The key is retrieved from the password file on restore and retrieve operations.

When the *save* option is specified for an API application (does not apply to Mac OS X), the initial key password must be provided by the application using the API in the *dsmInitEx* function call. The API itself does not issue a prompt to the user but relies on the application to prompt the user as necessary.

This is the default.

Restrictions:
- This option can only be used when passwordaccess *generate* is also specified.
- The root user or a Tivoli Storage Manager authorized user must specify the initial encryption key password.

prompt

The management of the encryption key password is provided by the user. The user is prompted for the encryption key password when the Tivoli Storage Manager client begins a backup or archive. A prompt for the same password is issued when restoring or retrieving the encrypted file. This password can be up to 63 bytes in length.

When the *prompt* option is specified for an API application (does not apply to Mac OS X), the key password must be provided by the application using the API in the *dsmInitEx* function call. The API itself does not issue a prompt to the user but relies on the application to prompt the user as necessary.

Restriction: This option can only be used by the root user or a Tivoli Storage Manager authorized user.

generate

An encryption key password is dynamically generated when the Tivoli Storage Manager client begins a backup or archive. This generated key password is used for the backups of files matching the include.encrypt specification. The generated key password, in an encrypted form, is kept on the Tivoli Storage Manager server. The key password is returned to the Tivoli Storage Manager client to enable the file to be decrypted on restore and retrieve operations.

Examples

Options file:

```
encryptkey prompt
```

Command line:

Does not apply.

Errorlogmax

The errorlogmax option specifies the maximum size of the error log, in megabytes.

You must be an authorized user to use the errorlogmax option.

If you change from errorlogmax to errorlogretention, all existing log entries are retained and the log is pruned using the new errorlogretention criteria.

If you change from errorlogretention to errorlogmax, all records in the existing log are copied to the pruned log dsmerlog.pru, the existing log is emptied, and logging begins under the new log wrapping criteria.

If you change the value of the errorlogmax option, the existing log is extended or shortened to accommodate the new size. If the value is reduced, the oldest entries are deleted to reduce the file to the new size.

Restriction: You cannot specify a non-zero errorlogmax value *and* enable errorlogretention.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **General** tab, **Select Error Log** button of the Preferences editor.

Syntax

```
►►──ERRORLOGMAX── ──size──────────────────────────────────────────────►◄
```

Parameters

size
> Specifies the maximum size, in megabytes, for the log file. The range of values is 0 to 2047; the default is 0, which disables log file wrapping and allows the log file to grow indefinitely.

Examples

Options file:
> errorlogmax 2000

Command line:
> -errorlogmax=2000

This option is valid only on the initial command line. It is not valid in interactive mode.

Errorlogname

This option specifies the fully qualified path and file name of the file that contains the error messages.

Authorized user: The value for this option overrides the `DSM_LOG` environment variable. The `dsmwebcl.log` and `dsmsched.log` files are created in the same directory as the error log file you specify with the errorlogname option.

For Mac OS X, the default location is one of the following:
```
~/Library/Logs/tivoli/tsm/
/Library/Logs/tivoli/tsm/
```

The dsmerror.log cannot be a symbolic link.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **General** tab, **Select Error Log** button of the Preferences editor.

Syntax

```
►►──ERRORLOGName─── ──filespec─────────────────────────────────────►◄
```

Parameters

filespec
> The fully qualified path and file name in which to store error log information. If any part of the path you specify does not exist, Tivoli Storage Manager attempts to create it.
>
> The dsmerror.log file *cannot* be a symbolic link.

Examples

Options file:
> ```
> errorlogname /tmp/tsmerror.log
> ```

Command line:
> ```
> -errorlogname=/tmp/tsmerror.log
> ```

This option is valid only on the initial command line. It is not valid in interactive mode.

Errorlogretention

The errorlogretention option specifies how many days to maintain error log entries before pruning, and whether to save the pruned entries.

The error log is pruned when the first error is written to the log after a Tivoli Storage Manager session is started. If the only session you run is the client scheduler, and you run it twenty-four hours a day, the error log might not be pruned according to your expectations. Stop the session and start it again to allow the scheduler to prune the error log.

You must be an authorized user to use the errorlogretention option.

If you change from errorlogretention to errorlogmax, all records in the existing log are copied to the pruned log dsmerlog.pru, the existing log is emptied, and logging begins under the new log wrapping criteria.

If you change from errorlogmax to errorlogretention, all existing log entries are retained and the log is pruned using the new errorlogretention criteria.

Restriction: You cannot specify errorlogretention *and* a non-zero errorlogmax value.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **General** tab, **Select Error Log** button of the Preferences editor.

Syntax

```
>>--ERRORLOGRetention--+--N----+--+--D--+------------------------><
                       |       |  |     |
                       +-days--+  +--S--+
```

Parameters

N **or** *days*
> Specifies how long to wait before pruning the error log.
>
> *N* Do not prune the error log. This permits the error log to grow indefinitely. This is the default.
>
> *days*
> > The number of days to keep log file entries before pruning the log. The range of values is zero through 9999.

D **or** *S*
> Specifies whether to save the pruned entries. Enter a space or comma to separate this parameter from the previous one.
>
> *D* Discard the error log entries when you prune the log. This is the default.
>
> *S* Save the error log entries when you prune the log.
>
> > The pruned entries are copied from the error log to the dsmerlog.pru file located in the same directory as the error log.

Examples

Options file:
> errorlogretention 400 S

Command line:
> -errorlogr=400,S

This option is valid only on the initial command line. It is not valid in interactive mode.

Determine compression and encryption processing

Tivoli Storage Manager evaluates exclude.dir and any other include-exclude options controlling backup and archive processing, and then the decision is made about which files undergo compression and encryption processing.

The following options determine which files undergo compression and encryption processing.

Table 90. Options for controlling compression and encryption

Option	Description	Page
Compression processing		
exclude.compression	Excludes files from compression processing if compression=*yes* is specified. This option applies to backups and archives.	"Exclude options"
include.compression	Includes files for compression processing if compression=*yes* is specified. This option applies to backups and archives.	"Include options" on page 351
Encryption processing		
exclude.encrypt	Excludes files from encryption processing.	"Exclude options"
include.encrypt	Includes files for encryption processing. The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received. **Important:** The include.encrypt option is the only way to enable encryption on the Backup-Archive client. If no include.encrypt statements are used encryption will not occur.	"Include options" on page 351

Exclude options

Use the exclude options to exclude objects from backup, image, or archive services.

Note: Image files do not apply to Mac OS X operating systems.

For example, you might want to exclude this type of information:
* All temporary files
* Any local caches of network files
* All files that contain compiled object code that you can easily reproduce using other methods
* Your operating system files

You can exclude specific files from encryption processing during a backup.

Note:
1. With the exception of exclude.fs, when you exclude a file that was previously included, existing backup versions become inactive during the next incremental backup.
2. The server can define exclude options with the inclexcl option.

Exclude any system files or images that could corrupt the operating system when recovered. Also exclude the directory containing the Tivoli Storage Manager client files.

Use wildcard characters to exclude a broad range of files.

To exclude an entire directory called /any/test, enter the following:

```
exclude.dir /any/test
```

To exclude subdirectories that begin with test under the /any directory, enter the following:

```
exclude.dir /any/test*
```

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set these options on the **Include-Exclude** tab, **Define Include-Exclude Options** section of the Preferences editor.

Syntax

▶▶──*options*─ ─*pattern*──◀◀

exclude, exclude.backup, exclude.file, exclude.file.backup
Use these options to exclude a file or group of files from backup services and space management services (if the HSM client is installed). The exclude.backup option only excludes files from normal backup, but not from HSM.

exclude.archive
Excludes a file or a group of files that match the pattern from archive services *only*.

exclude.attribute.symlink
Excludes a file or a group of files that are symbolic links or aliases (aliases apply to Mac OS X) from backup processing only.

Note: For Mac OS X aliases are excluded.

exclude.compression
Excludes files from compression processing if the compression option is set to *yes*. This option applies to backups and archives.

exclude.dedup
Excludes files from client-side data deduplication. To control a client-side data deduplication operation, specify ieobjtype as the value of the exclude.dedup option.

Valid ieobjtype parameters are:
    File
    Image

The default is File.

exclude.dir
Excludes a directory, its files, and all its subdirectories and their files from

backup processing. For example, the statement `exclude.dir /test/dan/data1` excludes the `/test/dan/data1` directory, its files, and all its subdirectories and their files.

If you exclude a directory that was previously included, Tivoli Storage Manager expires existing backup versions of the files and directories beneath it during the next incremental backup. Use this option to exclude a portion of your data that has no underlying files to back up.

Note: Avoid performing a selective backup of an individual file within an excluded directory. The next time that you perform an incremental backup, any files backed up in this manner is expired.

exclude.encrypt
Excludes the specified files from encryption processing. This option does not affect whether files are excluded from backup or archive processing, only whether they are excluded from encryption processing.

exclude.fs
Excludes file systems that match the specified pattern from backup, incremental image backup, and archive operations. If files from the excluded file systems were ever backed up, then management class rebinding and deleted file expiration does not occur. However, existing backup versions remain on the server subject to associated management class settings. The files that were previously archived from the excluded file system remain on the server as archive copies.

The exclude.fs option does NOT prevent the backup or archive of any virtual mount points that are subdirectories of the excluded file system.

Use exclude.image to exclude file systems from full image backup operations.

exclude.fs.nas
Excludes file systems on the NAS file server from an image backup when used with the **backup nas** command. The NAS node name must be prefixed to the file system name, for example: `netappsj1/vol/vol1`. To apply the exclude to all NAS nodes, replace the NAS node name with a wildcard, for example: `*/vol/vol1`. The **backup nas** command ignores all other exclude statements including exclude.fs and exclude.dir statements. This option is valid for AIX and Solaris clients *only*.

exclude.image

Restriction: This option does not apply to Mac OS X.
Excludes mounted file systems and raw logical volumes that match the specified pattern from full image backup operations. This option is valid for AIX, HP-UX, all Linux clients, and Solaris *only*. Use exclude.fs to exclude file systems from incremental image backup operations.

Parameters

pattern
Specifies the file or group of files that you want to exclude.

Note: For NAS file systems: You must prefix the NAS node name to the file specification to specify the file server to which the exclude statement applies. If you do not specify a NAS node name, the file system identified refers to the NAS node name specified in the client system-options file (dsm.sys) or on the command line.

If the pattern begins with a single or double quote or contains any embedded blanks or equal signs, you must surround the value in either single (') or double (") quotation marks. The opening and closing quotation marks must be the same type of quotation marks.

For the exclude.image option, the pattern is the name of a mounted file system or raw logical volume.

Examples

Options file:

```
exclude /unix/
exclude /.../core
exclude /home/jones/proj1/*
exclude.archive /.../core
exclude.backup /home/jones/proj1/devplan/
exclude.dir /home/jones/tmp
exclude.backup /users/home1/file1
exclude.image /usr/*/*
exclude.encrypt /users/home2/file1
exclude.compression /home/gordon/proj1/*
exclude.fs.nas netappsj/vol/vol0
exclude.attribute.symlink /.../*
exclude.dedup /Users/Administrator/Documents/Important/.../*
```

Command line:
Does not apply.

Related information

See "System files to exclude" on page 97 for a list of files that you should always exclude.

"Inclexcl" on page 349

See "Include and exclude groups of files with wildcard characters" on page 99 for a list of wildcard characters that you can use. Then, if necessary, use the include option to make exceptions.

Controlling symbolic link and alias processing

Tivoli Storage Manager treats symbolic links and aliases (aliases apply to Mac OS X only) as actual files and backs them up. However, the file referenced by the symbolic link is not backed up. In some cases symbolic links can be easily recreated and need not be backed up.

In addition, backing up these symbolic links can increase backup processing time and occupy a substantial amount of space on the Tivoli Storage Manager server. You can use the exclude.attribute.symlink option to exclude a file or a group of files that are symbolic links from backup processing. If necessary, you can use the include.attribute.symlink option to include symbolic links within broad group of excluded files for backup processing.

For example, to exclude all symbolic links from backup processing, except those that exist under the /home/spike directory, enter these statements in your dsm.sys file:

```
exclude.attribute.symlink  /.../*
include.attribute.symlink /home/spike/.../*
```

Controlling compression processing

This topic lists some items to consider if you want to exclude specific files or groups of files from compression processing during a backup or archive operation.

- Remember that Tivoli Storage Manager compares the files it processes against the patterns specified in the include-exclude statements, reading from the bottom to the top of the options file.

- You must set the compression option to *yes* to enable compression processing. If you do not specify the compression option or you set the compression option to *no*, Tivoli Storage Manager does not perform compression processing.

  If you set the compression option to *yes* and no exclude.compression statements exist, Tivoli Storage Manager considers all files for compression processing.

- Tivoli Storage Manager processes exclude.fs, exclude.dir, and other include-exclude statements first. Tivoli Storage Manager then considers any exclude.compression statements. For example, consider the following include-exclude list:

```
exclude /home/jones/proj1/*.*
exclude.compression /home/jones/proj1/file.txt
include /home/jones/proj1/file.txt
```

  Tivoli Storage Manager examines the statements (reading from bottom to top) and determines that the /home/jones/proj1/file.txt file is a candidate for backup, but is not a candidate for compression processing.

- Include-exclude compression processing is valid for backup and archive processing *only*. The exclude.compression option does not affect whether files are excluded from backup or archive processing, only whether they are excluded from compression processing.

Related reference
"Compression" on page 285

Processing NAS file systems

Use the exclude.fs.nas option to exclude file systems from NAS image backup processing.

Note: The exclude.fs.nas option does not apply to a snapshot difference incremental backup.

A NAS file system specification uses the following conventions:

- NAS nodes represent a unique node type. The NAS node name uniquely identifies a NAS file server and its data to Tivoli Storage Manager. You can prefix the NAS node name to the file specification to specify the file server to which the exclude statement applies. If you do not specify a NAS node name, the file system identified applies to all NAS file servers.

- Regardless of the client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.

For example, to exclude /vol/vol1 from backup services on all NAS nodes, specify the following exclude statement:

```
exclude.fs.nas */vol/vol1
```

Fbbranch

Use the fbbranch option with the **backup fastback** or **archive fastback** commands.

The fbbranch option specifies the branch ID of the remote FastBack server to back up or archive. The fbbranch option is only required when the a backup-archive client is installed on a FastBack DR Hub or when a dedicated proxy is connecting to a replicated DR Hub repository. Do not specify the fbbranch option when the backup-archive client is installed on the FastBack server.

Supported Clients

This option is valid for Linux x86/x86_64 clients.

Options File

None. You can specify this option only on the command line. The server can also define or override this option.

Syntax

```
►►──FBBranch=──branch_ID──────────────────────────────────────────◄◄
```

Parameters

branch_ID
> Specifies the FastBack server branch ID. The value is part of the disaster recovery configuration of the FastBack server.

Examples

Command line:
> -FBBranch=oracle
>
> On a Tivoli Storage Manager backup-archive client that is installed on a FastBack DR Hub:
>
> ```
> dsmc backup fastback -fbpolicyname=policy1 -fbserver=myFbServer
>   -fbbranch=oracle
> ```

Command line:
> On a Tivoli Storage Manager backup-archive client that is connecting to a repository on a remote FastBack DR Hub:
>
> ```
> dsmc backup fastback -fbpolicyname=policy1 -fbserver=server1
>   -Fbreposlocation=\\myDrHub.company.com\REP
>   -fbbranch=oracle
> ```

If the fbbranch option is specified on a Tivoli Storage Manager client workstation that is installed on a FastBack server, the fbbranch option is ignored.

Fbclientname

Use the fbclientname option with the **backup fastback** or **archive fastback** commands.

The fbclientname option is the name of one or more comma-separated FastBack clients to back up or archive from the backup proxy. The values for the fbclientname option are invalid if more than one policy is specified in the fbpolicyname option.

You cannot include spaces in the fbclientname option values.

If you do not specify any values for the fbvolumename option, all the volumes from all the FastBack clients in the policy that is specified are backed up. If you specify multiple FastBack clients in the fbclientname option, you cannot specify values for the fbvolumename option.

Supported Clients

This option is valid for Linux x86/x86_64 clients.

Options File

None. You can specify this option only on the command line.

Syntax

```
              ,
             ┌──────┐
             │  (1) │
►►─FBClientname─▼──client_name─┴──────────────────────────►◄
```
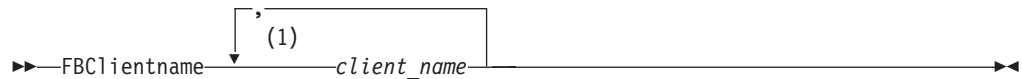
Notes:

1   Specify up to 10 client_name values for FBClientname.

Parameters

client_name
   Specifies the name of one or more FastBack clients. You can specify up to 10 FastBack client names.

Important:

When specifying the **archive fastback** or **backup fastback** command:
1. *At least* one FBpolicyName is always *required*.
2. You can specify up to 10 values for FBPolicyName, if no values are specified for both FBClientName *and* FBVolumeName.
3. When you specify a FBClientName value, there must be *only one* value for FBPolicyName.
4. You can specify up to 10 values for FBClientName if *only one* PolicyName is specified, and *no* values for FBVolumeName are specified.
5. When you specify the FBVolumeName option, you can have *only one* FBPolicy, and *only one* FBCLientName specified.
6. You can specify multiple FBVolumeNames if condition 5 is satisfied.
7. You must always specify the FBReposLocation option for Linux.

Examples

Command line:
```
dsmc backup fastback -fbpolicyname=Policy1
  -fbclientname=fbclient1,fbclient2
  -fbserver=myFbServer
  -fbreposlocation=/mnt/FBLocation
```

Backs up all volumes for FastBack clients fbclient1 and fbclient2 that are found in policy Policy1.

Fbpolicyname

Use the fbpolicyname option with the **backup fastback** or **archive fastback** commands.

The fbpolicyname option is the name of one or more comma-separated Tivoli Storage Manager FastBack policies that you want to back up or archive from the backup proxy. You must specify at least one policy name. Specify multiple policy names using a comma-delimited list of policies. There is no default value.

If one or more FB policy names contain spaces, you must specify them within quotation marks. Here is an example: `"FB Policy NAME1, FBPolicy Name 2"`.

If you do not specify any values for the fbclientname and fbvolumename options, all the volumes from all the FastBack clients in the policies that are specified are backed up. If you specify multiple policies in the fbpolicyname option, you cannot specify values for the fbclientname and fbvolumename options.

If a policy specification contains both Windows and Linux FastBack clients, only the Linux volumes will be backed up or archived to Tivoli Storage Manager by the Linux Tivoli Storage Manager backup-archive client.

At least one snapshot should exist in the Tivoli Storage Manager FastBack repository for the FastBack policies being archived or backed up prior to issuing the dsmc command

Supported Clients

This option is valid for Linux x86/x86_64 clients.

Options File

None. You can specify this option only on the command line.

Syntax

```
              ┌──,─────────────┐
              │       (1)      │
►►──FBPolicyname──▼──policy_name──┴────────────────────────►◄
```

Notes:

1 Specify up to 10 FastBack policy names.

Parameters

policy_name
  Specifies the name of the FastBack policies. You can specify up to 10 FastBack policy names.

Important:

When specifying the **archive fastback** or **backup fastback** command:
1. *At least* one FBpolicyName is always *required*.
2. You can specify up to 10 values for FBPolicyName, if no values are specified for both FBClientName *and* FBVolumeName.
3. When you specify a FBClientName value, there must be *only one* value for FBPolicyName.
4. You can specify up to 10 values for FBClientName if *only one* PolicyName is specified, and *no* values for FBVolumeName are specified.
5. When you specify the FBVolumeName option, you can have *only one* FBPolicy, and *only one* FBCLientName specified. You must specify exactly one FBClientName. It cannot be omitted.
6. You can specify multiple FBVolumeNames if condition 5 is satisfied.
7. You must always specify the FBReposLocation option for Linux.

Examples

Command line:
```
dsmc backup fastback -fbpolicyname=Policy1,Policy2,Policy3
  -fbserver=myFbServer
  -fbreposlocation=\\myFbServer.company.com\REP
```

Backs up all volumes for all FastBack clients found in policies Policy1, Policy2 and Policy3.

To specify policies with spaces, enclose them in double quotation marks, for example:
```
-fbpolicyname="Policy 1,Policy2,Policy3"
```

Fbreposlocation

Use the fbreposlocation option with the **backup fastback** or **archive fastback** commands.

The fbreposlocation option specifies the location of the Tivoli Storage Manager FastBack repository for the Tivoli Storage Manager backup-archive client proxy to connect to issue Tivoli Storage Manager FastBack shell commands necessary to mount appropriate snapshots.

This option is required on Linux systems. There is no default location.

It you specify the fbreposlocation option for a snapshot on the FastBack server, use the server_name@WORKGROUP format.

There are two ways to specify the FastBack repository location on the FastBack DR Hub.

- Specify the complete repository location via the option -fbreposlocation=\\ DR_Hub\rep_server. When using this format, "DR_Hub" is the FastBack DR Hub machine name and "rep_server" is the name of the replicated FastBack server repository on the DR Hub.
- Specify the repository location using a combination of the -fbreposlocation and -fbbranch options. When using this format, specify the DR Hub repository the location via the option -fbreposlocation=DR_Hub@WORKGROUP, and specify the name of the replicated FastBack server repository on the DR Hub using the -fbranch option.

If you use the format `-fbr=\\<fbserver>\REP`, specify two backslashes before `<fbserver>` and one backslash before REP when using the backup-archive client in interactive mode. If you are using this format as a Linux command `dsmc backup fastback -fbr=\\\\<fbserver>\\REP`, you must specify four backslashes before `<fbserver>` and two backslashes before REP. This is because the Linux shell interprets a backslash as an escape character; the first backslash is treated as an escape character for the following backslash.

Supported Clients

This option is valid for Linux x86/x86_64 clients.

Options File

None. You can specify this option only on the command line. The server can also define or override this option.

Syntax

```
►►──FBReposlocation──repository_location─────────────────────────────►◄
```

Parameters

repository_location
    Specifies the Tivoli Storage Manager FastBack repository location.

Examples

Command line:
```
dsmc backup fastback -fbpolicyname=Policy1
   -fbclientname=fbclient1,fbclient2  -fbserver=myFbDrHub
   -fbreposlocation=\\myFbDrHub\rep_myFbServer
```

   Note: Because Linux is supported only as a dedicated proxy configuration, a repository location is always required on Linux.

Command line:
```
dsmc backup fastback -fbpolicyname=Policy1
   -fbclientname=fbclient1,fbclient2  -fbserver=myFbDrHub
   -fbreposlocation=myFbDrHub  -fbbranch=rep_myFbServer
```

   Note: Because Linux is supported only as a dedicated proxy configuration, a repository location is always required on Linux.

Fbserver

Use the fbserver option with the **backup fastback** or **archive fastback** commands.

The fbserver option specifies the short host name of the Tivoli Storage Manager FastBack server workstation that owns the repository specified by the fbreposlocation option. For a DR Hub, the fbserver option specifies the short name of the FastBack server workstation whose branch repository the Tivoli Storage Manager backup-archive client is connecting to.

The fbserver option is a key to retrieving the necessary user credentials required to connect to the FastBack server repository or the DR Hub server repository for mount processing.

Supported Clients

This option is valid for Linux x86/x86_64 clients.

Options File

None. You can specify this option only on the command line.

Syntax

```
►►── -FBServer── ──server_name────────────────────────────────────►◄
```

Parameters

server_name
    Specifies the short hostname of the machine on which a FastBack server is installed.

Examples

Command line:
    The Tivoli Storage Manager backup-archive client is installed on a Linux proxy client machine. Use this command to archive all FastBack volumes for all Linux FastBack clients that are defined for FastBack policy1:

```
dsmc archive fastback -fbpolicyname=Policy1
  -fbserver=myfbserver
  -fbreposlocation=myfbserver@WORKGROUP
```

    The repository location is required. If you do not provide the repository location, the command will fail.

    The FastBack server name, -myfbserver, is the short host name of the FastBack server where the repository is located.

Command line:
    The repository, rep_server1, is located on a FastBack disaster hub, myFbDrHub.

```
dsmc archive fastback -fbpolicyname="Policy 1"
  -fbserver=myFbDrHub
  -fbreposlocation=\\myFbDrHub\rep_server1
```

    The FastBack server name, -myFbDrHub is the short host name of the FastBack disaster recovery hub server where the repository is located

    The -fbreposlocation specifies the location of the repository. The repository location is required. If you do not provide the repository location, the command fails.

FBServer should point to the short host name of the FastBack DR hub in this case.

Command line:

Archive all volumes protected by FastBack policy named policy1 from the FastBack server named basil:

```
dsmc archive fastback -Fbpolicyname=policy1
  -FBServer=basil -ARCHMC="my_tsm_mgmt_class"
  -fbreposlocation=basil@WORKGROUP
```

Fbvolumename

Use the fbvolumename option with the **backup fastback** or **archive fastback** commands.

The fbvolumename option is the name of one or more comma-separated Tivoli Storage Manager FastBack volumes to back up or archive from the backup proxy. Values for the fbvolumename option are not valid if more than one FastBack client is specified in the fbclientname option.

If you specify multiple FastBack clients in the fbclientname option, you cannot specify values for the fbvolumename option.

Supported Clients

This option is valid for Linux x86/x86_64 clients.

Options File

None. You can specify this option only on the command line.

Syntax

```
          ┌──,─────────────┐
          │        (1)     │
►►──FBVolumename─▼──volume_name─┴──────────────────────►◄
```
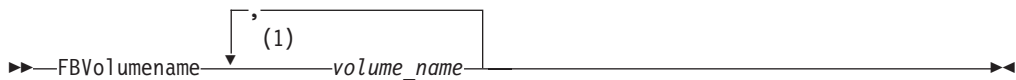
Notes:

1   Specify up to 10 FastBack volume names.

Parameters

volume_name

Specifies the name of the Tivoli Storage Manager FastBack volumes. You can specify up to 10 FastBack volume names.

Important:

When specifying the **archive fastback** or **backup fastback** command:

1. *At least* one FBpolicyName is always *required*.
2. You can specify up to 10 values for FBPolicyName, if no values are specified for both FBClientName *and* FBVolumeName.
3. When you specify a FBClientName value, there must be *only one* value for FBPolicyName.

4. You can specify up to 10 values for FBClientName if *only one* PolicyName is specified, and *no* values for FBVolumeName are specified.

5. When you specify the FBVolumeName option, you can have *only one* FBPolicy, and *only one* FBCLientName specified. You must specify exactly one FBClientName. It cannot be omitted.

6. You can specify multiple FBVolumeNames if condition 5 is satisfied.

7. You must specify the FBReposLocation option.

Examples

Command line:
```
dsmc backup fastback -fbpolicyname=Policy1 -fbclientname=client1
  -fbvolumename=data1,data2 -fbserver=myFbDrHub
  -fbreposlocation=\\myFbDrHub\rep_server1
```

Backs up volumes data1 and data2 from FastBack client client1 found in policy Policy1.

Filelist

Use the filelist option to process a list of files.

You can use the filelist option with the following commands:
- **archive**
- **backup group**
- **delete archive**
- **delete backup**
- **expire**
- **incremental**
- **query archive**
- **query backup**
- **restore**
- **retrieve**
- **selective**

The Tivoli Storage Manager client opens the file you specify with this option and processes the list of files within according to the specific command. With the exception of the **restore** and **retrieve** commands, when you use the filelist option, Tivoli Storage Manager ignores all other file specifications on the command line.

The files (entries) listed in the filelist must adhere to the following rules:
- Each entry must be a fully or partially qualified path to a file or directory or a relative path.
- Each entry must be on a new line.
- Do *not* use wildcard characters.
- Each entry results in the processing of only one object (file or directory).
- If the file name contains any spaces, enclose the file name with quotation marks.
- The filelist can be an MBCS file or a Unicode file with all Unicode entries. For Mac OS X, the filelist can be encoded in the current operating system language or UTF-16.
- Tivoli Storage Manager ignores any entry that is not valid.

The following is an example of a list of files within a filelist:

```
/home/dir/file1
/usr/tivoli/file2
/usr/avi/dir1
/fs1/dir2/file3
"/fs2/Ha Ha Ha/file.txt"
"/fs3/file.txt"
```

If an entry in the filelist indicates a directory, only that directory is processed and not the files within the directory.

If the file name (the `filelistspec`) you specify with the filelist option does not exist, the command fails. Tivoli Storage Manager skips any entries in the filelist that are not valid files or directories. Tivoli Storage Manager logs errors and processing continues to the next entry.

Use file specifications with the **restore** and **retrieve** commands to denote the destination for the restored filelist entries. For example, in the following **restore** command, the file specification `/user/record/` represents the restore destination for all entries in the filelist.

```
restore -filelist=/home/dir/file3 /usr/record/
```

However, in the following **selective** command, the file specification `/usr/record/` is ignored.

```
selective -filelist=/home/dir/file3 /usr/record/
```

If you specify a directory in a filelist for the **delete archive** or **delete backup** command, the directory is not deleted. Filelists that you use with the **delete archive** or **delete backup** command should not include directories.

The entries in the list are processed in the order they appear in the filelist. For optimal processing performance, pre-sort the filelist by file space name and path.

Note: Tivoli Storage Manager might back up a directory twice if the following conditions exist:
- The filelist contains an entry for the directory
- The filelist contains one or more entries for files within that directory
- No backup of the directory exists

For example, your filelist includes the entries `/home/dir/file1` and `/home/dir`. If the `/dir` directory does not exist on the server, the `/home/dir` directory is sent to the server a second time.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──FILEList =── ──filelistspec────────────────────────────────────────────►◄
```

Parameters

filelistspec
>   Specifies the location and name of the file that contains the list of files to process with the command.

Note: When you specify the filelist option on the command line, the subdir option is ignored.

Examples

Command line:
```
sel -filelist=/home/avi/filelist.txt
```

Filename

Use the filename option with the **query systeminfo** command to specify a file name in which to store information.

You can store information gathered from one or more of the following items:
- DSMOPTFILE - The contents of the dsm.opt file.
- DSMSYSFILE - The contents of the dsm.sys file.
- ENV - Environment variables.
- ERRORLOG - The Tivoli Storage Manager error log file.
- FILE - Attributes for the file name that you specify.
- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- OPTIONS - Compiled options.
- OSINFO - Name and version of the client operating system (includes ULIMIT information for UNIX and Linux).
- POLICY - Policy set dump.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually dsmsched.log.
- CLUSTER - AIX cluster information.

Note: The **query systeminfo** command is intended primarily as an aid for IBM support to assist in diagnosing problems, although users who are familiar with the concepts addressed by this information might also find it useful. If you use the console option, no special formatting of the output is performed to accommodate screen height or width. Therefore, the console output might be difficult to read due to length and line-wrapping. In this case, use the filename option with the **query systeminfo** command to allow the output to be written to a file that can subsequently be submitted to IBM support.

Supported Clients

This option is valid for all clients.

Syntax

```
►►──FILEName =── ──outputfilename──────────────────────────────────►◄
```

Parameters

outputfilename
　　Specifies a file name in which to store the information. If you do not specify a file name, by default the information is stored in the dsminfo.txt file.

Examples

Command line:
```
query systeminfo dsmoptfile errorlog -filename=tsminfo.txt
```

Related information

"Console" on page 286

Filesonly

The filesonly option restricts backup, restore, retrieve, or query processing to files *only*.

You cannot restore or retrieve directories from the Tivoli Storage Manager server when using the filesonly option with the **restore** or **retrieve** commands. However, directories with default attributes are created, if required, as placeholders for files that you restore or retrieve.

You can also use the filesonly option with the following commands:
- **archive**
- **incremental**
- **query archive**
- **query backup**
- **restore**
- **restore backupset**
- **restore group**
- **retrieve**
- **selective**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──FILESOnly──────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Command line:
```
dsmc incremental -filesonly
```

Followsymbolic

During a backup operation, the followsymbolic option specifies whether you want to use a symbolic link as a virtual mount point. During a restore or retrieve operation, the followsymbolic option specifies how the Tivoli Storage Manager client restores a directory whose name matches a symbolic link on the restore target file system.

For backup operations, the followsymbolic option can influence the virtualmountpoint option setting. If you use the virtualmountpoint option to specify a symbolic link as a virtual mount point, you must also set the followsymbolic option.

During restore and retrieve operations, followsymbolic can influence how the client handles a symbolic link on the file system. Set followsymbolic only when the client attempts to restore a directory whose name matches a symbolic link on the restore target file system.

If you specify `followsymbolic=no` (the default), the client does not restore the contents of the directory, but returns this error message:

```
ANS4029E Error processing 'filespace name path-name file-name':
unable to build a directory path; a file exists with the same name
as a directory.
```

If you specify `followsymbolic=yes`, the client restores the contents of the directory to the target of the symbolic link.

For example, assume the client backed up a file with this path: `/fs1/dir1/subdir1/file1`. Assume also that a symbolic link `/fs1/dir1`, that exists on the restore target file system, links to the directory `/fs88/dir88/subdir88`. Restore the file with the command:

```
restore /fs1/dir1/subdir1/file1
```

If you specify `followsymbolic=no`, the client does not restore the file, but returns the preceding error message. If you specify `followsymbolic=yes`, the client restores `file1` to the `/fs88/dir88/subdir88/subdir1/file1` directory.

If you restore a symbolic link (not a directory) whose name matches a symbolic link on the restore target file system, the client restores the symbolic link.

Use this option with the **incremental**, **restore**, and **retrieve** commands, or in the client user-options file (`dsm.opt`).

Supported Clients

This option is valid for all UNIX clients except Mac OS X.

Options File

Place this option in the client options file (`dsm.opt`).

Syntax

```
            ┌─No──┐
►►──FOLlowsymbolic─┼─────┼──────────────────────────────────────►◄
            └─Yes─┘
```

Parameters

No Do not back up a virtual mount point that is a symbolic link. Do not restore a directory if the restore target file system contains a symbolic link with matching name. This is the default.

Yes
> Restore the contents of a directory to the target of a symbolic link.

Examples

Options file:
> `followsymbolic Yes`

Command line:
> `-fol=Yes`

Related information

During archive, the"Archsymlinkasfile" on page 268 option determines how the client handles symbolic links.

For operating systems other than Mac OS X, see "Back up symbolic links" on page 171 for more information about how Tivoli Storage Manager handles symbolic links.

Fromdate

Use the fromdate option with the fromtime option to specify a date and time from which you want to search for backups or archives during a restore, retrieve, or query operation.

Files that were backed up or archived before this date and time are not included, although older directories might be included, if necessary, to restore or retrieve the files.

Use the fromdate option with the following commands:
- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──FROMDate =── ──date────────────────────────────────────────────────►◄
```

Parameters

date
> Specifies the date from which you want to search for backup copies or archived files. Enter the date in the format you selected with the dateformat option.
>
> When you include dateformat with a command, it must precede the fromdate, pitdate, and todate options.

Examples

Command line:
```
dsmc query backup -fromdate=12/11/2003 "/Users/van/Documents/*"
```

Command line:
```
dsmc query backup -fromdate=12/11/2003 /home/dilbert/*
```

Fromnode

The fromnode option permits one node to perform commands for another node. A user on another node must use the **set access** command to permit you to query, restore, or retrieve files for the other node.

Use the fromnode option with the following commands:
- **query archive**
- **query backup**
- **query filespace**
- **query image**
- **query mgmtclass**
- **restore**
- **restore group**
- **restore image**
- **retrieve**

Supported Clients

This option is valid for all clients.

Syntax

```
►►──FROMNode =── ──node──────────────────────────────────────────►◄
```

Parameters

node
> Specifies the node name on a workstation or a file server whose backup copies or archived files you want to access.

Examples

Command line:
```
dsmc query archive -fromnode=bob -subdir=yes  "/Users/van/Documents/
*"
```

Command line:
```
dsmc query archive -fromnode=bob -subdir=yes  "/home/jones/*"
```

Fromowner

The fromowner option specifies an alternate owner from which to restore backup versions or archived files or images. The owner must give access to another to use the files or images.

For example, to restore files from the /home/devel/proja directory belonging to *usermike* on system *puma*, and place the restored files in a directory you own named /home/id/proja, enter the following command:

```
dsmc restore -fromowner=usermike -fromnode=puma /home/devel/proja/
/home/id/proja/
```

Note: Archiving image restores does not apply to Mac OS X operating systems.

Non-root users can specify fromowner=*root* to access files owned by the root user if the root user has granted them access.

Note: If you specify the fromowner option without the fromnode option, the active user must be on the same node as the fromowner user.

Use the fromowner option with the following commands:
- **query archive**
- **query backup**
- **query group**
- **query image**
- **restore**
- **restore image**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

▶▶──FROMOwner =── ──*owner*──◀◀

Parameters

owner
    Name of an alternate owner.

Examples

Command line:
```
dsmc query archive "/home/id/proja/*" -fromowner=mark
```

Fromtime

Use the fromtime option with the fromdate option to specify a beginning time from which you want to search for backups or archives during a restore, retrieve, or query operation.

Tivoli Storage Manager ignores this option if you do not specify the fromdate option.

Use the fromtime option with the following commands:
- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──FROMTime =── ──time────────────────────────────────────────────────────►◄
```

Parameters

time

Specifies a beginning time on a specific date from which you want to search for backed up or archived files. If you do not specify a time, the time defaults to 00:00:00. Specify the time in the format you selected with the timeformat option.

When you include the timeformat option in a command, it must precede the fromtime, pittime, and totime options.

Examples

Command line:

```
dsmc q b -timeformat=4 -fromt=11:59AM -fromd=06/30/2003 -tot=11:59PM
-tod=06/30/2003 /home/*
```

Groupname

Use the groupname option with the **backup group** command to specify the name for a group. You can only perform operations on new groups or the current active version of the group.

Supported Clients

This option is valid for all UNIX and Linux clients except Mac OS X.

Syntax

```
►►──GROUPName =── ──name────────────────────────────────────────────────────►◄
```

Parameters

name

Specifies the name of the group which contains the files backed up using the filelist option. Directory delimiters are not allowed in the group name since the group name is not a file specification, but a name field.

Examples

Command line:

```
backup group -filelist=/home/dir1/filelist1 -groupname=group1
-virtualfsname=/virtfs -mode=full
```

Groups

Authorized User: The groups option specifies UNIX and Linux system groups on your workstation that you want to authorize to request Tivoli Storage Manager services from the server. You can use the groups option more than once to specify several group names.

If you do not specify group names with the groups option, or user IDs with the users option, *all* users can request Tivoli Storage Manager services. If you use both the groups option and the users option, only users specified with these options can request Tivoli Storage Manager services. A root user is always authorized to request services.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file *within* a server stanza.

Syntax

```
>>──┬─GRoups──┬──────groupname──┬─────────────────────────────────><
```

Parameters

groupname
    Specifies the name of a group you want to authorize to request Tivoli Storage Manager services.

Examples

Options file:

```
groups  dsmcdev  group1  test1  test2  design1
groups  endicott almaden qadev qadev1 tools23
```

Command line:
    Does not apply.

Guitreeviewafterbackup

The guitreeviewafterbackup option specifies whether the client returns to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client user-options file (dsm.opt) *or* the dsm.sys file. You can set this option on the **General** tab, **Return to tree window after function completed** check box of the Preferences editor.

Syntax

```
►►──GUITREEViewafterbackup──┬─No──┬──────────────────────────────────────►◄
                            └─Yes─┘
```

Parameters

No Returns you to the Tivoli Storage Manager main window after a successful operation completes. This is the default.

Yes
    Returns you to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.

Examples

Options file:
    guitreeviewafterbackup yes

Command line:
    Does not apply.

Httpport

The httpport option specifies a TCP/IP port address for the Web client.

You must be an authorized user to use the httpport option.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Web Client** tab → **HTTP Port** field of the Preferences editor.

Syntax

```
►►──HTTPport── ──port_address──────────────────────────────────────────────►◄
```

Parameters

port_address
    Specifies the TCP/IP port address that is used to communicate with the Web client. The range of values is 1000 through 32767; the default is 1581.

Examples

Options file:
>   httpport 1502

Command line:

Command line:
>   Does not apply.

Hsmreparsetag

The hsmreparsetag option specifies a unique reparse tag that is created by an HSM product installed on your system.

Many HSM products use reparse points to retrieve or recall migrated files. After a file is migrated, a small stub file, with the same name as the original file, is left on the file system. The stub file is a reparse point that triggers a recall of the original file when a user or application accesses the stub file. The reparse point includes a unique identifier called a *reparse tag* to identify which HSM product migrated the file.

If the Tivoli Storage Manager backup-archive client does not recognize the reparse tag in a stub file, the backup-archive client causes the HSM product to recall the original file. You can prevent files from being recalled if you specify the reparse tag with the hsmreparsetag option.

The backup-archive client recognizes the reparse tag of HSM products from the following companies:
- International Business Machines Corp.
- Wisdata System Co. Ltd.
- BridgeHead Software Ltd.
- CommVault Systems, Inc.
- Data Storage Group, Inc.
- Enigma Data Solutions, Ltd.
- Enterprise Data Solutions, Inc.
- Global 360
- GRAU DATA AG
- Hermes Software GmbH
- Hewlett Packard Company
- International Communication Products Engineering GmbH
- KOM Networks
- Memory-Tech Corporation
- Moonwalk Universal
- Pointsoft Australia Pty. Ltd.
- Symantec Corporation

If the HSM product you use is not in the preceding list, use the hsmreparsetag option to specify the reparse tag. Ask your HSM product vendor for the reparse tag used by the product.

Supported clients

This option is valid for all Windows clients.

Option file

Place this option in the client options file (dsm.opt).

Syntax

```
►►──HSMREPARSETAG──reparse_tag_value──────────────────────────────────────►◄
```

Parameters

reparse_tag_value
   A decimal (base 10) or hexadecimal (base 16) value that specifies the reparse
   tag.

Examples

Options file:
   Specify an HSM reparse tag in decimal format:
   ```
   hsmreparsetag 22
   ```

   Specify an HSM reparse tag in hexadecimal format:
   ```
   hsmreparsetag 0x16
   ```

Command line:
   Does not apply.

Ieobjtype

Use the ieobjtype option to specify an object type for a client-side data
deduplication operation within include-exclude statements.

The ieobjtype option is an additional parameter to the include.dedup or
exclude.dedup options.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API also
supports this option.

Options File

Place this option in the system-options file (dsm.sys). You can set this option on
the **Include/Exclude** tab of the Preferences editor. The option can be set in the
client option set on the Tivoli Storage Manager server.

Syntax

```
              ┌─File──┐
►►──IEObjtype──┴─Image─┴───────────────────────────────────────────────────►◄
```

Parameters

File

Specifies that you want to include files for, or exclude files from, client-side data deduplication processing. *File* is the default.

Image

Specifies that you want to include images for, or exclude images from, client-side data deduplication processing.

Examples

Options file:

```
exclude.dedup /home/*/* ieobjtype=image
```

Command line:

Does not apply.

Related reference

"Exclude options" on page 321

"Include options" on page 351

Ifnewer

The ifnewer option replaces an existing file with the latest backup version only if the backup version is newer than the existing file.

Only active backups are considered unless you also use the inactive or latest options.

Note: Directory entries are replaced with the latest backup version, whether the backup version is older or newer than the existing version.

Use the ifnewer option with the following commands:
* **restore**
* **restore backupset**
* **restore group**
* **retrieve**

Note: This option is ignored if the replace option is set to *No*.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►──IFNewer──►◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/Users/grover/Documents/*" -sub=y -rep=y -ifnewer
```

```
dsmc restore "/home/grover/*" -sub=y -rep=y -ifnewer
```

Imagegapsize

Use the imagegapsize option with the **backup image** command, in the options file, or with the include.image option to specify the minimum size of empty regions on a volume that you want to skip during image backup.

Use this option for LAN-based and LAN-free image backup.

For example, if you specify a gap size of 10, this means that an empty region on the disk that is larger than 10 k in size is not backed up. Gaps that are exactly 10 k is backed up. Empty regions that are exactly 10 k and that are smaller than 10 k is backed up, even though they do not contain data. However, an empty region that is smaller than 10 k is backed up, even though it does not contain data. A smaller image gap size means less data needs to be transferred, but with potentially decreased throughput. A larger image gap size results in more data being transferred, but with potentially better throughput.

Supported Clients

This option is valid for AIX, Linux, and JFS2 clients only. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the server stanza of the client systems options file (dsm.sys), or in the include.image statement in the dsm.sys file.

Syntax

```
►►──IMAGEGapsize── ──size─────────────────────────────────────────►◄
```

Parameters

size

Specifies the minimum size of empty regions in an AIX JFS2 file system that should be skipped during an image backup. You can specify k (kilobytes) m (megabytes) or g (gigabytes) qualifiers with the value. Without a qualifier, the value is interpreted in kilobytes. Valid values are 0 through 4294967295 KB. If you specify a value of 0, all blocks, including unused blocks at the end of the volume, is backed up. If you specify any value other than 0, unused blocks at the end of the volume will *not* be backed up. For LAN-based and LAN-free image backup the default value is 32 k. This option is applicable to both static and snapshot-based image backup.

Note: This option is valid for AIX JFS2 file systems. If you specify an imagegapsize that is greater than 0 for a file system other than AIX JFS2, you get a warning message.

Examples

Options file:

Add the following to the server stanza in the dsm.sys file: `imagegapsize 1m`

> **Include-exclude list example:** `include.image /kalafs1`
> `imagegapsize=-128k`

Command line:
> `-imagegapsize=64k`

Imagetofile

Use the imagetofile option with the **restore image** command to specify that you want to restore the source image to a file.

You might need to restore the image to a file if bad sectors are present on the target volume, or if you want to manipulate the image data. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume.

Linux supports mounting an image file as a logical volume, so you can get access to file data within the image. The following are some examples:

- The file system /usr has been backed up by Tivoli Storage Manager. The following command restores the file system image to the file /home/usr.img:
  ```
  # dsmc restore image /usr /home/usr.img -imagetofile
  ```
- To mount the image file at the /mnt/usr directory, the following mount command can be executed:
  ```
  # mount /home/usr.img /mnt/usr -o loop=/dev/loop0
  ```

Now the image contents are available from /mnt/usr as if a regular file system was mounted at that directory.

Supported Clients

This option is valid for AIX, HP-UX, all Linux, and Solaris clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──IMAGETOfile────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Command line:
> `dsmc restore image /usr /home/usr.img -imagetofile`

Inactive

Use the inactive option to display both active and inactive objects.

You can use the inactive option with the following commands:
- **delete group**
- **query backup**
- **query group**
- **query image**

- **query nas**
- **restore**
- **restore group**
- **restore image**
- **restore nas**

Important: When using the inactive option during a restore operation, also use the pick or some other filtering option because, unlike the latest option, all versions are restored in an indeterminate order. This option is implicit when pitdate is used.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──INActive────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Command line:
```
dsmc restore "/Users/zoe/Documents/*" -inactive -pick
```

Command line:
```
dsmc restore "/home/zoe/*" -inactive -pick
```

Inclexcl

The inclexcl option specifies the path and file name of an include-exclude options file.

You must be an authorized user to use the inclexcl option.

Multiple inclexcl statements are permitted. However, you must specify this option for each include-exclude file.

Ensure that you store your include-exclude options file in a directory to which all users have read access, such as /etc.

When processing occurs, the include-exclude statements within the include-exclude file are placed in the list position occupied by the inclexcl option, in the same order, and processed accordingly.

If you have the HSM client installed on your workstation, you can use an include-exclude options file to exclude files from backup and space management, from backup only or from space management only.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Include-Exclude** tab of the Preferences editor.

Syntax

```
►►──INCLExcl── ──filespec───────────────────────────────────────────────◄◄
```

Parameters

filespec
>   Specifies the path and file name of *one* include-exclude options file.

Examples

Options file:
>   INCLExcl /Users/user1/Documents/backup.excl
>
>   inclexcl /usr/dsm/backup.excl
>   inclexcl /etc/inclexcl.def

Command line:
>   Does not apply.

Related information

For more information about creating an include-exclude options file, see "Creating an include-exclude list" on page 93.

Considerations for Unicode-enabled clients

The include-exclude file can be in Unicode or non-Unicode format. If you specify a non-Unicode include-exclude file, that file must be in the same code page that the client is running.

For example, a non-Unicode include-exclude file on an English Mac OS X system cannot contain Japanese characters.

Note: This section applies to Mac OS X only.

A Unicode include-exclude file provides the following benefits:
*   Names with characters from another code page no longer have to be wildcarded.
*   File names and directories from any code page can be *fully specified* for the Unicode-enabled client to process.

To create an include-exclude file in Unicode format, perform the following steps:
1. Open TextEdit. Click **Format →Make PlainText**.
2. Enter your include and exclude statements.
3. Click **File** and then click **Save As**. The **Save As** window displays.
4. From **PlainText Encoding**, select **Unicode (UTF-16)**, specify the file and target directory, and then save the file.
5. Place an inclexcl option specifying the include-exclude file you just created in your dsm.sys file.
6. Restart the Tivoli Storage Manager client.

Include options

The include options specify objects that you want to include for backup and archive services.

You must be an authorized user to use the include options.

The include options specify any of the following:
- Objects within a broad group of excluded objects that you want to include for backup and archive services.
- Objects within a broad group of excluded objects that you want to include for backup, archive, image, and space management services.
- Files that are included for backup or archive processing that you want to include for encryption processing.
- Files that are included for backup or archive processing that you also want to include for compression processing.
- Objects to which you want to assign a specific management class.
- A management class to assign to all objects to which you do not explicitly assign a management class.
- File spaces to which you want to assign memory-efficient backup processing
- File spaces where you want to use the diskcachelocation option to cause specific file systems to use different, specific locations for their disk cache.

If you do not assign a specific management class to objects, Tivoli Storage Manager uses the default management class in the active policy set of your policy domain. Use the **query mgmtclass** command to display information about the management classes available in your active policy set.

Remember: Tivoli Storage Manager compares the files it processes against the patterns specified in the include-exclude statements, reading from the bottom to the top of the options file.

Note:
1. The exclude.fs and exclude.dir statements override all include statements that match the pattern.
2. The server can also define these options with the inclexcl option.

Supported Clients

This option is valid for all clients. The server can also define include.fs.nas.

Options File

Place this option in the dsm.sys file within a server stanza. You can set these options on the **Include-Exclude** tab, **Define Include-Exclude Options** section of the Preferences editor.

Syntax

►►—*options*— —*pattern*———————————————————————►◄
                └— —*optional_parameter*—┘

include, include.backup, include.file
> Use these options to include files or assign management classes for backup processing.
>
> The include option affects archive and backup processing. If you want to assign different management classes for archive and backup processing, always specify include.archive and include.backup with their own management classes. In this example, the `archmc` management class is assigned when an archive operation is performed. The management class is assigned when an archive operation is performed because include.backup is used only for backup processing, and not for archive processing.
>
> ```
> include.archive /home/test/* archmc
> include.backup /home/test/*
> ```

include.archive
> Includes files or assigns management classes for archive processing.

include.attribute.symlink
> Includes a file or a group of files that are symbolic links or aliases, within a broad group of excluded files for backup processing only.
>
> **Note:** For Mac OS X, aliases are included.

include.compression
> Includes files for compression processing if you set the compression option to *yes*. This option applies to backups and archives.

include.dedup
> Includes files for client-side data deduplication. To control a client-side data deduplication operation, specify ieobjtype as the value of the include.dedup option. By default, all data deduplication-eligible objects are included for client-side data deduplication.
>
> Valid ieobjtype parameters are:
> > File
> > Image
>
> The default is File.

include.encrypt
> Includes the specified files for encryption processing. By default, Tivoli Storage Manager does not perform encryption processing.
>
> **Important:** The include.encrypt option is the only way to enable encryption on the Backup-Archive client. If no include.encrypt statements are used, encryption does not occur.

include.fs

> For AIX JFS2 file systems: Use the snapshotcachesize option in the dsm.sys file or with the include.fs option, to specify an appropriate snapshot size so that all old data blocks can be stored while the snapshot-based file backup or archive occurs.
>
> To control how Tivoli Storage Manager processes your file space for incremental backup, you can specify these additional options in your dsm.sys file, or as values of the include.fs option: diskcachelocation and memoryefficientbackup. If these options appear both in the options file and an include.fs option, the include.fs values are used for the specified file space in place of any values in an option file or on the command line.
>
> **Examples:**

Options file or include/exclude file: Each of the `include.fs`, `memoryefficientbackup` and `diskcachelocation` options must be on the **same** line in the options file.

```
include.fs /home
     memoryefficientbackup=diskcachemethod
     diskcachelocation=/usr
include.fs /usr
     memoryefficientbackup=diskcachemethod
     diskcachelocation=/home
include.fs /Volumes/hfs3
     memoryefficientbackup=diskcachemethod
     diskcachelocation=/Volumes/hfs2
 AIX JFS2 filesystems only: include.fs
     /kalafs1 snapshotproviderfs=JFS2
```

include.fs.nas

Use the include.fs.nas option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system image backup, using the toc option with the include.fs.nas option in your dsm.sys file. This option is valid for AIX and Solaris clients *only*.

include.image

Includes a file space or logical volume, or assigns a management class when used with the **backup image** command. The **backup image** command ignores all other include options.

For Linux x86/x86_64 clients: Use the snapshotcachesize option in these situations:

- With the **backup image** command
- In the dsm.sys file
- With the include.image option

Using the snapshotcachesize option in these situations lets you specify an appropriate snapshot size, so that all old data blocks can be stored while the image backup occurs.

A snapshot size of 100 percent ensures a valid snapshot.

For AIX JFS2 file systems: Use the snapshotcachesize option in these situations:

- With the **backup image** command
- In the dsm.sys file
- With the include.image option

Using the snapshotcachesize option in these situations lets you specify an appropriate snapshot size, so that all old data blocks can be stored while the image backup occurs.

This option is valid for AIX, HP-UX, Solaris, and all Linux clients *only*.

Parameters

pattern

Specifies the objects to include for backup or archive processing or to assign a specific management class.

Note: For NAS file systems: You must prefix the NAS node name to the file specification to specify the file server to which the include statement applies. If

you do not specify a NAS node name, the file system identified refers to the NAS node name specified in the client system-options file (dsm.sys) or on the command line.

If the pattern begins with a single or double quotation mark, or contains any embedded blanks or equal signs, you must surround the value in either single (') or double (") quotation marks. The opening and closing quotation marks must be the same type of quotation marks.

For the include.image option, the pattern is the name of a mounted file system or raw logical volume.

optional_parameter

mgmtclassname
Specifies the name of the management class to assign to the objects. If a management class is not specified, the default management class is used.

Use the following command to bind a management class to a group:

```
include /myvfs/mygroupname ManagmentClass
```

Table 91. Other optional parameters

optional_parameter	Use with option
ieobjtype "Ieobjtype" on page 345	include.dedup
memoryefficientbackup "Memoryefficientbackup" on page 368	include.fs
diskcachelocation "Diskcachelocation" on page 302	include.fs
dynamicimage "Dynamicimage" on page 309	include.image
postsnapshotcmd "Postsnapshotcmd" on page 386	include.image
presnapshotcmd "Presnapshotcmd" on page 392	include.image
snapshotcachesize "Snapshotcachesize" on page 422	include.image
snapshotproviderfs "Snapshotproviderfs" on page 423	include.image
snapshotproviderimage "Snapshotproviderimage" on page 424	include.image

Examples

Options file:

```
include /home/proj/text/devel.*
include /home/proj/text/* textfiles
include * managall
include /WAS_ND_NDNODE mgmtclass
include /WAS_APPNODE mgmtclass
include.image /home/*/*
include.archive /home/proj/text/
  * myarchiveclass
include.backup /home/proj/text/
  * mybackupclass
include.compression /home/proj/text/
```

```
            devel.*
          include.encrypt /home/proj/gordon/*
          include.fs.nas netappsj/vol/vol0
            homemgmtclass

          include.dedup /Users/Administrator/Documents/Important/.../*

          AIX only:

          include.image /home
            MGMTCLASSNAME
            snapshotproviderimage=JFS2
            snapshotcachesize=40
          include.image /home
            snapshotproviderimage=NONE
          include.fs /kalafs1
            snapshotproviderfs=JFS2

          LINUX only:

          include.image /home
            snapshotproviderimage=LINUX_LVM
          include.image /myfs1 dynamicimage=yes
          include.image /home MGMTCLASSNAME
            snapshotproviderimage=NONE
          include.image /myfs1 dynamicimage=yes
          include.attribute.symlink /home/spike/.../*
          include.fs /usr
            memoryefficientbackup=diskcachemethod
```

Command line:
    Does not apply.

Related information

"Snapshotcachesize" on page 422

"Toc" on page 444

Controlling symbolic link and alias processing

Tivoli Storage Manager treats symbolic links and aliases (aliases apply to Mac OS X only) as actual files and backs them up. However, the file referenced by the symbolic link is not backed up.

In some cases symbolic links and aliases can be easily recreated and need not be backed up. In addition, backing up these symbolic links or aliases can increase backup processing time and occupy a substantial amount of space on the Tivoli Storage Manager server.

You can use the exclude.attribute.symlink option to exclude a file or a group of files that are symbolic links or aliases from backup processing. If necessary, you can use the include.attribute.symlink option to include symbolic links or aliases within broad group of excluded files for backup processing. For example, to exclude all symbolic links or aliases from backup processing, except those that exist under the /home/spike directory, enter these statements in your dsm.sys file:

```
  exclude.attribute.symlink  /.../*
  include.attribute.symlink /home/spike/.../*
```

Related reference

"Exclude options" on page 321

Compression and encryption backup processing

This topic lists some items to consider if you want to include specific files or groups of files for compression and encryption processing during a backup or archive operation.

- You must set the compression option to *yes* to enable compression processing. If you do not specify the compression option or you set the compression option to *no*, Tivoli Storage Manager does not perform compression processing.
- Tivoli Storage Manager processes exclude.fs, exclude.dir, and other include-exclude statements first. Tivoli Storage Manager then considers any include.compression and include.encrypt statements. For example, consider the following include-exclude list:

```
exclude /home/jones/proj1/file.txt
include.compression /home/jones/proj1/file.txt
include.encrypt /home/jones/proj1/file.txt
```

Tivoli Storage Manager examines the `exclude /home/jones/proj1/file.txt` statement first and determines that `/home/jones/proj1/file.txt` is excluded from backup processing and is, therefore, not a candidate for compression and encryption processing.

- Include-exclude compression and encryption processing is valid for backup and archive processing *only*.

Related reference

"Compression" on page 285

Processing NAS file systems

Use the include.fs.nas option to bind a management class to NAS file systems and to control if Table of Contents information is saved for the file system backup.

Note: The include.fs.nas option does not apply to incremental snapshot difference incremental backup.

A NAS file system specification uses the following conventions:

- NAS nodes represent a new node type. The NAS node name uniquely identifies a NAS file server and its data to Tivoli Storage Manager. You can prefix the NAS node name to the file specification to specify the file server to which the include statement applies. If you do not specify a NAS node name, the file system you specify applies to all NAS file servers.
- Regardless of the client platform, NAS file system specifications use the forward slash (/) separator, as in this example: `/vol/vol0`.

Use the following syntax:

```
►►──pattern── mgmtclassname─ toc=value────────────────────────────────────►◄
```

Parameters

pattern

Specifies the objects to include for backup services, to assign a specific management class, or to control TOC creation. You can use wildcards in the pattern.

mgmtclassname
> Specifies the name of the management class to assign to the objects. If a management class is not specified, the default management class is used.

toc=value
> See "Toc" on page 444 for more information.

For example, to assign a management class to the /vol/vol1 file system of a NAS node called netappsj, specify the following include statement:

```
include.fs.nas netappsj/vol/vol1 nasMgmtClass toc=yes
```

Incrbydate

Use the incrbydate option with the **incremental** command to back up new and changed files with a modification date later than the last incremental backup stored at the server, unless you exclude the file from backup.

Important: Files that are modified or created after their respective directory was processed by the Tivoli Storage Manager client, but before the incremental-by-date backup completes, are not backed up and will not be backed up in future incremental-by-date backups, unless the files are modified again. For this reason, a regular incremental backup should be run periodically, without specifying the incrbydate option.

An incremental-by-date updates the date and time of the last incremental at the server. If you perform an incremental-by-date on only part of a file system, the date of the last full incremental is not updated and the next incremental-by-date backs up these files again.

Important:

The last incremental backup time refers to the server time and the file modification time refers to the client time. If the client and server time are not synchronized, or the client and server are in different time zones, this affects incremental-by-date backup with mode=incremental.

The last incremental backup time refers to the server time and the file modification time refers to the client time. If the client and server time are not synchronized, or the client and server are in different time zones, this affects incremental-by-date backup and image backup with mode=incremental.

Both full incrementals and incrementals-by-date back up new and changed files. An incremental-by-date takes less time to process than a full incremental and requires less memory. However, unlike a full incremental, an incremental-by-date does not maintain current server storage of *all* your workstation files because:

- It does not expire backup versions of files that are deleted from the workstation.
- It does not rebind backup versions to a new management class if the management class has changed.
- It does not back up files with attributes that have changed, such as Access control list (ACL) data, unless the modification dates and times have also changed.
- It ignores the copy group frequency attribute of management classes.

Note: If you have limited time during the week to perform backups, but extra time on weekends, you can maintain current server storage of your workstation files by performing an incremental backup with the incrbydate option on weekdays and a full incremental backup on weekends.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──INCRbydate──────────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Command line:
```
dsmc incremental -incrbydate
```

Incremental

Use the incremental option with the **restore image** command to ensure that any changes that were made to the base image are also applied to the restored image.

If you also use the deletefiles option, changes include the deletion of files and directories that were in the original image but later deleted from the workstation.

Note: Using the incremental option with the **restore image** command to perform a dynamic image backup is not supported.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris *only*. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──INCREmental──────────────────────────────────────────────────────►◄
```

Examples

Command line:
```
res i "/home/devel/projecta/*" -incremental
```

Lanfreecommmethod

The lanfreecommmethod option specifies the communications protocol between the Tivoli Storage Manager client and Storage Agent. This enables processing between the client and the SAN-attached storage device.

You must be an authorized user to use the lanfreecommmethod option. If you are using LAN failover, you must have lanfreecommmethod in the dsm.sys file *within* a server stanza.

For AIX, HP-UX, Linux and Solaris, use the lanfreeshmport option to specify the shared memory port number where the Storage Agent is listening.

Supported Clients

This option is valid for AIX, HP-UX, Linux, and Solaris clients *only*.

Options File

Place this option in the dsm.sys file *within* a server stanza.

Syntax

►►──LANFREECommmethod── ──*commmethod*──────────────────────────────────◄►

Parameters

commmethod
    Specifies the supported protocol for your Tivoli Storage Manager client:

    TCPip
        The Transmission Control Protocol/Internet Protocol (TCP/IP) communication method.

        Use the lanfreetcpport option to specify the TCP/IP port number where the Storage Agent is listening. The TCP/IP communication method is the default for *non-root* users on all supported platforms. On Linux, this is the only communication method you can use.

    V6Tcpip
        Indicates that either TCP/IP Version 4 or Version 6 should be used, depending on the system configuration and results of a domain name service lookup. The only time this is not true is when 'dsmc schedule' is used *and* schedmode is prompt. A valid DNS environment must be available.

    SHAREdmem
        Use the shared memory communication method when the client and Storage Agent are running on the same system. Shared memory provides better performance than the TCP/IP protocol. This is the default communication method for AIX, HP-UX, and Solaris *root* users. When specifying this communication method on AIX, the backup-archive client user can be logged in as root or non-root, as long as the Storage Agent is running as root. If the Storage Agent is not running as root, the user ID running the backup-archive client must match the user ID running the Storage Agent. *SHAREdmem* is not valid for Linux clients.

Examples

Options file:
    `lanfreecommmethod tcp`

    Use only TCP/IP Version 4

    `lanfreecommmethod V6Tcpip`

Use both TCP/IP Version 4 or Version 6, depending on how the system is configured and the results of a domain name service lookup.

Command line:
```
-lanfreec=tcp
```
```
-lanfreec=V6Tcpip
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Related information

"Lanfreeshmport"

"Lanfreetcpport" on page 361

Lanfreeshmport

Use the lanfreeshmport option when lanfreecommmethod=*SHAREdmem* is specified for communication between the Tivoli Storage Manager client and Storage Agent. This enables processing between the client and the SAN-attached storage device.

You must be an authorized user to use the lanfreeshmport option.

Supported Clients

This option is valid for AIX, HP-UX, and Solaris clients *only*.

Options File

Place this option in the client system-options file (dsm.sys) *within* a server stanza.

Syntax

```
►►──LANFREEShmport── ──port_address────────────────────────────────────►◄
```

Parameters

port_address
     Specifies the number that is used to connect to the storage agent. The range of values is 1 through 32767; the default is 1.

Examples

Options file:
```
lanfrees 1520
```

Command line:
```
-lanfrees=1520
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Related information

Lanfreetcpport

The lanfreetcpport option specifies the TCP/IP port number where the Tivoli Storage Manager Storage Agent is listening.

You must be an authorized user to use the lanfreetcpport option.

Use this option when you specify lanfreecommmethod=*TCPip* for communication between the Tivoli Storage Manager client and Storage Agent. Do not specify the lanfreetcpport option if you want to use the NAMedpipes communication method for LAN-free communication.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris clients *only*.

Options File

Place this option in the dsm.sys file *within* a server stanza.

Syntax

▶▶──LANFREETCPport── ──*port_address*──────────────────────────────────────◀◀

Parameters

port_address
> Specifies the TCP/IP port number where the Storage Agent is listening. The range of values is 1 through 32767; the default is 1500.

> **Note:** The client lanfreetcpport value must match Storage Agent tcpport value for communications with the Storage Agent (virtual server). The client tcpport value must match the server tcpport value for communications with the actual server.

Examples

Options file:
```
lanfreetcp 1520
```

Command line:
```
-lanfreetcp=1520
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Related information

Lanfreetcpserveraddress

The lanfreetcpserveraddress option specifies the TCP/IP address for a Tivoli Storage Manager Storage Agent.

Use this option when you specify lanfreecommmethod=*TCPip* or *V6Tcpip* for communication between the Tivoli Storage Manager client and Storage Agent.

Overriding the default for this option is useful when configuring LAN-free in an environment where the client and storage agent are running on different systems. You can obtain this Storage Agent address from your administrator.

You must be an authorized user to use the lanfreetcpserveraddress option.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris clients *only*.

Options File

Place this option in the client system-options file.

Syntax

►►──LANFREETCPServeraddress── ──*stagent_address*──────────────────────────────►◄

Parameters

stagent_address
Specifies a 1 to 64 character TCP/IP address for a server. Specify a TCP/IP domain name or a numeric IP address. The numeric IP address can be either a TCP/IP v4 or TCP/IP v6 address. The default value is 127.0.0.1 (localhost).

Examples

Options file:

        LANFREETCPServeraddress stagent.bayamon.ibm.com

        LANFREETCPServeraddress 192.168.1.50

Command line:
        Does not apply.

Latest

Use the latest option to restore the most recent backup version of a file, even if the backup is inactive.

You can use the latest option with the following commands:
* **restore**
* **restore group**

If you are performing a point-in-time restore (using the pitdate option), it is not necessary to specify latest since this option is implicit when pitdate is used.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──LATest──────────────────────────────────────────────────────────────◄◄
```

Parameters

There are no parameters for this option.

Examples

Command line:
```
dsmc restore "/Users/devel/projecta/*" -latest
```

Command line:
```
dsmc restore "/home/devel/projecta/*" -latest
```

Localbackupset

The localbackupset option specifies whether the Tivoli Storage Manager GUI bypasses initial logon with the Tivoli Storage Manager server to restore a local backup set on a standalone workstation.

If you set the localbackupset option to *yes*, the GUI does not attempt initial logon with the server. In this case, the GUI only enables the restore functionality.

If you set the localbackupset option to *no* (the default), the GUI attempts initial logon with the server and enables all GUI functions.

Note: The **restore backupset** command supports restore of local backup sets on a standalone workstation *without* using the localbackupset option.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza.

Syntax

```
            ┌─No──┐
►►──LOCALbackupset─┤     ├─────────────────────────────────────────────◄◄
            └─Yes─┘
```

Parameters

No Specifies that the GUI attempts initial logon with the server and enables all functions. This is the default.

Yes
   Specifies that the GUI does not attempt initial logon with the server and enables only the restore functionality.

Examples

Options file:
```
localbackupset yes
```

This option is not valid with the **dsmc** command-line client. (Does not apply to Mac OS X)

Related information

"Restore Backupset" on page 529

Makesparsefile

Use the makesparsefile option with the **restore** or **retrieve** commands to specify how sparse files are recreated.

Sparse files do not have disk space allocated for every block in the whole address space, leading to holes within the file. The Tivoli Storage Manager client detects sparse files during a backup operation and marks them as sparse on the Tivoli Storage Manager server. Holes are detected by their content, which is always zeros.

If you set the makesparsefile option to *yes* (default), holes within the file are not written to disk so no additional disk space is allocated during a restore.

If you set the makesparsefile option to *no*, holes are not recreated, leading to disk blocks allocated for the whole address space. This might result in a larger amount of used disk space. Ensure that you have enough disk space to restore all data.

On some UNIX and Linux systems, it might be necessary to back up system specific files as non-sparse files. Use the makesparsefile option for files where the existence of physical disk blocks is required, such as ufsboot on Solaris, which is executed during boot time. The boot file loader of the operating system accesses physical disk blocks directly and does not support sparse files.

Supported Clients

This option is valid for all UNIX and Linux clients except Mac OS X.

Options File

Place this option in the client user options file (dsm.opt).

Syntax

```
▶▶──MAKesparsefile──┬─Yes─┬──────────────────────────────▶◀
                    └─No──┘
```

Parameters

Yes
> Specifies that holes within the file are not written so that no additional disk space is allocated during a restore. This is the default.

No Specifies that holes are not recreated leading to disk blocks allocated for the
whole address space.

Examples

Options file:
```
makesparsefile no
```

Command line:
```
-makesparsefile=no
```

This option is valid only on the initial command line. It is not valid in interactive
mode.

Managedservices

The managedservices option specifies whether the Tivoli Storage Manager client
acceptor service manages the scheduler, the Web client, or both.

Restriction: You cannot use the dsmcad for scheduling when you set the
sessioninitiation option to *serveronly*.

The client acceptor daemon serves as an external timer for the scheduler. When the
scheduler is started, it queries the server for the next scheduled event. The event is
either executed immediately or the scheduler exits. The client acceptor daemon
restarts the scheduler when it is time to execute the scheduled event.

Note:
1. If you set the schedmode option to *prompt*, the server prompts the client
   acceptor daemon when it is time to run the schedule. The scheduler connects to
   and disconnects from the server when the client acceptor daemon is first
   started.

   The dsmc schedule command cannot be used when both schedmode *prompt*
   and commmethod *V6Tcpip* are specified.
2. For Mac OS X, if you do not specify the managedservices option, the client
   acceptor daemon manages the scheduler program by default.
3. Set the passwordaccess option to *generate* in your dsm.sys file and generate a
   password, so Tivoli Storage Manager can manage your password automatically.

Using the client acceptor daemon to manage the scheduler service can provide the
following benefits:
* Memory retention problems that can occur when using traditional methods of
  running the scheduler are resolved. Using the client acceptor daemon to manage
  the scheduler requires very little memory between scheduled operations.
* The client acceptor daemon can manage both the scheduler program and the
  Web client, reducing the number of background processes on your workstation.
* To use the Web client, you must specify this option in the client system-options
  file.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not
support this option. The server can also define this option.

Options File

Place this option in the client system-options file (dsm.sys) *within* a server stanza. You can set this option on the **Web Client** tab of the Preferences editor.

Syntax

```
►►──MANAGEDServices──mode────────────────────────────────────────────►◄
```

Parameters

mode

Specifies whether the client acceptor daemon manages the scheduler, the Web client, or both.

webclient

Specifies that the client acceptor daemon manages the Web client. This is the default for UNIX and Linux. Both *webclient* and *schedule* are the defaults for Mac OS X.

schedule

Specifies that the client acceptor daemon manages the scheduler. This is the default for Mac OS X.

none

For Mac OS X, specifies that the client acceptor daemon not manage the Web client or schedules. Set *managedservices none* to enable the **dsmc schedule** command.

Examples

Options file:

The following are examples of how you might specify the managedservices option in your client system-options file (dsm.sys).

Task   Specify that the client acceptor daemon manages the Web client *only*.

```
managedservices webclient
```

Task   Specify that the client acceptor daemon manages the scheduler *only*.

```
managedservices schedule
```

Task   Specify that the client acceptor daemon manages both the Web client and the scheduler.

```
managedservices schedule webclient
```

Note: The order in which these values are specified is not important.

Task   For Mac OS X, to use the **dsmc schedule** command, specify:

```
managedservices none
```

Command line:

Does not apply.

Related information

"Passwordaccess" on page 380

See "Configuring the client scheduler" on page 59 for instructions to set up the client acceptor daemon to manage the scheduler.

"Sessioninitiation" on page 414

Maxcmdretries

The maxcmdretries option specifies the maximum number of times the client scheduler (on your workstation) attempts to process a scheduled command that fails.

The command retry starts *only if* the client scheduler has not yet backed up a file, never connected to the server, or failed before backing up a file. This option is only used when the scheduler is running.

You must be an authorized user to use the maxcmdretries option.

Your Tivoli Storage Manager administrator can also set this option. If your Tivoli Storage Manager administrator specifies a value for this option, that value overrides what you specify in the client options file *after* your client node successfully contacts the server.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Scheduler** tab, **Maximum command retries** field of the Preferences editor.

Syntax

►►──MAXCMDRetries── ──*maxcmdretries*───────────────────────────────────────►◄

Parameters

maxcmdretries
   Specifies the number of times the client scheduler can attempt to process a scheduled command that fails. The range of values is zero through 9999; the default is 2.

Examples

Options file:
   maxcmdr 4

Command line:
   -maxcmdretries=4

This option is valid only on the initial command line. It is not valid in interactive mode.

Memoryefficientbackup

The memoryefficientbackup option specifies the memory-conserving algorithm to use for processing full file space backups.

One method backs up one directory at a time, using less memory. The other method uses much less memory, but requires more disk space.

Use the memoryefficientbackup option with the **incremental** command when your workstation is memory constrained. You can also use this option as a parameter to the include.fs option in order to select the algorithm Tivoli Storage Manager uses on a per-filespace basis.

Use memoryefficientbackup=*diskcachemethod* for any file space that has too many files for Tivoli Storage Manager to complete the incremental backup with either the default setting, memoryefficientbackup=*no*, or with memoryefficientbackup=*yes*.

The actual amount of disk space required for the disk cache file created by disk cache incremental backups depends on the number of files and directories included in the backup and on the average path length of the files and directories to be backed up. For UNIX and Linux estimate 1 byte per character in the path name. For Mac OS X, estimate 4 bytes per character in the path name. For example, if there are 1 000 000 files and directories to be backed up and the average path length is 200 characters, then the database occupies approximately 200 MB for UNIX and Linux, and 800 MB for Mac OS X clients. Another way to estimate for planning purposes is to multiply the number of files and directories by the length of the longest path to establish a maximum database size.

A second disk cache file is created for the list of migrated files when backing up an HSM managed file system. The combined disk cache files, created by disk cache incremental backups and HSM managed file system backups, can require above 400 MB of disk space for each million files being backed up. The disk cache file can become very large. Large file support must be enabled on the file system that is being used for the disk cache file.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

This option is allowed in dsm.opt and within a server stanza in dsm.sys, but the value in dsm.opt is ignored if it also appears in dsm.sys. You can also place this option on the initial command line. In interactive mode, this option can be used with the **incremental** command. You can set this option on the **Performance Tuning** tab, **Use memory-saving algorithm** check box of the Preferences editor.

Syntax

```
                                   ┌─No──────────────────┐
►►──MEMORYEFficientbackup──────────┼─────────────────────┼──────────────────►◄
                                   ├─Yes─────────────────┤
                                   └─DISKCACHEMethod──────┘
```

Parameters

No Your client node uses the faster, more memory-intensive method when processing incremental backups. This is the default.

Yes

Your client node uses the method that requires less memory when processing incremental backups.

Diskcachemethod

Your client node uses the method that requires much less memory but more disk space when processing incremental backups for full file systems.

Examples

Options file:

```
memoryefficientbackup yes
memoryefficientbackup diskcachem
```

Command line:
```
-memoryef=no
```

Related information

"Include options" on page 351

Mode

Use the mode option with the **backup group**, **backup image**, and **backup nas** to specify whether you want to perform a full or differential group backup containing a list of files from one or more file spaces.

You can use the mode option with the following commands:

backup image

To specify whether to perform a selective or incremental image backup of client file systems.

backup nas

To specify whether to perform a full or differential image backup of NAS file systems.

backup group

To specify whether to perform a full or differential group backup containing a list of files from one or more file space origins.

The mode option has no effect on a raw logical device backup.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, Linux on System z and Solaris *only*. The Tivoli Storage Manager client API does not support this option.

Syntax

For image backup of client file systems (AIX, Solaris, HP-UX, Linux x86/x86_64, and Linux on POWER).

```
           ┌─Selective──┐
►►──MODE = ─┤            ├──────────────────────────────────────────────►◄
           └─Incremental─┘
```

For image backup of NAS file systems (AIX and Solaris).

```
           ┌─differential─┐
►►──MODE = ─┤              ├────────────────────────────────────────────►◄
           └─full─────────┘
```
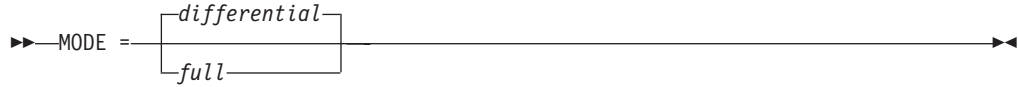
For group backups AIX, Solaris, Linux x86/x86_64.

```
           ┌─full─────────┐
►►──MODE = ─┤              ├────────────────────────────────────────────►◄
           └─differential─┘
```
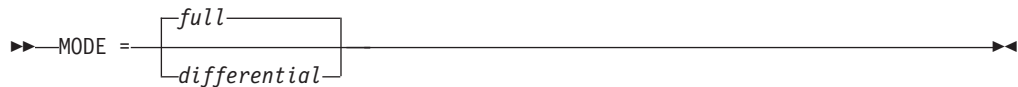
Parameters

selective
> Specifies that you want to perform a full (selective) image backup. This is the default for image backup of client file systems.

incremental
> Specifies that you want to back up only new and changed files after the most recent image backup (full or incremental). If an image backup has not yet been made, then the first backup is a full image backup, regardless of the mode option.
>
> Deleted files are not marked inactive on the server.

full
> Specifies that you want to perform a full backup of NAS or group objects. This is the default for group backups.

differential
> This is the default for NAS objects. Specifies that you want to perform a NAS or group backup of files that changed since the last full backup. If there is no copy of a full image stored on the Tivoli Storage Manager server, a full backup occurs. If a full image exists, whether it is restorable, or expired and being maintained because of dependent differential images, specifying **MODE** =differential sends a differential image backup. If a full image is sent during a differential backup, it is reflected as a full image using the **QUERY NASBACKUP** server command.
>
> A full image can be eligible for expiration based on versioning or retention (verexists retextra), but still be maintained on the Tivoli Storage Manager server to allow for restoring dependent differential images. A full image that is eligible for expiration *cannot* be selected for restore, so it is not displayed using the **QUERY NASBACKUP** server command. The differential image backups that depend on an "expired" full image can be restored.

Examples

Task   Perform the NAS image backup of the entire file system.

> **Command:** `dsmc backup nas -mode=full -nasnodename=nas1 /vol/vol0 /vol/vol1`

Task Back up the /home/test file space using an image incremental backup that backs up only new and changed files after the last full image backup.

> **Command:** `dsmc backup image /home/test -mode=incremental`

Task Perform a full backup of all the files in filelist /home/dir1/filelist1 to the virtual file space name /virtfs containing the group leader /home/group1 file.

> **Command:**
>
> ```
> backup group -filelist=/home/dir1/filelist1 -groupname=group1
> -virtualfsname=/virtfs -mode=full
> ```

Monitor

The monitor option specifies whether to monitor an image backup or restore of file systems belonging to a Network Attached Storage (NAS) file server.

If you specify monitor=*yes*, Tivoli Storage Manager monitors the current NAS image backup or restore operation and displays processing information on your screen. This is the default.

If you specify monitor=*no*, Tivoli Storage Manager does not monitor the current NAS image backup or restore operation and is available to process the next command.

Use this option with the **backup nas** or **restore nas** commands.

Supported Clients

This option is valid for AIX, Linux, and Solaris clients *only*.

Syntax

```
►►──MONitor =─┬─Yes─┬─────────────────────────────────────────────►◄
              └─No──┘
```

Parameters

Yes
　　Specifies that you want to monitor the current NAS image backup or restore operation and display processing information on your screen. This is the default.

No Specifies that you do not want to monitor the current NAS image backup or restore operation.

Examples

Command line:
```
backup nas -mode=full -nasnodename=nas1 -monitor=yes
/vol/vol0 /vol/vol1
```

Nasnodename

The nasnodename option specifies the node name for the NAS file server when processing NAS file systems.

The node name identifies the NAS file server to the Tivoli Storage Manager server. The server must register the NAS file server.

You can specify this option on the command line or in the client system-options file (dsm.sys).

You can override the default value in the dsm.sys file by entering a different value on the command line. If you do not specify the nasnodename option in the dsm.sys file, you *must* specify this option on the command line when processing NAS file systems.

You can use the nasnodename option with the following commands:
- **backup nas**
- **delete filespace**
- **query backup**
- **query filespace**
- **restore nas**

You can use the **delete filespace** command to interactively delete NAS file spaces from server storage.

Use the nasnodename option to identify the NAS file server. When using an interactive command-line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID. Place the nasnodename option in your client system-options file (dsm.sys). The value in the client system-options file is the default, but this value can be overridden on the command line. If the nasnodename option is not specified in the client system-options file, you must specify this option on the command line when processing NAS file systems.

Use the class option to specify the class of the file space to delete. To display a list of file spaces belonging to a NAS node so that you can choose one to delete, use the *-class=nas* option.

To delete NAS file spaces using the Web client, see the topic for backing up your data.

Supported Clients

This option is valid for the AIX, Linux, and Solaris clients *only*. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **General** tab of the Preferences editor.

Syntax

►►──NASNodename── ──*nodename*──►◄

Parameters

nodename
    Specifies the node name for the NAS file server.

Examples

Options file:
    nasnodename nas2

Command line:
    -nasnodename=nas2

Nfstimeout

Authorized user: The nfstimeout option specifies the number of seconds the client waits for a status system call on an NFS file system before it times out.

You can use this option to mitigate the default behavior of status calls on file systems. For example, if an NFS file system is stale, a status system call is timed out by NFS (softmounted) or hang the process (hardmounted).

When the value of this option is changed to a value other than zero, a new thread is created by a caller thread to issue the status system call. The new thread is timed out by the caller thread and the operation can continue.

Note: On Solaris and HP-UX, the nfstimeout option can fail if the NFS mount is hard. If a hang occurs, deactivate the nfstimeout option and mount the NFS file system soft mounted, as follows:

    mount -o soft,timeo=5,retry=5 machine:/filesystem /mountpoint

The parameters are defined as follows:

soft   Generates a soft mount of the NFS file system. If an error occurs, the stat() function returns with an error. If the option hard is used, stat() does not return until the file system is available.

timeo=n
         Sets the time out for a soft mount error to n seconds

retry=n
         Set the internal retries and the mount retries to 5, the default is 10000.

Supported Clients

This option is for all UNIX and Linux clients. The server can also define this option.

Options File

Place this option in the dsm.sys file within a server stanza *or* the client options file (dsm.opt).

Syntax

►►──NFSTIMEout── ──*number*──►◄

Parameters

number
    Specifies the number of seconds the client waits for a status system call on a file system before timing out. The range of values is 0 through 120; the default is 0 seconds.

Chapter 10. Processing options   **373**

Examples

Options file:
```
nfstimeout 10
```

Command line:
```
-nfstimeout=10
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Related information

See "NFS hard and soft mounts" on page 174 for a discussion of how NFS hard and soft mounts are handled.

Nodename

Use the nodename option in your client options file to identify your workstation to the server. You can use different node names to identify multiple operating systems on your workstation.

You must be an authorized user to use the nodename option.

When you use the nodename option, Tivoli Storage Manager prompts for the password assigned to the node you specify, if a password is required.

If you want to restore or retrieve files from the server while you are working from a different workstation, use the virtualnodename option. You can also use the asnodename option, if it is set up by the administrator.

When connecting to a server, the client must identity itself to the server. This login identification is determined in the following manner:

- In the absence of a nodename entry in the dsm.sys file, or a virtualnodename entry in the client user-options file (dsm.opt), or a virtual node name specified on a command line, the default login ID is the name that the **hostname** command returns.
- If a nodename entry exists in the dsm.sys file, the nodename entry overrides the name that the **hostname** command returns.
- If a virtualnodename entry exists in the client system-options file (dsm.sys), or a virtual node name is specified on a command line, it cannot be the same name as the name returned by the **hostname** command. When the server accepts the virtual node name, a password is required (if authentication is on), even if the passwordaccess option is *generate*. When a connection to the server is established, access is permitted to any file that is backed up using this login ID.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **General** tab, **Node Name** field of the Preferences editor.

Syntax

►►──NODename── ──*nodename*──►◄

Parameters

nodename

Specifies a 1 to 64 character node name for which you want to request Tivoli Storage Manager services. The default is the value returned with the **hostname** command.

Not specifying a node name permits the node name to default to the host name of the workstation

Examples

Options file:

    nodename cougar

Nojournal

Use the nojournal option with the **incremental** command to specify that you want to perform a traditional full incremental backup, instead of the default journal-based backup.

Journal-based incremental backup differs from the traditional full incremental backup in the following ways:

- Tivoli Storage Manager does not enforce non-default copy frequencies (other than 0).
- UNIX special file changes are not detected by the Journal daemon and are not, therefore, backed up.

For these reasons, you want to use the nojournal option periodically to perform a traditional full incremental backup.

Supported Clients

This option is valid for the AIX backup-archive client.

Syntax

►►──NOJournal──►◄

Parameters

There are no parameters for this option.

Examples

Command line:

    dsmc incr /home -nojournal

Related information

See "Journal-based backup" on page 496 for more information about journal-based incremental backup.

Noprompt

The noprompt option suppresses the confirmation prompt that is presented by the **delete group**, **delete archive**, **expire**, **restore image**, and **set event** commands.
- **delete archive**
- **delete backup**
- **delete group**
- **expire**
- **restore image**

Note: The **restore image** command does not apply to Mac OS X operating systems.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►──NOPrompt──►◄

Parameters

There are no parameters for this option.

Examples

Command line:
      dsmc delete archive -noprompt "/Users/van/Documents/*"

Command line:
      dsmc delete archive -noprompt "/home/project/*"

Numberformat

The numberformat option specifies the format you want to use to display numbers.

The AIX, Solaris, and HP-UX clients support locales other than English that describe every user interface that varies with location or language.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time the client is called. Consult the documentation on your local system for details about setting up your locale definition.

Note: The numberformat option does not affect the Web client. The Web client uses the number format for the locale that the browser is running in. If the browser is not running in a supported locale, the Web client uses the number format for US English.

You can use the numberformat option with the following commands:

- **delete archive**
- **delete backup**
- **expire**
- **query archive**
- **query backup**
- **query image**
- **query nas**
- **restore**
- **restore image**
- **restore nas**
- **retrieve**
- **set event**

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Regional Settings** tab, **Number Format** field of the Preferences editor.

Syntax

►►──NUMberformat── ──*number*──►◄

Parameters

number

Displays numbers using any one of the following formats. Specify the number (0–6) that corresponds to the number format you want to use.

0 Use the locale-specified date format. This is the default (does not apply to Mac OS X).

1 1,000.00

This is the default for the following available translations:
- US English
- Japanese
- Chinese (Traditional)
- Chinese (Simplified)
- Korean

2 1,000,00

3 1 000,00

This is the default for the following available translations:
- French
- Czech
- Hungarian
- Polish
- Russian

4 1 000.00

5 1.000,00

This is the default for the following available translations:
- Brazilian Portuguese
- German
- Italian
- Spanish

6   1'000,00

For AIX, HP-UX, and Solaris: To define number formats, modify the following lines in the source file of your locale. Whatever format you select applies both to output and to input.

decimal_point
> The character that separates the whole number from its fractional part.

thousands_sep
> The character that separates the hundreds from the thousands from the millions.

grouping
> The number of digits in each group that is separated by the thousands_sep character.

Examples

Options file:
> ```
> num 4
> ```

Command line:
> ```
> -numberformat=4
> ```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

Optfile

The optfile option specifies the client options file to use when you start a Tivoli Storage Manager session.

Supported Clients

This option is valid for all clients.

Syntax

```
►►──OPTFILE =── ──file_name──────────────────────────────────────────────►◄
```

Parameters

file_name
> Specifies an alternate client options file, if you use the fully qualified path name. If you specify only the file name, Tivoli Storage Manager assumes the file name specified is located in the current working directory. The default is dsm.opt.

Restriction: Specify the full path when you use this option with the client acceptor daemon (dsmcad), because the client acceptor daemon changes its working directory to root ("/") after initialization.

Examples

Command line:

```
dsmc query session -optfile=myopts.opt
```

client acceptor daemon:

```
dsmcad -optfile=/usr/tivoli/tsm/client/ba/bin/myopts.opt
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Password

The password option specifies a Tivoli Storage Manager password.

If you do not specify this option and your administrator has set authentication to *On*, you are prompted for a password when you start a Tivoli Storage Manager session.

You must be an authorized user to use the password option.

Note:
1. If the server prompts for a password, the password is not displayed as you enter it. However, if you use the password option on the command line, your password is displayed as you enter it.
2. If the Tivoli Storage Manager server name changes or Tivoli Storage Manager clients are directed to a different Tivoli Storage Manager server, all clients must re-authenticate with the server because the stored encrypted password must be regenerated.

The password option is ignored when the passwordaccess option is set to *generate*.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client user-options file (dsm.opt).

Syntax

```
►►──PASsword── ──password────────────────────────────────────────────►◄
```

Parameters

password

Specifies a 1 to 63 character password. A password is not case-sensitive. Valid characters include:

Characters
      Description

A–Z	Any letter, A through Z, uppercase or lowercase
0–9	Any number, 0 through 9
+	Plus
.	Period
_	Underscore
-	Hyphen
&	Ampersand

Examples

Options file:
```
password secretword
```

Command line:
```
-password=secretword
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Passwordaccess

The passwordaccess option specifies whether you want to generate your password automatically or set as a user prompt.

Your administrator can require a password for your client node by enabling the authentication feature. Ask your administrator if a password is required for your client node.

You must be an authorized user to use the passwordaccess option.

If a password is required, you can choose one of the following methods:
- Set the password for your client node yourself and have Tivoli Storage Manager prompt for it each time you request services.
- Let Tivoli Storage Manager automatically generate a new password for your client node each time it expires, encrypt and store the password in a file, and retrieve the password from that file when you request services. You are not prompted for the password.

When the passwordaccess option is set to *generate* and you specify the password option, the password option is ignored.

Setting the passwordaccess option to *generate* is required in the following situations:
- When using the HSM client.
- When using the Web client.
- When performing NAS operations.
- When using Tivoli Continuous Data Protection for Files.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Authorization** tab, **Password Access** section of the Preferences editor.

Syntax

```
►►──PASSWORDAccess──┬──prompt──┬─────────────────────────────────────►◄
                    └─generate─┘
```

Parameters

prompt

You are prompted for your Tivoli Storage Manager client node password each time a client connects to the server. This is the default.

To keep your client node password secure, enter commands without the password and wait for Tivoli Storage Manager to prompt you for the password.

Each user must know the Tivoli Storage Manager password for your client node. Any user who knows the password for your client node can gain access to *all* backups and archives that originate from your client node. For example:

- If the user enters the node name and password for your client node from a different client node, the user becomes a virtual root user.

API applications must supply the password when a session is initiated. The application is responsible for obtaining the password (does not apply to Mac OS X).

generate

Encrypts and stores your password locally and generates a new password when the old password expires. The new password is randomly generated by the Tivoli Storage Manager client

A password prompt is displayed when registering a workstation with a server using open registration or if your administrator changes your password manually.

When logging in locally, users do not need to know the Tivoli Storage Manager password for the client node. However, by using the nodename option at a remote node, users can access files they own and files to which another user grants access.

Examples

Options file:
        passwordaccess generate

Command line:
        Does not apply.

Related information

For information on where the password is stored, see "Passworddir" on page 382.

Passworddir

The passworddir option specifies the directory location in which to store an encrypted password file.

Authorized user: The default directory for AIX is `/etc/security/adsm` and for other UNIX and Linux platforms it is `/etc/adsm`. Regardless of where it is stored, the password file created by Tivoli Storage Manager is always named `TSM.PWD`.

Supported Clients

This option is valid for all UNIX clients.

Options File

Place this option in the dsm.sys file within a server stanza.

Syntax

```
►►──PASSWORDDIR── ──directoryname───────────────────────────────────►◄
```

Parameters

directoryname
Specifies the path in which to store the encrypted password file. The name of the password file is `TSM.PWD`. If any part of the specified path does not exist, Tivoli Storage Manager attempts to create it.

Examples

Options file:
```
passworddir "/Users/user1/Library/Preferences/Tivoli Storage Manager/"
passworddir /etc/security/tsm
```

Command line:
Does not apply.

Pick

The pick option creates a list of backup versions or archive copies that match the file specification you enter.

From the list, you can select the versions to process. Include the inactive option to view both active and inactive objects.

For images, if you do not specify a source file space and destination file space, the pick list contains all backed up images. In this case, the images selected from the pick list are restored to their original location. If you specify the source file space and the destination file space, you can select only one entry from the pick list.

Use the pick option with the following commands:
- **delete archive**
- **delete backup**
- **delete group**
- **expire**

- **restore**
- **restore group**
- **restore image**
- **restore nas**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──PIck──────────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Command line:
      dsmc restore "/Users/van/Documents/*" -pick -inactive

Command line:
      dsmc restore "/home/project/*" -pick -inactive

Pitdate

Use the pitdate option with the pittime option to establish a point in time to display or restore the latest version of your backups.

Files that were backed up *on or before* the date and time you specify, and which were not deleted *before* the date and time you specify, are processed. Backup versions that you create after this date and time are ignored.

Use the pitdate option with the following commands:
- **delete backup**
- **query backup**
- **query group**
- **query image**
- **restore**
- **restore group**
- **restore image**
- **restore nas**

When pitdate is used, the inactive and latest options are implicit.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──PITDate =── ──date──────────────────────────────────────────────◄◄
```

Parameters

date

Specifies the appropriate date. Enter the date in the format you selected with the dateformat option.

When you include dateformat with a command, it must precede the fromdate, pitdate, and todate options.

Examples

Command line:

```
dsmc restore "/Volumes/proj4/myproj/*" -sub=y -pitdate=08/01/2003
-pittime=06:00:00
```

Command line:

```
dsmc restore "/fs1/*" -sub=y -pitdate=08/01/2003 -pittime=06:00:00
```

Pittime

Use the pittime option with the pitdate option to establish a point in time to display or restore the latest version of your backups.

Files that were backed up *on or before* the date and time you specify, and which were not deleted *before* the date and time you specify, are processed. Backup versions that you create after this date and time are ignored. This option is ignored if you do not specify pitdate option.

Use the pittime option with the following commands:
* **delete backup**
* **query backup**
* **query image**
* **restore**
* **restore image**
* **restore nas**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──PITTime =── ──time──────────────────────────────────────────────◄◄
```

Parameters

time

Specifies a time on a specified date. If you do not specify a time, the time defaults to 23:59:59. Specify the time in the format you selected with the timeformat option.

When you include the timeformat option in a command, it must precede the fromtime, pittime, and tottime options.

Examples

Command line:
```
dsmc query backup -pitt=06:00:00 -pitd=08/01/2003
"/Volumes/proj5/myproj/*"
```

Command line:
```
dsmc q b "/fs1/*" -pitt=06:00:00 -pitd=08/01/2003
```

Postschedulecmd/Postnschedulecmd

The postschedulecmd/postnschedulecmd option specifies a command that the client program processes after it runs a schedule.

The client program waits for the command to complete before it continues with other processing.

You must be an authorized user to use the postschedulecmd option.

If you do not want to wait, specify postnschedulecmd.

Note:

1. If the postschedulecmd command does not complete with return code 0, the client reports that the scheduled event completed with return code 8 (unless the scheduled operation encounters a more severe error yielding a higher return code). If you do not want the postschedulecmd command to be governed by this rule, you can create a script or batch file that invokes the command and exits with return code 0. Then configure postschedulecmd to invoke the script or batch file. The return code for the postnschedulecmd command is not tracked, and does not influence the return code of the scheduled event.

2. The server can also define the postschedulecmd option (and the postnschedulecmd option).

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Scheduler** tab, **Schedule Command** button of the Preferences editor.

Syntax

```
►►──┬─POSTSChedulecmd──┬─ ──"cmdstring"──────────────────────────────────►◄
    └─POSTNSChedulecmd─┘
```

Parameters

"cmdstring"
Specifies the command to process. You can enter a command to be run after a schedule with this option. Use only one postschedulecmd option.

If the command string contains blanks, enclose the command string in quotation marks. If you placed quotation marks within the command string, then enclose the entire command string in single quotation marks.

Use a blank, or null, string for *cmdstring* if you want to prevent any commands from running that the Tivoli Storage Manager server administrator uses for postschedulecmd or preschedulecmd. If you specify a blank or null string on *either* option, it prevents the administrator from using a command on *both* options.

If your administrator uses a blank or null string on the postschedulecmd option, you cannot run a post-schedule command.

For Mac OS X, if the postschedulecmd schedule command is a UNIX shell script, enter the name of the script to run. For example, if *"presched.sh"* is a UNIX shell script, enter this command:

    postschedulecmd "/Volumes/La Pomme/Scripting/presched.sh"

Examples

Options file:
    postschedulecmd "/Volumes/La Pomme/Scripting/postsched.sh"

Options file:
    postschedulecmd "restart database"

The command string is a valid command for restarting your database.

Command line:
    -postschedulecmd="/Volumes/La Pomme/Scripting/postsched.sh"

Command line:
    -postschedulecmd="'restart database'"

This option is valid only on the initial command line. It is not valid in interactive mode.

Postsnapshotcmd

The postsnapshotcmd option allows you to run operating system shell commands or scripts after the Tivoli Storage Manager client starts a snapshot during a snapshot-based backup operation.

AIX only: This option is only applicable to JFS2 snapshot-based file backup or archive and snapshot-based image backup. For a snapshot-based file backup or archive, use this option with the **backup** command, the include.fs option, or in the dsm.sys file.

Linux only: This option is only valid if the LVM is installed and configured on your system, allowing you to perform a snapshot-based image backup operation.

AIX and Linux only: For a snapshot-based image backup, use this option with the **backup image** command, the include.image option, or in the dsm.sys file.

If the postsnapshotcmd fails the operation continues, but appropriate warnings are logged.

Supported Clients

This option is valid for AIX clients and Linux x86/x86_64 clients only. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can also set this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

```
►►──POSTSNAPshotcmd── ──"cmdstring"──────────────────────────────────────────►◄
```

Parameters

"cmdstring"
>   Specifies a command to process.
>
>   Use the srvprepostsnapdisabled option to prevent the Tivoli Storage Manager server administrator from executing operating system commands on the client system.
>
>   If the command string contains blanks, enclose the command string in quotation marks:
>
>   ```
>   "resume database myDb"
>   ```
>
>   If you placed quotation marks within the command string, then enclose the entire command string in single quotation marks:
>
>   ```
>   'resume database "myDb"'
>   ```

Examples

Options file:
>   ```
>   postsnapshotcmd "any command"
>   ```
>
>   The command string is a valid command for restarting your application.

Command line:
>   ```
>   backup image -postsnapshotcmd="any command"
>   ```

This option is valid only on the initial command line. It is not valid in interactive mode.

Related information

"Include options" on page 351

"Srvprepostscheddisabled" on page 428

Preschedulecmd/Prenschedulecmd

The preschedulecmd option specifies a command that the client program processes before it runs a schedule.

The client program waits for the command to complete before it starts the schedule. If you do not want it to wait, specify prenschedulecmd.

You must be an authorized user to use this option.

Note:

1. Successful completion of the preschedulecmd command is considered to be a prerequisite to running the scheduled operation. If the preschedulecmd command does not complete with return code 0, the scheduled operation and any postschedulecmd and postnschedulecmd commands will not run. The client reports that the scheduled event failed, and the return code is 12. If you do not want the preschedulecmd command to be governed by this rule, you can create a script or batch file that invokes the command and exits with return code 0. Then configure preschedulecmd to invoke the script or batch file. The return code for the prenschedulecmd command is not tracked, and does not influence the return code of the scheduled event.

2. The server can also define the preschedulecmd option (and the prenschedulecmd option).

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Scheduler** tab, **Schedule Command** button of the Preferences editor.

Syntax

```
►►──┬─PRESChedulecmd──┬──── "cmdstring" ──────────────────────────────►◄
    └─PRENSChedulecmd─┘
```
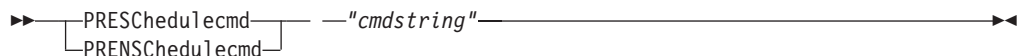
Parameters

"cmdstring"
Specifies the command to process. Use only one preschedulecmd option. You can enter a command to be executed before a schedule using this option.

If the command string contains blanks, enclose the command string in quotation marks. If you placed quotation marks within the command string, then enclose the entire command string in single quotation marks.

Use a blank or null string for *cmdstring* if you want to prevent any commands from running that the Tivoli Storage Manager server administrator uses for postschedulecmd and preschedulecmd. If you specify a blank or null string on *either* option, it prevents the administrator from using a command on *both* options.

If your administrator uses a blank or null string on the preschedulecmd option, you cannot run a pre-schedule command.

For Mac OS X, if the preschedulecmd schedule command is an AppleScript, you must use the osascript command to run the script. For example, if *"Database Script"* is an apple script, enter this command:

```
preschedulecmd osascript "/Volumes/La Pomme/Scripting/
  Database Script"
```

Examples

Options file:
```
preschedulecmd "<the quiesce command of your database product>
database"
```

The command string is a valid command for quiescing your database.

Command line:
```
-preschedulecmd="'quiesce database'"
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Preservelastaccessdate

Use the preservelastaccessdate option during a backup or archive operation to specify whether to reset the last access date of any specified files to their original value following the backup or archive operation.

By default, the Tivoli Storage Manager client *will not* reset the last access date of any backed up or archived files to their original value following the backup or archive operation. This option requires extra processing time during backup and archive for each file that is sent to the Tivoli Storage Manager server, so it should only be enabled when necessary.

If you have enabled open file support, the last access date for files is always preserved regardless of the setting for preservelastaccessdate. When using open file support, leave this option disabled.

Use this option with the **incremental**, **selective**, or **archive** commands.

Note:
1. This option only applies to files; it does not apply to directories.
2. Resetting the last access date incurs additional overhead that can impact backup and archive performance. The last access date should be reset only if you are using another application, such as a Storage Resource Management (SRM) or Hierarchical Storage Management that relies on accurate last access dates.
3. You cannot reset the last access date of read-only files. The preservelastaccessdate option ignores read-only files and does not change their date.

Supported Clients

This option is valid for all UNIX and Linux clients except Mac OS X.

The server can also define this option.

Options File

Place this option in the client user options file (dsm.opt). You can set this option on the **Backup** tab of the Preferences editor.

Syntax

```
►►──PRESERVELAstaccessdate──┬──No──┬─────────────────────────────────────────────►◄
                            └──Yes─┘
```

Parameters

No Specifies that the Tivoli Storage Manager client *will not* reset the last access date of any backed up or archived files to their original value following the backup or archive operation. This is the default.

Yes

Specifies that the Tivoli Storage Manager client *will* reset the last access date of any backed up or archived files to their original value following the backup or archive operation.

Examples

Options file:
```
preservelastaccessdate yes
```

Command line:
```
Incremental /proj/test/test_file -preservelastaccessdate=yes
```

Preservepath

The preservepath option specifies how much of the source path to reproduce as part of the target directory path when you restore or retrieve files to a new location.

Use the -subdir=*yes* option to include the entire subtree of the source directory (directories and files below the lowest-level source directory) as source to be restored. If a required target directory does not exist, it is created. If a target file has the same name as a source file, it is overwritten. Use the -replace=*prompt* option to have Tivoli Storage Manager prompt you before files are overwritten.

Use the preservepath option with the following commands:
- **restore**
- **restore backupset**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Syntax

```
►►──PRESERvepath =──┬──Subtree───┬──────────────────────────────────────────────►◄
                    ├──Complete──┤
                    ├──NOBase────┤
                    └──NONe──────┘
```

Parameters

Subtree
Creates the lowest-level source directory as a subdirectory of the target directory. Files from the source directory are stored in the new subdirectory. This is the default.

Complete
Restores the entire path, starting from the root, into the specified directory. The entire path includes all the directories *except* the file space name.

NOBase
Restores the contents of the source directory without the lowest level, or base directory, into the specified destination directory.

NONe
Restores all selected source files to the target directory. No part of the source path at or above the source directory is reproduced at the target.

If you specify SUBDIR=*yes*, Tivoli Storage Manager restores all files in the source directories to the single target directory.

Examples

Command line:
For the examples below, assume that the server file space contains the following backup copies:
```
/fs/h1/m1/file.a
/fs/h1/m1/file.b
/fs/h1/m1/l1/file.x
/fs/h1/m1/l1/file.y
```

This command:
```
dsmc res /fs/h1/m1/ /u/ann/ -preser=complete
```

Restores these directories and files:
```
/u/ann/h1/m1/file.a
/u/ann/h1/m1/file.b
```

This command:
```
dsmc res /fs/h1/m1/ /u/ann/ -preser=nobase
```

Restores these directories and files:
```
/u/ann/file.a
/u/ann/file.b
```

This command:
```
dsmc res backupset /fs/h1/m1/ /u/ann/ -su=yes
-preser=nobase -loc=file
```

Restores these directories and files:
```
/u/ann/file.a
/u/ann/file.b
/u/ann/file.x
/u/ann/file.y
```

This command:
```
dsmc res /fs/h1/m1/ /u/ann/ -preser=subtree
```

Restores these directories and files:

```
/u/ann/m1/file.a
/u/ann/m1/file.b
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -preser=none
```

Restores these directories and files:

```
/u/ann/file.a
/u/ann/file.b
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -su=yes -preser=complete
```

Restores these directories and files:

```
/u/ann/h1/m1/file.a
/u/ann/h1/m1/file.b
/u/ann/h1/m1/l1/file.x
/u/ann/h1/m1/l1/file.y
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -su=yes -preser=nobase
```

Restores these directories and files:

```
/u/ann/file.a
/u/ann/file.b
/u/ann/l1/file.x
/u/ann/l1/file.y
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -su=yes -preser=subtree
```

Restores these directories and files:

```
/u/ann/m1/file.a
/u/ann/m1/file.b
/u/ann/m1/l1/file.x
/u/ann/m1/l1/file.y
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -su=yes -preser=none
```

Restores these directories and files:

```
/u/ann/file.a
/u/ann/file.b
/u/ann/file.x
/u/ann/file.y
```

Presnapshotcmd

The presnapshotcmd option allows you to run operating system commands before the Tivoli Storage Manager client starts a snapshot.

This allows you to quiesce an application before the Tivoli Storage Manager client starts the snapshot during a snapshot-based backup or archive.

AIX only: This option is only applicable to JFS2 snapshot-based file backup or archive and snapshot-based image backup. For a snapshot-based file backup or archive, use this option with the **backup** command, the include.fs option, or in the dsm.sys file.

Linux only: This option is only valid if the LVM is installed and configured on your system, allowing you to perform a snapshot-based image backup.

AIX and Linux only: For a snapshot-based image backup, use this option with the **backup image** command, the include.image option, or in the dsm.sys file.

If the presnapshotcmd fails it is assumed that the application is not in a consistent state and the client stops the operation and display the appropriate error message.

Supported Clients

This option is valid for AIX JFS2 and Linux x86/x86_64 clients *only*. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set also this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

```
►►──PRESNAPshotcmd── ──"cmdstring"──────────────────────────────────►◄
```

Parameters

"cmdstring"
> Specifies a command to process.
>
> Use the srvprepostsnapdisabled option to prevent the Tivoli Storage Manager server administrator from running operating system commands on the client system.
>
> If the command string contains blanks, enclose the command string in quotation marks:
>
> > "quiesce database myDb"
>
> If you placed quotation marks within the command string, then enclose the entire command string in single quotation marks:
>
> > 'resume database "myDb"'

Examples

Options file:

> presnapshotcmd "any shell command or script"

Command line:
> backup image -presnapshotcmd="any shell command or script"

This option is valid only on the initial command line. It is not valid in interactive mode.

Related information

"Include options" on page 351

Queryschedperiod

The queryschedperiod option specifies the number of hours you want the client scheduler to wait between attempts to contact the server for scheduled work.

You must be an authorized user to use the queryschedperiod option.

This option applies only when you set the schedmode option to *polling*. This option is used only when the scheduler is running.

Your administrator can also set this option. If your administrator specifies a value for this option, that value overrides the value set in your client options file after your client node successfully contacts the server.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Scheduler** tab, **Query Schedule Interval** field of the Preferences editor.

Syntax

►►──QUERYSCHedperiod── ──*hours*──►◄

Parameters

hours
> Specifies the number of hours the client scheduler waits between attempts to contact the server for scheduled work. The range of values is 1 through 9999; the default is 12.

Examples

Options file:
> querysch 6

Command line:
> -queryschedperiod=8

This option is valid only on the initial command line. It is not valid in interactive mode.

Querysummary

The querysummary option provides statistics about files, directories and objects that are returned by the **query backup** or **query archive** commands.

The following statistics are provided by the querysummary option:
- The aggregate number of files and directories that are returned by the query backup or query archive command

- The aggregate amount of data of the objects that are returned by the query backup or query archive command
- The classic restore memory-utilization estimate to restore objects that are returned by the query backup or query archive command
- The total number of unique server volumes where the objects that are returned by the query command reside

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──QUERYSUMMARY────────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Command line:
```
        dsmc q ba  '/usr/fs1/*'   -sub=yes  -querysummary
```

```
[root@kaveri:/home/cpark] $ dsmc q ba '/kalafs1/*' -sub=yes -querysummary
IBM Tivoli Storage Manager
Command Line Backup-Archive Client Interface
  Client Version 6, Release 1, Level 0.0 0804FB
  Client date/time: 08/19/08   08:09:48
(c) Copyright by IBM Corporation and other(s) 1990, 2008. All Rights Reserved.

Node Name: KAVERI
Session established with server TEMPLAR: AIX-RS/6000
  Server Version 5, Release 4, Level 3.0
  Server date/time: 08/19/08   08:09:02  Last access: 08/19/08   07:48:59

      Size      Backup Date          Mgmt Class      A/I File
      ----      -----------          ----------      --- ----
     4,096  B  08/07/08   12:07:30      BASVT2        A  /kalafs1/
       256  B  08/07/08   12:07:30      BASVT2        A  /kalafs1/dir1
10,485,760  B  08/07/08   12:07:30      DEFAULT       A  /kalafs1/info1
 5,242,880  B  08/07/08   12:07:30      DEFAULT       A  /kalafs1/info2
     1,044  B  08/07/08   12:07:30      DEFAULT       A  /kalafs1/dir1/subfile1
     1,044  B  08/07/08   12:07:30      DEFAULT       A  /kalafs1/dir1/subfile2


Summary Statistics

Total Files Total Dirs  Avg. File Size Total Data  Memory Est.
----------- ----------  -------------- ----------  ----------
          4          2     3.75 MB       15.00 MB    1.07 KB


Estimated Number of Volumes: 2

[root@kaveri:/home/cpark] $
```

Quiet

The quiet option limits the number of messages that are displayed on your screen during processing..

For example, when you run the **incremental**, **selective**, or **archive** commands, information might appear about each file that is backed up. Use the quiet option if you do not want to display this information

When you use the quiet option, error and processing information appears on your screen, and messages are written to log files. If you do not specify quiet, the default option, verbose is used.

Supported Clients

This option is valid for all clients. The server can also define the quiet option, overriding the client setting. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Command Line** tab, **Do not display process information on screen** checkbox of the Preferences editor.

Syntax

```
►►──QUIET──────────────────────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Options file:
```
quiet
```

Command line:
```
-quiet
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Removeoperandlimit

The removeoperandlimit option specifies that Tivoli Storage Manager removes the 20-operand limit.

If you specify the removeoperandlimit option with the **incremental**, **selective**, or **archive** commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits.

The removeoperandlimit option can be useful if you generate scripts which can invoke the command-line client with a large number of operands. For example, you can prescan a directory tree looking for files to back up. As each *eligible* file is

discovered, it is added to the operand list of a **selective** command. Later, this **selective** command is submitted by a controlling script. In this case, specifying the removeoperandlimit option removes the 20-operand limit.

Note:

1. The removeoperandlimit option *must* be placed immediately after the **incremental**, **selective**, or **archive** command before any file specifications.

2. This option does not accept a value. If this option is specified on a command, the 20-operand limit is removed.

3. Because it adversely affects performance to allow the shell to expand wild cards, use the removeoperandlimit option in backup or archive operations in which wild cards are not used.

4. The removeoperandlimit option is valid only on the **incremental**, **selective**, or **archive** commands in batch mode. It is not valid in the client options file (dsm.opt) or dsm.sys file.

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

```
►►──REMOVEOPerandlimit──────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Command line:
```
-removeoperandlimit
```

Replace

The replace option specifies whether to overwrite existing files on your workstation, or to prompt you for your selection when you restore or retrieve files.

Important: The replace option does not affect recovery of directory objects. Directory objects are always recovered, even when specifying replace=*no*. To prevent overwriting existing directories, use the filesonly option.

You can use this option with the following commands:
- **restore**
- **restore backupset**
- **restore group**
- **retrieve**

Note: Replace prompting does not occur during a scheduled operation. If you set the replace option to prompt, Tivoli Storage Manager skips files without prompting during a scheduled operation.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Restore** tab, **Action for files that already exist** section of the Preferences editor.

Syntax

```
►►──REPlace──┬─Prompt─┬──────────────────────────────────────────────►◄
             ├─All────┤
             ├─Yes────┤
             └─No─────┘
```

Parameters

Prompt
> For nonscheduled operations, you specify whether to overwrite existing files. For scheduled operations, existing files are not overwritten and no prompts are displayed. This is the default.

All
> All existing files are overwritten, including read-only files. If access to a file is denied, you are prompted to skip or overwrite the file. No action is taken on the file until there is a response to the prompt.

Yes
> Existing files are overwritten, *except* read-only files. For nonscheduled operations, you specify whether to overwrite existing read-only files. For scheduled operations, existing read-only files are not overwritten and no prompts are displayed. If access to a file is denied, the file is skipped.

No Existing files are not overwritten. No prompts are displayed.

Examples

Options file:
> ```
> replace all
> ```

Command line:
> ```
> -replace=no
> ```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

Resourceutilization

Use the resourceutilization option in your option file to regulate the level of resources the Tivoli Storage Manager server and client can use during processing.

You must be an authorized user to use the resourceutilization option.

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **General** tab, **Resource Utilization** field of the Preferences editor.

Syntax

```
►►──RESOURceutilization── ──number───────────────────────────────────►◄
```

Parameters

number
> Specifies the level of resources the Tivoli Storage Manager server and client can use during processing. The range of values that you can specify is 1 through 10.

Examples

Options file:
```
        resourceutilization 7
```

Command line:
```
        -resourceutilization=7
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Regulating backup and archive sessions

When you request a backup or archive, the client can use more than one session to the server.

The default is to use a maximum of two sessions; one to query the server and one to send file data. The client can use only one server session if you specify a resourceutilization setting of 1.

When you request a backup or archive, the client can use more than one session to the server. The default is to use a maximum of two sessions; one to query the server and one to send file data. The client can use only one server session if you specify a resourceutilization setting of 1. The client is also restricted to a single session if a user who is not an authorized user invokes a UNIX or Linux client with passwordaccess=*generate* specified.

A client can use more than the default number of sessions when connecting to a server that is Version 3.7 or higher. For example, resourceutilization=10 permits up to eight sessions with the server. Multiple sessions can be used for querying the server and sending file data.

Multiple query sessions are used when you specify multiple file specifications with a backup or archive command. For example, if you enter the following commands and you specify resourceutilization=5, the client might start a second session to query files on file space B.

```
inc /Volumes/filespaceA /Volumes/filespaceB
```

Whether or not the second session starts depends on how long it takes to query the server about files backed up on file space A. The client might also try to read data from the file system and send it to the server on multiple sessions.

Note: During a backup operation, if you enter multiple file specifications, the result might be that files from one file specification are stored on multiple tapes and interspersed with files from different file specifications. This can decrease restore performance. Setting the collocatebyfilespec option to *yes* eliminates interspersing of files from different file specifications, by limiting the client to one server session per file specification. Therefore, if you store the data to tape, files for each file specification are stored together on one tape (unless another tape is required for more capacity).

Related reference

"Collocatebyfilespec" on page 280

Regulating restore sessions

When you request a restore, the default is to use a maximum of one session.

Additional restore sessions are based on:
- resourceutilization value
- how many tapes on which the requested data is stored
- how many tape drives are available
- the maximum number of mount points allowed for the node

Note:

1. If all of the files are on disk, only one session is used. There is no multi-session for a pure disk storage pool restore. However, if you are performing a restore in which the files reside on 4 tapes and some on disk, you could use up to 5 sessions during the restore.

2. The Tivoli Storage Manager server can set the maximum number of mount points a node can use on the server using the MAXNUMMP parameter. If the resourceutilization option value exceeds the value of the MAXNUMMP on the server for a node, the backup can fail with an *Unknown System Error* message.

3. You can get a multi-session restore from your single **restore** command, and from a single volume on the server, if that volume is device class FILE.

For example, if the data you want to restore is on 5 different tape volumes, the maximum number of mount points is 5 for your node, and resourceutilization is set to 3, then 3 sessions are used for the restore. If you increase the resourceutilization setting to 5, then 5 sessions are used for the restore. There is a 1 to 1 relationship to the number of restore sessions allowed for the resourceutilization setting. Multiple restore sessions are only allowed for no query restore operations.

Multiple client session considerations

This topic lists some items to consider when working with multiple client sessions.

The following factors can affect the throughput of multiple sessions:
- The ability of the server to handle multiple client sessions. Is there sufficient memory, multiple storage volumes, and processor cycles to increase backup throughput?

- The ability of the client to drive multiple sessions (sufficient processor cycles, memory, etc.).
- The configuration of the client storage subsystem. File systems that are striped across multiple disks, using either software striping or RAID-5 can better handle an increase in random read requests than a single drive file system. Additionally, a single drive file system might not see performance improvement if it attempts to handle many random concurrent read requests.
- Sufficient bandwidth in the network to support the increased traffic.

Potentially undesirable aspects of running multiple sessions include:
- The client could produce multiple accounting records.
- The server might not start enough concurrent sessions. To avoid this, the server *maxsessions* parameter must be reviewed and possibly changed.
- A query node command might not summarize client activity.

Retryperiod

The retryperiod option specifies the number of minutes the client scheduler waits between attempts to process a scheduled command that fails, or between unsuccessful attempts to report results to the server. Use this option only when the scheduler is running.

You must be an authorized user to use the retryperiod option.

Your administrator can also set this option. If your administrator specifies a value for this option, that value overrides the value in your client system options file after your client node successfully contacts the server.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Scheduler** tab, **Retry period** field of the Preferences editor.

Syntax

```
►►──RETRYPeriod── ──minutes──────────────────────────────────────────►◄
```

Parameters

minutes
Specifies the number of minutes the client scheduler waits between attempts to contact the server, or to process a scheduled command that fails. The range of values is 1 through 9999; the default is 20.

Examples

Options file:
```
retryp 10
```

Command line:
```
-retryperiod=10
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Revokeremoteaccess

The revokeremoteaccess option restricts an administrator with client access privilege from accessing a client workstation that is running the Web client.

This option does not restrict administrators with client owner, system, or policy privilege from accessing your workstation through the Web client.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Web Client** tab of the Preferences editor.

Syntax

```
►►──REVOKEremoteaccess──┬──None──┬──────────────────────────►◄
                        └─Access─┘
```

Parameters

None
    Does not revoke access to administrators who have client access authority for the client. This is the default.

Access
    Revokes access to administrators who have client access authority for the client.

Examples

Options file:
```
revokeremoteaccess none
```

Command line:
    Does not apply.

Schedcmddisabled

The schedcmddisabled option specifies whether to disable the scheduling of commands by the server action=*command* option on the **define schedule** server command.

You must be an authorized user to use the schedcmddisabled option.

This option does not disable the preschedulecmd and postschedulecmd commands. However, you can specify preschedulecmd or postschedulecmd with a blank or a null string to disable the scheduling of these commands.

You can disable the scheduling of commands defined by your Tivoli Storage Manager administrator by setting the schedcmddisabled option to *yes*.

Use the **query schedule** command to query the schedules defined by your administrator.

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file *within* a server stanza.

Syntax

```
                      ┌─No──┐
►►──SCHEDCMDDisabled──┼─────┼──────────────────────────────────────────────────►◄
                      └─Yes─┘
```

Parameters

Yes
    Specifies that Tivoli Storage Manager disables the scheduling of commands by the server using the action=*command* option on the DEFINE SCHEDULE server command.

No Specifies that Tivoli Storage Manager does not disable the scheduling of commands by the server using the action=*command* option on the DEFINE SCHEDULE server command. This is the default.

Examples

Options file:
    schedcmddisabled no

Command line:
    Does not apply.

Related information

"Query Schedule" on page 521

Schedcmdexception

The schedcmexception option is used in conjunction with the schedcmddisabled option to disable the scheduling of commands by the server action=*command* option on the DEFINE SCHEDULE server command, except for specific command strings.

You must be an authorized user to use the schedcmddisabled option.

You must specify the exact string that matches the "objects" definition in the schedule for the scheduled server command to be accepted. If the string does not match exactly (for example, there is an extra space or the capitalization is different), the scheduled command action is blocked.

You can provide multiple schedcmexception options in the options file. This option is not honored if schedcmddisabled is not enabled. The placement of this option in the options file is independent of the placement of the schedcmddisabled option.

Supported Clients

This option is valid for all clients. This option is not valid in the Tivoli Storage Manager Server client options set.

Options File

Place this option in the dsm.sys file *within* a server stanza.

Syntax

```
►►──SCHEDCMDException──string──────────────────────────────────────────►◄
```

Parameters

string
    For commands scheduled by the action=*command* option on the DEFINE SCHEDULE server command, this parameter indicates the objects pattern to enable if the schedcmddisabled=*yes* option is specified. This parameter is case sensitive, and must match the command string on the Tivoli Storage Manager Server schedule definition exactly.

Examples

Options file:
    schedcmddisabled yes

    schedcmdexception "start echo hello, world!"

Examples:
    Disable the scheduling of commands by the server action=*command* option on the DEFINE SCHEDULE server command, except for two schedules defined on the Tivoli Storage Manager Server:

Related information

"Schedcmddisabled" on page 402

Schedlogmax

The schedlogmax option specifies the maximum size of the schedule log and Web client log, in megabytes.

Log records are added to the end of the file, until the maximum specified size is reached. When the maximum specified size is reached, a log record saying "Continued at beginning of file" is placed as the last record in the file. Subsequent logging is resumed at the top of the file. The end of the wrapped log is indicated

by a record saying "END OF DATA". Log messages that are overwritten by wrapping are not saved in a prune file, as they are with the pruning method of log size management.

You must be an authorized user to use the schedlogmax option.

If you change from schedlogmax to schedlogretention, all existing log entries are retained and the log is pruned using the new schedlogretention criteria.

If you change from schedlogretention to schedlogmax, all records in the existing log are copied to the file of pruned entries (`dsmsched.pru` or `dsmwebcl.pru`), the existing log is emptied, and logging begins under the new log wrapping criteria.

If you change the value of the schedlogmax option, the existing log is extended or shortened to accommodate the new size. If the value is reduced, the oldest entries are deleted to reduce the file to the new size.

Restriction: You cannot specify a non-zero schedlogmax value *and* enable schedlogretention.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Scheduler** tab, **Schedule Log** button of the Preferences editor.

Syntax

```
▶▶──SCHEDLOGMAX── ──size──────────────────────────────────────────────────▶◀
```

Parameters

size
> Specifies the maximum size, in megabytes, for the log file. The range of values is 0 to 2047; the default is 0, which disables log file wrapping and allows the log file to grow indefinitely.

Examples

Options file:
> ```
> schedlogmax 100
> ```

Command line:
> ```
> -schedlogmax=100
> ```

This option is valid only on the initial command line. It is not valid in interactive mode.

Schedlogname

The schedlogname option specifies the path and file name where you want to store schedule log information.

Use this option only when you want to store schedule log information. This option applies only when the scheduler is running.

If this option is not used, the `dsmsched.log` file is created in the same directory as the `dsmerror.log` file.

You must be an authorized user to use the schedlogname option.

When you run the **schedule** command, output from scheduled commands appears on your screen. Output is also sent to the file you specify with this option. If any part of the path you specify does not exist, Tivoli Storage Manager attempts to create it

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Scheduler** tab, **Schedule Log** button of the Preferences editor.

Important: Set the `DSM_LOG` environment variable to name a directory where the log is to be placed. The directory specified must have permissions which allow write access from the account under which the client is run. The root directory is *not* a valid value for `DSM_LOG`.

Syntax

```
►►──SCHEDLOGName─── ─filespec──────────────────────────────────────────────►◄
```

Parameters

filespec
> Specifies the path and file name where you want to store schedule log information when processing scheduled work. If any part of the path you specify does not exist, Tivoli Storage Manager attempts to create it.
>
> If you specify a file name only, the file is stored in your current directory. The default is the current working directory with a file name of `dsmsched.log`. The `dsmsched.log` file *cannot* be a symbolic link.
>
> For Mac OS X, if you specify a file name only, the file is stored in your default folder. The default directories are:
> ```
> ~/Library/Logs/tivoli/tsm
> /Library/Logs/tivoli/tsm
> ```

Examples

Options file:
```
SCHEDLOGN /Users/user1/Library/Logs/schedlog.jan
schedlogname /home/mydir/schedlog.jan
```

Command line:
```
-schedlogname=/Users/user1/Library/Logs/schedlog.jan
```

Command line:

```
-schedlogname=/home/mydir/schedlog.jan
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Related information

See "Errorlogname" on page 318 for more information on placement of the `dsmsched.log` file.

Schedlogretention

The schedlogretention option specifies the number of days to keep entries in the schedule log and the Web client log, and whether to save the pruned entries.

The schedule log is pruned during the initial start of scheduler and after a scheduled event completes.

The Web client log is pruned during the initial start of scheduler.

You must be an authorized user to use the schedlogretention option.

If you change from schedlogretention to schedlogmax, all records in the existing log are copied to the pruned log dsmsched.pru, the existing log is emptied, and logging begins under the new log wrapping criteria.

If you change from schedlogmax to schedlogretention, all existing log entries are retained and the log is pruned using the new schedlogretention criteria.

Restriction: You cannot specify a non-zero schedlogmax value *and* enable schedlogretention.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Scheduler** tab, **Schedule Log** button of the Preferences editor.

Syntax

```
►►──SCHEDLOGRetention─┬──N───┬──┬──D──┬──────────────────────►◄
                      └─days─┘  └──S──┘
```

Parameters

N **or** *days*
    Specifies how long to wait before pruning the log.

    N   Do not prune the log. This permits the log to grow indefinitely. This is the default.

days
> Specifies the number of days to keep log file entries before pruning. The range of values is zero through 9999.

D **or** *S*
> Specifies whether to save the pruned entries. Use a space or comma to separate this parameter from the previous one.

> *D* Discards the log entries when pruning the log. This is the default.

> *S* Saves the log entries when pruning the log.

> Pruned entries are copied to the file of pruned entries (`dsmsched.pru` or `dsmsched.pru`), which is stored in the same directory as the log.

Examples

Options file:
> schedlogretention 30 S

Command line:
> -schedlogretention=30,S

This option is valid only on the initial command line. It is not valid in interactive mode.

Schedmode

The schedmode option specifies whether you want to use the *polling* mode (your client node periodically queries the server for scheduled work), or the *prompted* mode (the server contacts your client node when it is time to start a scheduled operation).

All communication methods can use the client polling mode, but only TCP/IP can use the server prompted mode.

You must be an authorized user to use the *polling* mode.

This option applies *only if* you are using the TCP/IP communication method, and the **schedule** command is running.

Your administrator can specify that the server support both modes or just one mode. If your administrator specifies that both modes are supported, you can select either schedule mode. If your administrator specifies only one mode, you must specify that mode in your dsm.sys file or scheduled work is not processed.

If you specify *prompted* mode, you should consider supplying values for the tcpclientaddress and tcpclientport options in your dsm.sys file or on the schedule command; the client can then be contacted at either an address or a port of your choice (useful for client systems with multiple network interface cards).

Note:
1. When changing the setting of this option in the dsm.sys file you must stop and restart the scheduler service for the setting to take effect.
2. The server can also define this option.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Scheduler** tab, **Schedule Mode** section of the Preferences editor.

Syntax

```
►►──SCHEDMODe───┬──POlling──┬──────────────────────────────────►◄
                └──PRompted──┘
```

Parameters

POlling
> The client scheduler queries the server for scheduled work at prescribed time intervals. This is the default. You can set the time intervals using the queryschedperiod option.

PRompted
> The client scheduler waits for the server to contact your client node when scheduled work needs to be done.

> **Note:**
> 1. If you use the **dsmc schedule** command and both schedmode *prompted* and commmethod *V6Tcpip* are specified, the client and Tivoli Storage Manager server must be configured for IPv6. Additionally, the client host name must be set up for the IPv6 address.

Examples

Options file:
>   schedmode prompted

Command line:
>   -schedmod=po

This option is valid only on the initial command line. It is not valid in interactive mode.

Related reference

"Tcpclientaddress" on page 437

"Tcpclientport" on page 438

Schedrestretrdisabled

The schedrestretrdisabled option specifies whether to disable the execution of restore or retrieve schedule operations.

You must be an authorized user to use the schedrestretrdisabled option.

Supported Clients

This option is valid for all clients. The server cannot define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza for the scheduler. You can set this option on the **Scheduler** tab, **Schedule Command** button of the Preferences editor.

Syntax

```
►►──SCHEDRESTRETRDisabled──┬──No──┬──────────────────────────────────►◄
                           └──Yes──┘
```

Parameters

No Specifies that Tivoli Storage Manager does not disable the execution of restore and retrieve schedule operations. This is the default.

Yes
    Specifies that Tivoli Storage Manager disable the execution of restore and retrieve schedule operations.

Examples

Options file:
    schedrestretrdisabled yes

Command line:
    Does not apply.

Scrolllines

The scrolllines option specifies the number of lines of information that are displayed on your screen at one time.

Use this option when you set the scrollprompt option to *Yes*.

You can use the scrolllines option with the following commands only:
* **delete filespace**
* **query archive**
* **query backup**
* **query backupset**
* **query filespace**
* **query group**
* **query image**
* **query nas**
* **query node**
* **query options**

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Command Line** tab, **Number of lines to display on screen** field of the Preferences editor.

Syntax

►►──SCROLLLines── ──*number*──►◄

Parameters

number
> Specifies the number of lines of information that are displayed on your screen at one time. The range of values is 1 through 80; the default is 20.

Examples

Options file:
> scrolllines 25

Command line:
> -scroll=25

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

Scrollprompt

The scrollprompt option specifies whether you want Tivoli Storage Manager to stop and wait after displaying the number of lines of information you specified with the scrolllines option, or scroll through and stop at the end of the information list.

You can use the scrollprompt option with the following commands only:
- **delete filespace**
- **query archive**
- **query backup**
- **query backupset**
- **query filespace**
- **query group**
- **query image**
- **query nas**
- **query node**
- **query options**

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Command Line** tab, **Pause after displaying the following number of lines** field of the Preferences editor.

Syntax

```
►►──SCROLLPrompt──┬──No──┬──────────────────────────────────────►◄
                  └─Yes─┘
```

Parameters

No Scrolls to the end of the list and stops. This is the default.

Yes

Stops and waits after displaying the number of lines you specified with the scrolllines option. The following prompt is displayed at the bottom of the screen:

```
Press 'Q' to quit, 'C' to continuous scroll, or 'Enter' to
continue.
```

Examples

Options file:
```
scrollprompt yes
```

Command line:
```
-scrollp=yes
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

Servername

In your dsm.sys file, the servername option specifies the name you want to use to identify a server and to begin a stanza containing options for that server. You can name and specify options for more than one server.

The following example demonstrates how to specify options for two different servers:

```
SErvername      server_a
  COMMMethod              TCPip
  TCPPort                 1500
  TCPServeraddress        server_hostname2.domain.company.com
  PASSWORDAccess          prompt
  GRoups                  tsm
  USERs                   sullivan  mushock  tallan
  INCLExcl                /adm/tsm/backup.excl

SErvername      server_b
  COMMMethod              SHAREdmem
  shmport                 1520
  PASSWORDAccess          generate
  GRoups                  system tsm
  INCLExcl                /adm/tsm/archive.excl
```

In your client user-options file (dsm.opt), the servername option specifies which server, of those named in your dsm.sys file, to contact for backup-archive services. When specified in a client user-options file (dsm.opt) or on the command line, the servername option overrides the default server specified in your client system options file.

Note:

1. You cannot use the servername option to override the server that is specified for migration in your client system options file.
2. If the Tivoli Storage Manager server name changes or Tivoli Storage Manager clients are directed to a different Tivoli Storage Manager server, all clients must have a new password initialized for the new server name.

Supported Clients

This option is for all UNIX and Linux clients.

Options File

Place this option in the client user options file (dsm.opt) and the dsm.sys file. In the dsm.sys file, the servername option is the beginning of a server stanza.

Syntax

▶▶──SErvername── ──*servername*──▶◀

Parameters

servername
    In your dsm.sys file, specify the name you want to assign to a particular server. In your client user-options file (dsm.opt) or on the command line, specify the name of the server you want to contact for backup-archive services. A server name is not case sensitive; it can have up to 64 characters.

Examples

Options file:
    servername server_a

Command line:
    -se=server_b

This option is valid only on the initial command line. It is not valid in interactive mode.

Sessioninitiation

Use the sessioninitiation option to control whether the server or client initiates sessions through a firewall. The default is that the client initiates sessions. You can use this option with the **schedule** command.

You must be an authorized user to use the sessioninitiation option.

For the client scheduler, it is unnecessary to open *any* ports on the firewall. If you set the sessioninitiation option to *serveronly*, the client will not attempt to contact the server. *All sessions must be initiated by server prompted scheduling* on the port defined on the client with the tcpclientport option. The sessioninitiation option only affects the behavior of the client scheduler running in the prompted mode. If you set the sessioninitiation option to *serveronly*, with the exception of client acceptor daemon-managed schedulers, the command-line client, backup-archive client Java GUI, and Web client GUI still attempts to initiate sessions.

Important: You cannot use the dsmcad for scheduling when you set the sessioninitiation option to *serveronly*.

Note: If you set the sessioninitiation option to *serveronly*, the client setup wizard and scheduler service are unable to authenticate to the Tivoli Storage Manager server. In this case, you can execute the scheduler from the command line (dsmc schedule) and enter the password for your node when prompted.

A similar problem can occur if an encryption key is required for backup operations. In this case, you can execute the scheduler from the command line (dsmc schedule) and enter the encryption key when prompted. After the password and encryption key are updated, you must restart the scheduler.

If you set the sessioninitiation option to *client*, the client initiates sessions with the server by communicating on the TCP/IP port defined with the *server* option tcpport. This is the default. Server prompted scheduling can be used to prompt the client to connect to the server.

Note:
1. The Tivoli Storage Manager server can specify SESSIONINITiation=*clientorserver* or SESSIONINITiation=*serveronly* on the **register node** and **update node** commands. If the server specifies SESSIONINITiation=*clientorserver*, the client can decide which method to use. If the server specifies SESSIONINITiation=*serveronly*, all sessions are initiated by the server.
2. If sessioninitiation is set to *serveronly*, the value for the tcpclientaddress client option must be the same as the value for the *HLAddress* option of the **update node** or **register node** server command. The value for the tcpclientport client option must be the same as the value for the *LLAddress* option of the **update node** or **register node** server command.
3. The Tivoli Storage Manager client API does not support this option.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Scheduler** tab, **Session Initiation** field of the Preferences editor.

Syntax

```
►►──SESSIONINITiation──┬──────────────┬──────────────────►◄
                       ├──Client──────┤
                       └──SERVEROnly──┘
```
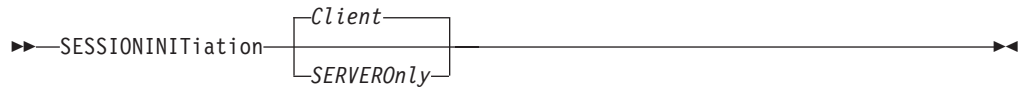
Parameters

Client
> Specifies that the client initiates sessions with the server by communicating on the TCP/IP port defined with the *server* option TCPPORT. This is the default. Server prompted scheduling can be used to prompt the client to connect to the server.

SERVEROnly
> Specifies that the server will not accept client requests for sessions. All sessions must be initiated by server prompted scheduling on the port defined on the client with the tcpclientport option. Except for client acceptor daemon-managed schedulers, the command-line client, backup-archive client Java GUI, and Web client GUI still attempts to initiate sessions.

Examples

Options file:
> ```
> sessioninitiation serveronly
> ```

Command line:
> ```
> schedule -sessioninitiation=serveronly
> ```

This option is valid only on the initial command line. It is not valid in interactive mode.

Related information

"Configuring the client scheduler" on page 59

"Tcpclientport" on page 438

Shmport

The shmport option specifies the TCP/IP port address of a server when using shared memory. All shared memory communications start with a TCP/IP connection.

You must be an authorized user to use the shmport option.

Note: The value specified for the shmport option in the dsm.sys file must match the value specified for shmport in the server options file.

Supported Clients

This option is valid for AIX, HP-UX, Linux, and Solaris clients *only*.

Options File

Place this option in the dsm.sys file within a server stanza.

Syntax

►►──SHMPort── ──*port_number*──►◄

Parameters

port_number
> Specifies the port number. You can specify a value from 1000 to 32767. The
> default value is 1510.

Examples

Options file:
> shmport 1580

Command line:
> Does not apply.

Showmembers

Use the showmembers option to display all members of a group.

You can use the showmembers option with the **query group**, and **restore group**
commands.

The showmembers option is not valid with the inactive option. If you want to
display members of a group that are not currently active, use the pitdate and
pittime options.

Supported Clients

This option is valid for all UNIX and Linux clients except Mac OS X.

Syntax

►►──SHOWMembers──►◄

Parameters

There are no parameters for this option.

Examples

Command line:
> restore group /virtfs/* -pick -showmembers

Skipacl

The skipacl option allows you to back up or restore ACL data.

When set to *yes* (default is *no*), the Tivoli Storage Manager client skips ACL processing completely and no ACL data is backed up or restored. The skipacl=*yes* option should be used only when ACLs are not defined on the file system or when the loss of data contained in the ACLs during restore is acceptable. The file mode permissions (rwx) continues to be backed up and restored as before.

If a prior version of the file being restored has ACL data backed up, ACLs are restored even when the skipacl option is set.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

Place this option in the client user options (dsm.opt) file.

Syntax

```
►►──SKIPACL──┬──No──┬──────────────────────────────────────────────────►◄
             └─Yes─┘
```

Parameters

No If you specify *No*, the ACL data is backed up. This is the default.

Yes
>   If you specify *Yes*, the ACL data is not backed up. skipacl=*yes* overrides skipaclupdatecheck settings.

Examples

Options file:
```
      skipacl yes
```

Skipaclupdatecheck

The skipaclupdatecheck option disables checksum and size comparisons of ACL data.

When set to *yes* (default is *no*), the Tivoli Storage Manager client will not perform checksum and size comparisons before or after backup and during incremental processing (ACL checksum from previous backup and current ACL) to detect ACL updates. However, current ACL data is backed up if the file is selected for backup due to other reasons. If only ACLs are updated on a file, the next incremental backup will not recognize this ACL update, and the file is not backed up.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

Place this option in the client user options (dsm.opt) file.

Syntax

```
►►──SKIPACLUPdatecheck──┬──No──┬─────────────────────────────────────────────◄
                        └──Yes──┘
```

Parameters

No If you specify *No*, the Tivoli Storage Manager client performs checksum and size comparisons of the ACL data, before and after backup and during incremental processing. This is the default.

Yes

If you specify *Yes*, the Tivoli Storage Manager client does not perform checksum and size comparisons of the ACL data.

Examples

Options file:
        skipaclup yes

Snapdiff

Using the snapdiff option with the **incremental** command streamlines the incremental backup process. The command runs an incremental backup of the files that were reported as changed by NetApp instead of scanning all of the volume for changed files.

You must configure a user ID and password on the Tivoli Storage Manager client to enable snapshot difference processing. For more information about setting up the snapdiff option, see "Configuring the file server connection for Tivoli Storage Manager" on page 91.

Use this option with an incremental backup of a NAS file server volume instead of a simple incremental or incremental with snapshotroot whenever the NAS file server is running ONTAP 7.3.0, or later, for performance reasons. Do not use the snapdiff and snapshotroot options together.

Restriction: Incremental backup using snapshot difference is only available with the Tivoli Storage Manager 64 bit AIX client and the Tivoli Storage Manager Linux x86/86_64 client.

The first time you perform an incremental backup with the snapshot difference option, a snapshot is created (the base snapshot) and a traditional incremental backup is run using this snapshot as the source. The name of the snapshot that is created is recorded in the Tivoli Storage Manager database.

The second time an incremental backup is run with this option, a newer snapshot is either created, or an existing one is used to find the differences between these two snapshots. The second snapshot is called the diffsnapshot. Tivoli Storage Manager then incrementally backs up the files reported as changed by NetApp to the Tivoli Storage Manager server. The filesystem that you select for snapshot difference processing must be mounted to the root of the volume. You cannot use the snapdiff option for any filesystem that is not mounted to the root of the volume. After backing up data using the snapdiff option, the snapshot that was used as the base snapshot is deleted from the `.snapshot` directory. Tivoli Storage

Manager does not delete the snapshot if it was not created by Tivoli Storage Manager. You can also run a snapshot-difference incremental-backup by specifying the -diffsnapshot=latest option.

For NAS and N-Series file servers that are running ONTAP 7.3.0, or later, you can use the createnewbase option to backup any files that might have been skipped due to one of the following reasons:

- A file is excluded because the include-exclude file has an exclude rule in effect. A file is excluded when you have not changed the include-exclude file, but you have removed the rule that excluded the file. The NetApp API only detects file changes between two snapshots, not changes to the include-exclude file.
- If you have added an include statement to the option file, that include option does not take effect unless NetApp detects that the file has changed. This is because Tivoli Storage Manager does not inspect each file on the volume during backup.
- You have used the **dsmc delete backup** command to explicitly delete a file from the Tivoli Storage Manager inventory. NetApp does not detect that a file has been manually deleted from Tivoli Storage Manager. Therefore, the file remains unprotected in Tivoli Storage Manager storage until it is changed on the volume and the change is detected by NetApp, signalling Tivoli Storage Manager to back it up again.
- Policy changes such as changing the policy from mode=modified to mode=absolute are not detected.
- The entire file space is deleted from the Tivoli Storage Manager inventory. This causes the snapshot difference option to create a snapshot to use as the source, and runs a full incremental backup.
- A file is excluded from backup because the file name contains a character that is not in the 7 bit-ASCII character set. The createnewbase option creates a new base snapshot and uses it as a source to run a full incremental backup. NetApp controls what constitutes a changed object.

Note: In the list of options used by the traditional **incremental** command, the last column shows the interaction of each option with the snapdiff option. The definitions of *valid*, *not valid*, and *no effect* are:

Valid   Processing runs normally when the option is used.

Not valid
         If the option is used with the snapdiff option, an error message is generated.

No effect
         The option can be used, but it is ignored.

Table 92. Incremental command: Related options

Option	Where specified	With snapdiff
asnodename "Asnodename" on page 268	Client system options file (dsm.sys) or command line.	Valid
automount "Automount" on page 276	Client options file (dsm.opt).	No effect
changingretries "Changingretries" on page 278	Client system options file (dsm.sys) or command line.	No effect
compressalways "Compressalways" on page 284	Client options file (dsm.opt) or command line.	Valid

Table 92. Incremental command: Related options (continued)

Option	Where specified	With snapdiff
compression "Compression" on page 285	Client system options file (dsm.sys) *within* a server stanza, or command line.	Valid
createnewbase "Createnewbase" on page 287	Command line only.	Valid
diffsnapshot "Diffsnapshot" on page 297	Command line only.	Valid
dirsonly "Dirsonly" on page 300	Command line only.	Valid
domain "Domain" on page 303	Client system options file (dsm.sys), client user-options file (dsm.opt), or command line.	Valid
efsdecrypt "Efsdecrypt" on page 310	Client system options file (dsm.sys), client user-options file (dsm.opt), or command line.	No effect
enablelanfree "Enablelanfree" on page 314	Client system options file (dsm.sys) or command line.	Valid
encryptiontype "Encryptiontype" on page 315	system-options file (dsm.sys) within a server stanza.	Valid
encryptkey "Encryptkey" on page 316	System-options file (dsm.sys) within a server stanza.	Valid
exclude.fs.nas "Exclude options" on page 321	Client system options file (dsm.sys).	No effect
filelist "Filelist" on page 333	Command line only.	Not valid
filesonly "Filesonly" on page 336	Command line only.	Valid
followsymboliclink "Followsymbolic" on page 336	Client options file (dsm.opt).	No effect
include.fs.nas "Include options" on page 351	Client system options file (dsm.sys) or command line.	No effect
inclexcl "Inclexcl" on page 349	Client system options file (dsm.sys).	Valid, but only when a file change is detected by NetApp.
incrbydate "Incrbydate" on page 357	Command line only.	Not valid
memoryefficientbackup "Memoryefficientbackup" on page 368	This option is allowed in both dsm.sys and dsm.opt, but the value in dsm.opt is ignored if it also appears in dsm.sys. You can also place this option within a server stanza, or on the initial command line.	No effect
monitor "Monitor" on page 371	Command line only.	Not valid
nojournal "Nojournal" on page 375	Command line only.	Not valid
postsnapshotcmd "Postsnapshotcmd" on page 386	Client system options file (dsm.sys) or with the include.fs option.	Valid
preservelastaccessdate "Preservelastaccessdate" on page 389	Client user-options file (dsm.opt) or command line.	Valid

Table 92. Incremental command: Related options (continued)

Option	Where specified	With snapdiff
presnapshotcmd "Presnapshotcmd" on page 392	Client system options file (dsm.sys) or with the include.fs option.	Valid
removeoperandlimit "Removeoperandlimit" on page 396	Command line only.	Valid
skipaclupdatecheck "Skipaclupdatecheck" on page 417	Client options file (dsm.opt).	Valid
snapshotcachesize "Snapshotcachesize" on page 422	Client system options file (dsm.sys) or with the include.fs option.	No effect
snapshotproviderfs "Snapshotproviderfs" on page 423	System-options file (dsm.sys) within a server stanza or with the include.fs option.	Not valid
snapshotproviderimage "Snapshotproviderimage" on page 424	Client system options file (dsm.sys) or with the include.image option.	Not valid
snapshotroot "Snapshotroot" on page 425	Command line only.	Not valid
subdir "Subdir" on page 432	Client options file (dsm.opt) or command line.	Not valid
tapeprompt "Tapeprompt" on page 433	Client options file (dsm.opt) or command line.	Valid
toc "Toc" on page 444	Command line only.	Not valid
virtualfsname "Virtualfsname" on page 452	Command line only.	Not valid
virtualmountpoint "Virtualmountpoint" on page 452	Client system options file (dsm.sys).	Not valid

Supported Clients

This option is valid for the following clients:

* AIX 64 bit clients
* Linux x86/86_64 clients

Syntax

```
►►──SNAPDiff──────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Command line:
      Perform a snapdiff incremental backup of an NFS mounted file system /vol/vol1 hosted on the file sever `homestore.example.com`, where /net/home1 is the mount point of /vol/vol1.

```
incremental -snapdiff -diffsnapshot=latest /net/home1
```

Command line:

> Perform a one-time full incremental backup after detecting that the NetApp server has migrated to a unicode-enabled file server from a server that did not support unicode file names.
>
> `dsmc incremental -snapdiff -createnewbase=migrate /net/home1`
>
> Perform a snapshot difference incremental backup after detecting that the NetApp server has migrated to a unicode-enabled file server from a server that did not support unicode file names. This command suppresses the warning message.
>
> `dsmc incremental -snapdiff -createnewbase=ign /net/home1`
>
> Perform a full incremental backup because you have made some include or exclude changes:
>
> `dsmc incremental -snapdiff -createnewbase=yes /net/home1`

Related information

"Configuring the file server connection for Tivoli Storage Manager" on page 91

"Diffsnapshot" on page 297

"Set Password" on page 555

"Configuring the file server connection for Tivoli Storage Manager" on page 91

Snapshotcachesize

Use the snapshotcachesize option to specify an appropriate size to create the snapshot.

The size estimation is needed for storing the original data blocks for modified and deleted data for the point in time when the snapshot was taken.

For snapshot-based file backup or archive, use the snapshotcachesize option with the include.fs option, or in the server stanza in the dsm.sys file.

For snapshot-based image backups, use the snapshotcachesize option with the **backup image** command, the include.image option, or in your dsm.sys file.

Supported Clients

This option is valid for AIX and Linux clients *only*. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the server stanza in the dsm.sys file. You can set this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

►►──SNAPSHOTCACHESize── ─*size*──►◄

Parameters

size

Specifies an appropriate size to create the snapshot for storing the original data blocks for modified and deleted data for the point in time when the snapshot was taken. The value is the percent of the file system size that is changed due to file system activity. The range of values is 1 to 100 percent. For AIX JFS2 and Linux the default value is 100 percent of the file system size. If a sufficient amount of free space is not available to create the snapshot, the command fails with an error message. You can then either increase the size of the volume group or retry the operation. If based on your experience with your AIX JFS2 file system activity, you find that a snapshot size of 100 percent is not necessary, you can fine-tune the value.

Examples

Options file:

```
snapshotcachesize 95
AIX only: include.fs /kalafs1
  snapshotproviderfs=JFS2 snapshotcachesize=95
AIX only: include.image /kalafs2
  snapshotcachesize=95
Linux only: include.image /linuxfs1
  snapshotcachesize=100
```

Command line:
```
-snapshotcachesize=95
```

Related information

See "Include options" on page 351 for more information about include.fs.

Snapshotproviderfs

Use the snapshotproviderfs option to enable snapshot-based file backup and archive operations, and to specify a snapshot provider.

You must be a root user to perform a snapshot-based file backup or archive operation. If you are not a root user, the operation fails with an error message.

Supported Clients

This option is valid for AIX clients *only*. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Specify this option in the server stanza of the system-options file, dsm.sys, to enable snapshots for all JFS2 file systems on the client. You can override the client-wide option for a specific operation by specifying this option on the command line for the backup and archive commands. You can also override the client-wide option for a specific file system by using the include.fs statement in the dsm.sys file. You can also set this option using the Preferences editor.

Syntax

```
►►──SNAPSHOTPROVIDERFs── ──value────────────────────────────────────────────►◄
```

Parameters

value
>   Specifies one of the following values:

JFS2
>   Specifies that you want to perform a snapshot-based file backup or archive
>   while the file system is available to other system applications. Valid for JFS2
>   file systems on AIX clients *only*.

NONE
>   Specifies that no snapshots should be used. A file backup or archive operation
>   is performed using the specified file system. This is the default.

Examples

Options file:

```
        snapshotproviderfs JFS2
        include.fs  /kalafs1   snapshotproviderfs=JFS
```

Command line:
```
        -SNAPSHOTPROVIDERFs=JFS2
```

Snapshotproviderimage

Use the snapshotproviderimage option to enable snapshot-based image backup,
and to specify a snapshot provider.

You must be a root user to perform a snapshot-based image backup operation. If
you are not a root user, the operation fails with an error message.

Supported Clients

This option is valid for AIX and Linux clients *only*. The Tivoli Storage Manager API
does not support this option. The server can also define this option.

Options File

Specify this option in the server stanza of the system-options file, dsm.sys, to
enable snapshots for all the file systems on the client. You can override the
client-wide option for a specific operation by specifying this option on the
command line for the **backup image** command. You can also override the
client-wide option for a specific file system using the include.image statement in
the dsm.sys file. You can also set this option using the Preferences editor.

Syntax

```
►►──SNAPSHOTPROVIDERImage── ──value──────────────────────────────────────────►◄
```

Parameters

value

Specifies one of the following values:

JFS2

Specifies that you want to perform a snapshot-based image backup while the file system is available to other system applications. This is the default for JFS2 file systems. Valid for AIX clients *only*.

LINUX_LVM

Specifies that you want to perform a snapshot-based image backup while the file system is available to other system applications. This is the default for file systems residing on logical volumes created by the Linux Logical Volume Manager. Valid for Linux clients *only*.

NONE

Specifies that you do not want to perform a snapshot-based image backup operation. This performs a static image backup operation using the specified file system. This is the default for file systems other than AIX JFS2 and Linux LVM.

Examples

Options file:

```
snapshotprovideri JFS2
include.image   /kalafs1   snapshotprovideri=JFS2
```

Command line:
```
-SNAPSHOTPROVIDERImage=NONE
```

Snapshotroot

Use the snapshotroot option with the **incremental**, **selective**, or **archive** commands in conjunction with a independent software vendor application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server.

The snapshotroot option can be used to back up NFS mounted file systems. Both the back up specification (source) and the snapshotroot value can be a NFS mounted file specification. For example, the snapshotroot option can be used to back up a NFS file system hosted on a network-attached storage (NAS) that supports snapshot.

This option should be used with an incremental backup of a NAS file server volume instead of a simple incremental or incremental with snapshotroot option whenever the NAS file server is running ONTAP V7.3 for performance reasons. The snapdiff and snapshotroot options should not be used together.

In the following example, `filesystem test495` is NFS-mounted from a NAS file server `philo` and `/philo/test945/.snapshot/backupsnap` represents the snapshot created at the NAS file server.

```
dsmc i /philo/test945 -snapshotroot=/philo/test945/.snapshot/backupsnap
```

You can also specify a directory with the snapshotroot option when backing up each fileset as a separate file space.

The snapshotroot option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot.

For example, consider an application that takes a snapshot of the /usr file system and mounts it as /snapshot/day1. If you back up this data using the following command, a unique file space called /snapshot/day1 is created on the server.

```
dsmc incremental /snapshot/day1
```

However, you might want to associate the snapshot data with the data already processed for the /usr file system. Using the snapshotroot option, you can associate the data with the file space corresponding to the /usr file system on the Tivoli Storage Manager server:

```
dsmc incremental /usr -snapshotroot=/snapshot/day1
```

On a subsequent day, you can back up a snapshot that was written to an alternate location, but managed under the same file space on the server:

```
dsmc incremental /usr -snapshotroot=/snapshot/day2
```

You can perform incremental backups, selective backups, or archives of a single directory, directory structure, or single file using the snapshotroot option. In all instances, the snapshotroot option must identify the root of the logical volume that was created by the snapshot. For example:

```
dsmc incremental /usr/dir1/* -subdir=yes
  -snapshotroot=/snapshot/day1
dsmc selective /usr/dir1/sub1/file.txt
  -snapshotroot=/snapshot/day1
dsmc archive /usr/dir1/sub1/*.txt
  -snapshotroot=/snapshot/day1
```

If you want to include or exclude specific file specifications, the include and exclude statements should contain the name of the file system that was the *source* of the snapshot (the /usr file system), and *not* the name of the target of the snapshot (/snapshot/day1). This allows you to preserve a set of include and exclude statements regardless of the name of the logical volume to which the snapshot is written. The following are examples of include and exclude statements.

```
include /usr/dir1/*.txt 1yrmgmtclass
exclude /usr/mydocs/*.txt
```

The following include-exclude statements are not valid because they contain the name of the snapshot:

```
include /snapshot/day1/dir1/*.txt 1yrmgmtclass
exclude /snapshot/day1/mydocs/*.txt
```

You must use the snapshotroot option in conjunction with a single file specification for a incremental, selective, or archive operation. You cannot specify multiple file specifications or no file specifications. For example, these commands are valid:

```
dsmc incremental /usr -snapshotroot=/snapshot/day1
dsmc incremental /usr/dir1/* -snapshotroot=/snapshot/day1
```

The following command is invalid because it contains two file specifications:

```
dsmc incremental /usr/dir1/* /home/dir2/*
  -snapshotroot=/snapshot/day1
```

The following command is invalid because it contains no file specification:

```
dsmc incremental -snapshotroot=/snapshot/day1
```

Note:

1. Ensure that the snapshotroot references a snapshot of the correct volume. Ensure that snapshotroot refers to the root of the snapshot. If these rules are not followed, unintended results such as files expiring incorrectly can result.
2. You cannot use the snapshotroot option in conjunction with the filelist option.
3. You cannot use the snapshotroot option with the snapdiff option.
4. You can use the snapshotroot option in conjunction with the preschedulecmd and postschedulecmd options, or in a automated script that you execute with the Tivoli Storage Manager client scheduler.

Supported Clients

This option is valid for the following clients:
- UNIX and Linux clients except Mac OS X.

Syntax

```
►►──SNAPSHOTRoot =── ──snapshot_volume_name───────────────────────────►◄
```

Parameters

snapshot_volume_name
    Specifies the root of the logical volume created by the independent software vendor snapshot application.

Examples

Command line:
```
dsmc incremental /usr -SNAPSHOTRoot=/snapshot/day1
```

Srvoptsetencryptiondisabled

The srvoptsetencryptiondisabled option allows the client to ignore encryption options in a client options set from a Tivoli Storage Manager server.

If the option is set to yes in the client options file, the client ignores the following options in a client options set from a Tivoli Storage Manager server:
- encryptkey generate
- exclude.encrypt
- include.encrypt

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (dsm.sys) within a server stanza.

Syntax

```
►►──SRVOPTSETENCryptiondisabled──┬──no──┬──────────────────────────────────◄◄
                                 └──yes──┘
```
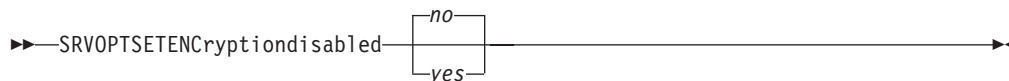
Parameters

yes
> Tivoli Storage Manager client ignores the values of the listed encryption options in a client options set from a Tivoli Storage Manager server.

no Tivoli Storage Managerclient processes the setting of the listed encryption options in a client options set from a Tivoli Storage Manager server. This is the default.

Examples

Options file:
> srvoptsetencryptiondisabled no

Command line:
> Does not apply.

Srvprepostscheddisabled

The srvprepostscheddisabled option specifies whether to prevent the pre-schedule and post-schedule commands specified by the Tivoli Storage Manager administrator from executing on the client system, when performing scheduled operations.

The srvprepostscheddisabled option can be used in conjunction with the schedcmddisabled and srvprepostscheddisabled options to disable the execution of any unwanted Operating System command by a Tivoli Storage Manager administrator on a Tivoli Storage Manager client node.

You must be an authorized user to use the srvprepostscheddisabled option.

Supported Clients

This option is valid for all Tivoli Storage Manager clients that use the Tivoli Storage Manager client scheduler. The server cannot define this option.

Options File

Place this option in the dsm.sys file within a server stanza for the scheduler. You can set this option on the **Scheduler** tab of the Preferences editor.

Syntax

```
►►──SRVPREPOSTSCHeddisabled──┬──No──┬──────────────────────────────────────◄◄
                             └──Yes──┘
```

Parameters

No Specifies that Tivoli Storage Manager allows pre-schedule and post-schedule commands defined by the Tivoli Storage Manager administrator to execute on the client system, when performing scheduled operations. If a pre-schedule or

a post-schedule command is defined by both the client and the Tivoli Storage Manger administrator, the command defined by the administrator overrides the corresponding command defined in the client option file. This is the default.

Yes

Specifies that Tivoli Storage Manager prevents pre-schedule and post-schedule commands defined by the Tivoli Storage Manager administrator to execute on the client system, when performing scheduled operations. If a pre-schedule or a post-schedule command is defined by both the client and the Tivoli Storage Manger administrator, the command defined by the administrator will *not* override the corresponding command defined in the client option file. This option can be used in conjunction with the schedcmddisabled and srvprepostscheddisabled options.

Examples

Options file:
    srvprepostscheddisabled yes

Command line:
    Does not apply.

Srvprepostsnapdisabled

The srvprepostsnapdisabled option specifies whether to prevent the pre-snapshot and post-snapshot commands specified by the Tivoli Storage Manager administrator from executing on the client system, when performing scheduled image snapshot backup operations.

The srvprepostsnapdisabled option can be used in conjunction with the schedcmddisabled and srvprepostsnapdisabled options to disable the execution of any unwanted Operating System command by a Tivoli Storage Manager administrator on a Tivoli Storage Manager client node.

You must be an authorized user to use the srvprepostsnapdisabled option.

Supported Clients

This option is valid for Linux clients that support the image snapshot backup command. The server cannot define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza for the scheduler. You can set this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

```
              ┌─No──┐
►►──SRVPREPOSTSNApdisabled──┼─────┼──────────────────────────────────────►◄
              └─Yes─┘
```

Parameters

No Specifies that Tivoli Storage Manager allows pre-snapshot and post-snapshot

commands defined by the Tivoli Storage Manager administrator to execute on the client system, when performing scheduled image snapshot backup operations. If a pre-snapshot or a post-snapshot command is defined by both the client and the Tivoli Storage Manger administrator, the command defined by the administrator overrides the corresponding command defined in the client option file. This is the default.

Yes

Specifies that Tivoli Storage Manager allows pre-snapshot and post-snapshot commands defined by the Tivoli Storage Manager administrator to execute on the client system, when performing scheduled image snapshot backup operations. If a pre-snapshot or a post-snapshot command is defined by both the client and the Tivoli Storage Manger administrator, the command defined by the administrator will *not* override the corresponding command defined in the client option file. This option can be used in conjunction with the schedcmddisabled and srvprepostsnapdisabled options.

Examples

Options file:
        srvprepostsnapdisabled yes

Command line:
        Does not apply.

Ssl

Use the ssl option to enable Secure Sockets Layer (SSL), to provide secure client and server communications.

You must be an authorized user to use the ssl option.

Supported Clients

This option is valid for AIX, HP-UX, Linux, and Solaris clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Communications** tab of the Preferences editor.

Syntax

```
►►──SSL──┬──No──┬──────────────────────────────────────────────────►◄
         └─Yes─┘
```

Parameters

No Specifies that the Tivoli Storage Manager client disables SSL. *No* is the default.

Yes

Specifies that the Tivoli Storage Manager client enables SSL.

To enable SSL, specify ssl=*yes* and change the value of the tcpport option. Changing the value of the tcpport option is needed, because the Tivoli Storage Manager server is typically set up to listen for SSL connections on a separate port.

Examples

Options file:
> ssl yes

Command line:
> Does not apply.

Related information

"Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer" on page 65.

Stagingdirectory

The stagingdirectory option defines the location where the client can keep any data it generates in order to perform its operations. The data is deleted when processing is complete.

This option is used by Active Directory objects query and restore operations, and during snapshot difference backup operations on the AIX client.

Supported Clients

This option is valid for AIX clients. The server can also define this option.

Options File

Place this option in the dsm.sys file within a server stanza. When stagingdirectory is specified on the command line, it overrides the values specified in the options file.

Syntax

```
►►──STAGINGDIRectory──path──────────────────────────────────────►◄
```

Parameters

path
> Specifies the directory path where the client writes staging data. If you do not specify a staging directory, Tivoli Storage Manager stores temporary data in the temporary file system (typically /tmp).

Examples

Options file:
> stagingdirectory /usr/tsmdata
>
> stagingdirectory /private/tmp

Command line:
> -stagingdir="/tmp/tsmtempdata"

Related information

"Diffsnapshot" on page 297

"Snapdiff" on page 418

Subdir

The subdir option specifies whether you want to include subdirectories of named directories for processing.

You can use the subdir option with the following commands:
- **archive**
- **delete archive**
- **delete backup**
- **incremental**
- **query archive**
- **query backup**
- **restore**
- **restore backupset**
- **restore group**
- **retrieve**
- **selective**

For example, if you set the subdir option to *yes* when backing up a specific path and file, Tivoli Storage Manager recursively backs up *all* subdirectories under that path, and any instances of the specified file that exist under *any* of those subdirectories.

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client user-options file (dsm.opt).

Syntax

```
            ┌─No──┐
►►──SUbdir──┤     ├────────────────────────────────────────────►◄
            └─Yes─┘
```

Parameters

No Subdirectories are not processed. This is the default.

Yes

    Subdirectories are processed. Because the client program searches all subdirectories of a directory that is being processed, processing can take longer to complete. Specify *Yes* only when necessary.

    Note: If you use the preservepath option in addition to subdir=*yes*, it can affect which subdirectories are processed.

    If a subdirectory is a mounted file system, it is not processed even if you specify subdir=*yes*.

Examples

Options file:
```
        subdir no
```

Command line:

To restore the structure:

```
/Users/mike/dir1
/Users/mike/dir1/file1
/Users/mike/dir1/dir2
/Users/mike/dir1/dir2/file1
```

enter any of the following commands:

```
dsmc rest "/Users/van/dir1/*" /Users/mike/ -su=yes
dsmc rest "/Users/van/dir1/file*" /Users/mike/ -su=yes
dsmc rest "/Users/van/dir1/file1*" /Users/mike/ -su=yes
```

To restore the structure:

```
/path2/dir1
/path2/dir1/file1
/path2/dir1/dir2
/path2/dir1/dir2/file1
```

enter any of the following commands:

```
dsmc rest "/path/dir1/*" /path2/ -su=yes
dsmc rest "/path/dir1/file*" /path2/ -su=yes
dsmc rest "/path/dir1/file1*" /path2/ -su=yes
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

Related information

"Preservepath" on page 390

Tapeprompt

The tapeprompt option specifies whether you want Tivoli Storage Manager to wait for a tape mount if it is required for a backup, archive, restore, or retrieve process, or to be prompted for a choice.

In the Tivoli Storage Manager GUI, the Media Mount dialog can display the **Information Not Available** value in the Device and Volume Label fields if you perform a standard (also known as classic) restore or retrieve operation. This value means that this information is only available for no query restore or retrieve operations; not a standard restore or retrieve operation. The Device field displays the name of the device on which to mount the media needed to process an object. The Volume Label field displays the name of the volume needed to process an object.

Tape prompting does not occur during a scheduled operation regardless of the setting for the tapeprompt option.

The tapeprompt option can be used with the following commands:
* **archive**
* **delete archive**
* **delete backup**
* **incremental**
* **restore**

- **retrieve**
- **selective**

Note: The server can also define this option.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **General** tab, **Prompt before mounting tapes** checkbox of the Preferences editor.

Syntax

```
►►──TAPEPrompt──┬─No──┬──────────────────────────────────────►◄
                └─Yes─┘
```
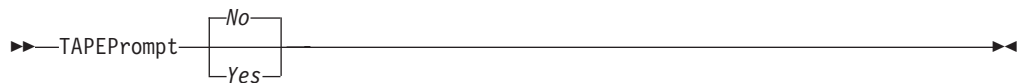
Parameters

No You are not prompted for your choice. The server waits for the appropriate tape to mount. This is the default.

    Note: For API applications, this permits backup directly to tape.

Yes

    You are prompted when a tape is required to back up, archive, restore, or retrieve data. At the prompt, you can wait for the appropriate tape to be mounted, always wait for a tape to be mounted, skip a particular object, skip all objects on a single tape, skip all objects on all tapes, or cancel the entire operation.

Examples

Options file:
      `tapeprompt yes`

Command line:
      `-tapep=yes`

This option is valid only on the initial command line. It is not valid in interactive mode.

Tcpadminport

Use the tcpadminport option to specify a separate TCP/IP port number on which the server waits for requests for administrative client sessions, allowing secure administrative sessions within a private network.

You must be an authorized user to use the tcpadminport option.

The client tcpadminport setting depends on how the Tivoli Storage Manager server tcpadminport and adminonclientport options are configured. The Tivoli Storage

Manager server has a tcpadminport setting that indicates on which port the server listens for administrative sessions, and the adminonclientport setting, which can be either *yes* or *no*.

If tcpadminport is not set on the server, then administrative sessions are allowed on the same port as client sessions.

If tcpadminport is set on the server, then administrative sessions are allowed on the port specified by that setting. In this case, if adminonclientport *yes* is in effect, then administrative sessions can connect on either the regular client port *or* the port specified by tcpadminport. If adminonclientport *no* is in effect, then administrative sessions can connect only on the port specified by tcpadminport.

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Communication** tab, **Admin Port** field of the Preferences editor.

Syntax

```
►►──TCPADMINPort─────────────────────────────────────────────────►◄
                  └─admin_port_address─┘
```
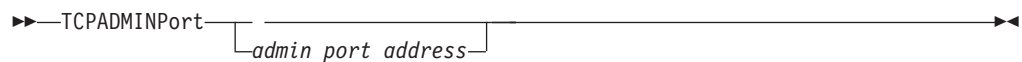
Parameters

admin_port_address
    Specifies the port number of the server. The default value is the value of the tcpport option.

Examples

Options file:
```
tcpadminport 1502
```

Tcpbuffsize

The tcpbuffsize option specifies the size of the internal TCP/IP communication buffer used to transfer data between the client node and server. Although it uses more memory, a larger buffer can improve communication performance.

You must be an authorized user to use the tcpbuffsize option.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Communication** tab, **Buffer Size** field of the Preferences editor.

Syntax

```
►►──TCPBuffsize── ──size────────────────────────────────────────────────────►◄
```

Parameters

size

    Specifies the size, in kilobytes, that you want to use for the internal TCP/IP communication buffer. The range of values is 1 through 512; the default is 32.

    Depending on the operating system communication settings, your system might not accept all values in the range of 1 through 512.

Examples

Options file:

       `tcpb 32`

Command line:

       `-tcpbuffsize=32`

This option is valid only on the initial command line. It is not valid in interactive mode.

Tcpcadaddress

The tcpcadaddress option specifies a TCP/IP address for dsmcad. Normally, this option is not needed. Use this option only if your client node has more than one TCP/IP address, or if TCP/IP is not the default communication method.

You must be an authorized user to use the tcpcadaddress option.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza.

Syntax

```
►►──TCPCADAddress── ──cad_address────────────────────────────────────────────►◄
```

Parameters

cad_address

    Specifies a TCP/IP Internet domain name or a numeric IP address. You can specify TCPip or V6Tcpip. This option uses TCP/IP Version 4 if TCPip is specified and uses Version 6 if V6Tcpip is specified. You can only use IPv6 addresses if you specified the commmethod *V6Tcpip* option.

Examples

Options file:

       `tcpcada dsmclnt.sanjose.ibm.com`

Command line:

```
-tcpcadaddress=128.33.10.249

-tcpcadaddress=khoyt.mycompany.mydomain.com

-tcpcadaddress=2002:92b:111:221:128:33:10:249
```

This option is valid only on the initial command line of the dsmcad program. It is not valid with other dsm modules.

Related information

See "Commmethod" on page 281 to determine if your client node has more than one TCP/IP address, or if TCP/IP is not the default communication method.

Tcpclientaddress

The tcpclientaddress option specifies a TCP/IP address if your client node has more than one address, and you want the server to contact an address other than the one that was used to make the first server contact.

The server uses this address when it begins the server prompted scheduled operation.

You must be an authorized user to use the tcpclientaddress option.

Use this option only if you use the *prompted* parameter with the schedmode option.

If sessioninitiation is set to *serveronly*, the value for the tcpclientaddress client option should be the same as the value for the *HLAddress* server option.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Scheduler** tab, **Your TCP/IP address** field of the Preferences editor.

Syntax

```
►►──TCPCLIENTAddress── ──client_address────────────────────────────►◄
```

Parameters

client_address
    Specifies the TCP/IP address you want the server to use to contact your client node. Specify a TCP/IP Internet domain name or a numeric IP address. The numeric IP address can be either a TCP/IPv4 or TCP/IPv6 address. You can only use IPv6 addresses if you specified the commmethod *V6Tcpip* option.

Examples

Options file:

```
        tcpclienta dsmclnt.mycompany.mydomain.com
        or
        tcplienta 192.168.0.1
```

Options file:
        tcpclienta dsmclnt.sanjose.example.com

This option is valid only on the initial command line. It is not valid in interactive mode.

Tcpclientport

The tcpclientport option specifies a TCP/IP port number for the server to contact the client when the server begins the server prompted scheduled operation.

You must be an authorized user to use the tcpclientport option

Use this option only if you specify the *prompted* parameter with the schedmode option.

If sessioninitiation is set to *serveronly*, the value for the tcpclientport client option should be the same as the value for the **LLAddress** server option.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Scheduler** tab, **Your TCP/IP port** field of the Preferences editor.

Syntax

►►──TCPCLIENTPort── ──*client_port_address*──────────────────────────────────►◄

Parameters

client_port_address
    Specifies the TCP/IP port address you want the server to use to contact your client node. The range of values is 1 through 32767; the default is 1501.

Examples

Options file:
        tcpclientp 1502

Command line:
        -tcpclientport=1492

This option is valid only on the initial command line. It is not valid in interactive mode.

Tcpnodelay

The tcpnodelay option specifies whether the client disables the delay of sending successive small packets on the network, per transaction.

You must be an authorized user to use the tcpnodelay option.

Change the value from the default of *yes* only under one of the following conditions:

- You are directed to change the option by IBM technical support.
- You fully understand the effects of the TCP Nagle algorithm on network transmissions. Setting the option to *no* enables the Nagle algorithm, which delays sending small successive packets.

Supported Clients

This option is valid for all UNIX and Linux clientst.

Options File

Place this option in the client system-options file (dsm.sys) *within* a server stanza. You can set this option on the **Communication** tab of the Preferences editor.

Syntax

```
                        ┌─Yes─┐
►►──TCPNodelay──────────┴─────┴────────────────────────────────────►◄
                        └─No──┘
```

Parameters

No   Specifies that the server does not allow successive small packets to be sent immediately over the network. Setting this option to *no* can degrade performance.

Yes
> Specifies that the server or client allows successive small packets to be sent immediately over the network. The default is *yes*.

Examples

Options file:
> ```
> tcpnodelay yes
> ```

Command line:
> Does not apply.

Tcpport

The tcpport option specifies a TCP/IP port address for a Tivoli Storage Manager server. You can obtain this address from your administrator.

You must be an authorized user to use the tcpport option.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Communication** tab, **Server Port** field of the Preferences editor.

Syntax

►►──TCPPort── ──*port_address*───►◄

Parameters

port_address
>   Specifies the TCP/IP port address that is used to communicate with a server. The range of values is 1 through 32767; the default is 1500.

Examples

Options file:
>   tcpp 1501

Command line:
>   Does not apply.

Tcpserveraddress

The tcpserveraddress option specifies the TCP/IP address for a Tivoli Storage Manager server. You can obtain this server address from your administrator.

You must be an authorized user to use the tcpserveraddress option.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Communication** tab, **Server Address** field of the Preferences editor.

If this option is not specified, the client attempts to contact a Tivoli Storage Manager server running on the same computer as the Tivoli Storage Manager client.

Syntax

►►──TCPServeraddress── ──*server_address*───────────────────────────────────►◄

Parameters

server_address
>   Specifies a 1 to 64 character TCP/IP address for a server. Specify a TCP/IP domain name or a numeric IP address. The numeric IP address can be either a

TCP/IP v4 or TCP/IP v6 address. You can only use IPv6 addresses if you
specified the commmethod *V6Tcpip* option.

Examples

Options file:
> tcps dsmchost.endicott.ibm.com

Command line:
> Does not apply.

Tcpwindowsize

Use the tcpwindowsize option to specify, in kilobytes, the size you want to use for
the TCP/IP sliding window for your client node.

The sending host cannot send more data until it receives an acknowledgment and
a TCP receive window update. Each TCP packet contains the advertised TCP
receive window on the connection. A larger window allows the sender to continue
sending data and can improve communication performance.

You must be an authorized user to use the tcpwindowsize option.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option
on the **Communication** tab, **Window Size** field of the Preferences editor.

Syntax

```
►►──TCPWindowsize─ ─window_size─────────────────────────────────────────►◄
```

Parameters

window_size
> Specifies the size, in kilobytes, to use for your client node TCP/IP sliding
> window. The range of values is 0 through 2048. A value of 0 allows Tivoli
> Storage Manager to use the operating system default TCP window size. Values
> from 1 to 2048 indicate that the window size is in the range of 1KB to 2MB.
> For Mac OS X, the range of values is 1 through 128. If you specify a value less
> than 1, the TCP window size defaults to 1. If you specify a value greater than
> 2048, the TCP window size defaults to 2048.
>
> **Note:**
> 1. The TCP window acts as a buffer on the network. It is not related to the
>   tcpbuffsize option, or to the send and receive buffers allocated in client or
>   server memory.
> 2. A window size larger than the buffer space on the network adapter might
>   degrade throughput due to resending packets that were lost on the adapter.
> 3. Depending on the operating system communication settings, your system
>   might not accept all values in the range of values.

Examples

Options file:
>   tcpwindowsize 63

Command line:
>   -tcpw=63

This option is valid only on the initial command line. It is not valid in interactive mode.

Timeformat

The timeformat option specifies the format in which you want to display system time.

The AIX, Solaris, and HP-UX clients support locales other than English that describe every user interface that varies with location or language.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time the client is called. Consult the documentation on your local system for details about setting up your locale definition.

Note: The timeformat option does not affect the Web client. The Web client uses the time format for the locale that the browser is running in. If the browser is not running in a locale that Tivoli Storage Manager supports, the Web client uses the time format for US English.

You can use the timeformat option with the following commands:
- **delete archive**
- **delete backup**
- **expire**
- **query archive**
- **query backup**
- **query filespace**
- **query image**
- **query nas**
- **restore**
- **restore image**
- **restore nas**
- **retrieve**
- **set event**

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Regional Settings** tab, **Time Format** field of the Preferences editor.

Syntax

Parameters

format_number

> Displays time in one of the formats listed below. Select the format number that corresponds to the format you want to use. When you include the timeformat option in a command, it must precede the fromtime, pittime, and totime options.

0

> Use the locale-defined time format (does not apply to Mac OS X).
>
> For AIX, HP-UX, and Solaris: This is the default if the locale-specified format consists of digits, separator characters, and, if applicable, the AM or PM string.

1   23:00:00

> For AIX, HP-UX, and Solaris: This is the default if the locale-specified format does not consist of digits, separator characters, and, if applicable, the AM or PM string.

2   23,00,00

3   23.00.00

4   12:00:00 A/P

5   A/P 12:00:00

> For AIX, HP-UX, and Solaris: To set a particular time format, edit the source file for your locale and modify the t_fmt line to support your needs. Whatever time format you select applies both to output and to input.

"%H:%M:%S"

> Displays time in the form **hh:mm:ss** with **hh** ranging from 0 through 23.

"%H,%M,%S"

> Displays time in the form **hh,mm,ss** with **hh** ranging from 0 through 23.

"%I,%M,%S,0p"

> Displays time in the form **hh,mm,ssA/P** with **hh** ranging from 1 through 12 and **A/P** is the local abbreviation for ante-meridian (AM in English) or post-meridian (PM in English).

Examples

Options file:

```
timeformat 4
```

Command line:

```
-time=3
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

Toc

Use the toc option with the **backup nas** command or the include.fs.nas option to specify whether Tivoli Storage Manager saves table of contents (TOC) information for each file system backup.

You should consider the following when deciding whether you want to save TOC information:

- If you save TOC information, you can use the QUERY TOC server command to determine the contents of a file system backup in conjunction with the RESTORE NODE server command to restore individual files or directory trees.
- You can also use the Tivoli Storage Manager Web client to examine the entire file system tree and select files and directories to restore.
- Creation of a TOC requires that you define the TOCDESTINATION attribute in the backup copy group for the management class to which this backup image is bound. Note that TOC creation requires additional processing, network resources, storage pool space, and possibly a mount point during the backup operation.
- If you do not save TOC information, you can still restore individual files or directory trees using the RESTORE NODE server command, provided that you know the fully qualified name of each file or directory and the image in which that object was backed up.

Supported Clients

This option is valid for AIX and Solaris clients *only*. The Tivoli Storage Manager client API does not support this option.

Options File

Place the include.fs.nas statement containing the toc value in the dsm.sys file within a server stanza.
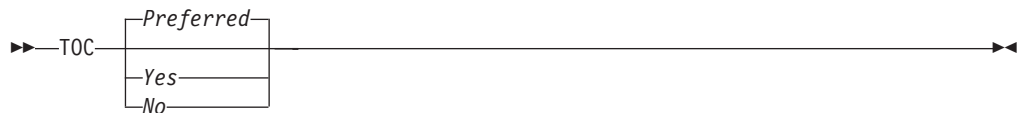
Syntax

```
         ┌─Preferred─┐
►►──TOC───┼─Yes───────┼──────────────────────────────────────────►◄
         └─No────────┘
```

Parameters

Yes
   Specifies that Tivoli Storage Manager saves TOC information during a NAS file system image backup. However, the backup fails if an error occurs during creation of the TOC.

No Specifies that Tivoli Storage Manager does not save TOC information during a NAS file system image backup.

Preferred
   Specifies that Tivoli Storage Manager saves TOC information during a NAS file system image backup. The backup does not fail if an error occurs during creation of the TOC. This is the default.

Note: If the mode option is set to *differential* and you set the toc option to *preferred* or *yes*, but the last full image does not have a TOC, Tivoli Storage Manager performs a full image backup and creates a TOC.

Examples

Options file:
```
include.fs.nas netappsj/vol/vol0 homemgmtclass toc=yes
```

Command line:
```
backup nas -nasnodename=netappsj /vol/vol0 -toc=yes
```

Todate

Use the todate option with the totime option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation

Use the todate and totime options with the fromtime and fromdate options to request a list of backed up or archived files within a period of time. For example, you might request a list of files that were backed up between 6:00 AM on July 1, 2002 and 11:59 PM on July 30, 2002.

Use the todate option with the following commands:
- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──TODate =── ──date──────────────────────────────────────────────►◄
```

Parameters

date
>   Specifies an ending date. Enter the date in the format you selected with the dateformat option.
>
>   When you include dateformat with a command, it must precede the fromdate, pitdate, and todate options.

Examples

Command line:
```
dsmc restore "/Users/agordon/Documents/*" -todate=12/11/2003
```

Command line:
```
dsmc restore "/home/user1/*" -todate=12/11/2003
```

Totime

Use the totime option with the todate option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation. Tivoli Storage Manager ignores this option if you do not specify the todate option.

Use the totime and todate options with the fromtime and fromdate options to request a list of files that were backed up within a period of time. For example, you might request a list of files that were backed up between 6:00 AM on July 1, 2003 and 11:59 PM on July 30, 2003.

Use the totime option with the following commands:
- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──TOTime =── ──time──────────────────────────────────────────────►◄
```

Parameters

time
> Specifies an ending time. If you do not specify a time, the time defaults to 23:59:59. Specify the time in the format you selected with the timeformat option.
>
> When you include the timeformat option in a command, it must precede the fromtime, pittime, and totime options.

Examples

Command line:
> ```
> dsmc restore "/Users/van/Documents/myfiles/*" -todate=09/17/2003
> -totime=23:00:00
> ```

Command line:
> ```
> dsmc restore "/home/user1/*" -todate=09/17/2003 -totime=23:00:00
> ```

Txnbytelimit

The txnbytelimit option specifies the number of kilobytes the client program buffers before it sends a transaction to the server.

A *transaction* is the unit of work exchanged between the client and server. A transaction can contain more than one file or directory, called a *transaction group*.

You must be an authorized user to use the txnbytelimit option.

You can control the amount of data sent between the client and server, before the server commits the data and changes to the server database, using the txnbytelimit option. Controlling the amount of data sent changes the speed of the client to perform the transactions. The amount of data sent applies when files are batched together during backup or when receiving files from the server during a restore procedure.

After the txngroupmax number is reached, the client sends the files to the server, even if the transaction byte limit is not reached.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **General** tab **Transaction Buffer Size** field of the Preferences editor.

Syntax

```
►►──TXNBytelimit── ──number───────────────────────────────────────►◄
```

Parameters

number
> Specifies the number of kilobytes the client program sends to the server before committing the transaction. The range of values is 300 through 34359738368 (32 GB). The default is 25600 KB. The number can be specified as an integer or as an integer with one of the following unit qualifiers:
> > K or k (kilobytes)
> > M or m (megabytes)
> > G or g (gigabytes)
>
> If no unit qualifier is specified, the integer is in kilobytes.
>
> **Restriction:** The txnbytelimit option does not support decimal numbers, and only one-unit letters are allowed. For example: K, M, or G.

Examples

Options file:
```
txnb 25600

txnb 2097152

txnb 2097152k

txnb 2048m

txnb 2g

txnb 32G
```
Command line:
```
-txnb=25600

-txnb=16G
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Type

Use the type option with the **query node** command to specify the type of node to query. Use this option with the **set event** command to activate, hold, or release.

Supported Clients

This option is also valid for the **set password** command with *TSM* on AIX clients.

This option is valid for AIX and Solaris clients *only*. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──TYpe =──┬─any────┬──────────────────────────────────►◄
            ├─nas────┤
            ├─server─┤
            └─client─┘
```

Parameters

nas
    Specifies all NAS nodes registered at the server.

server
    Specifies client nodes that are other Tivoli Storage Manager servers.

client
    Specifies client nodes that are backup-archive clients.

Examples

Command line:
```
query node -type=nas
```

Users

Authorized User: The users option authorizes specific users on your workstation to request services from a server.

You can use this option more than once to specify a large number of user IDs. If you do not specify group names with the groups option, or user IDs with the users option, all users can request Tivoli Storage Manager services. If you use the groups option, the users option, or both, only users included in one of the specified groups, or included in the list of users, can request Tivoli Storage Manager services.

Define your root user name only with the users option to exclude all other users from accessing the server.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file *within* a server stanza.

Syntax

```
        ┌─────────────────────────┐
        │        ┌───────────┐     │
►►──────▼─USERs──▼──username──┴─────┴──────────────────────────►◄
```

Parameters

username
Names a user that you want to authorize to request Tivoli Storage Manager services.

Examples

Options file:
```
    users  carol  larry  davecd  kathyba  michelle
    users  amyb  tkaspar  kbsmith  egray  srjames
```

Command line:
Does not apply.

V2archive

Use the v2archive option with the **archive** command to archive only files to the server.

Tivoli Storage Manager will not process directories that exist in the path of the source file specification.

This option differs from the filesonly option in that the filesonly option archives the directories that exist in the path of the source file specification.

The v2archive and dirsonly options are mutually exclusive and an error message is displayed if you use both options in the same **archive** command.

If you use this option, you might want to consider the following:
- You might experience performance problems when retrieving large amounts of data archived with this option.
- You might want to use this option only if you are concerned about expiration performance on a server that already contains extremely large amounts of archived data.
- If there are multiple files with the same name for the v2archive option, the files are archived multiple times, with their directory structure. The v2archive option archives only the files.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──V2archive──────────────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

This command:
>   dsmc archive "/Users/user2/Documents/*" -v2archive -su=y.

>   **Archives these files:**

>   /Users/user2/Documents/file1
>   /Users/user2/Documents/file2
>   /Users/user2/Documents/file3
>   /Users/user2/Documents/dir2/file4
>   /Users/user2/Documents/dir2/file5

>   Note: Tivoli Storage Manager does not archive /Users/user2/Documents and /Users/user2/Documents/dir2.

This command:
>   dsmc archive "/home/relx/dir1/*" -v2archive -su=y.

>   **Archives these files:**

>   /home/relx/dir1/file1
>   /home/relx/dir1/file2
>   /home/relx/dir1/file3
>   /home/relx/dir1/dir2/file4
>   /home/relx/dir1/dir2/file5

>   Note: Tivoli Storage Manager does not archive /home/relx/dir1 and /home/relx/dir1/dir2.

Verbose

The verbose option specifies that you want to display detailed processing information on your screen. This is the default.

When you run the **incremental**, **selective**, or **archive** commands, information is displayed about each file that is backed up. Use the quiet option if you do not want to display this information.

The following behavior applies when using the verbose and quiet options:
- If the server specifies either the quiet or verbose option in the server client option set, the server settings override the client values, even if **force** is set to *no* on the server.
- If you specify quiet in your dsm.opt file, and you specify *-verbose* on the command line, *-verbose* prevails.

- If you specify both *-quiet* and *-verbose* on the same command, the last option encountered during options processing prevails. If you specify *-quiet -verbose*, *-verbose* prevails. If you specify *-verbose -quiet*, *-quiet* prevails.

The information is displayed on your screen in the Scheduler Status window. This option only applies when you are running the scheduler and Tivoli Storage Manager is performing scheduled work.

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Command Line** tab, **Do not display process information on screen** checkbox of the Preferences editor.

Syntax

```
►►──VErbose───────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Options file:
        verbose

Command line:
        -verbose

This option is valid only on the initial command line. It is not valid in interactive mode.

Verifyimage

Use the verifyimage option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume.

If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log.

Supported Clients

This option is valid for AIX, HP-UX, all Linux clients, and Solaris *only*. The Tivoli Storage Manager client API does not support this option.

Syntax

```
►►──VERIFYImage──────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this option.

Examples

Command line:
>   dsmc restore image /usr -verifyimage

Virtualfsname

Use the *virtualfsname* option with the **backup group** command to specify the name of the virtual file space for the group on which you want to perform the operation. The *virtualfsname* cannot be the same as an existing file space name.

Supported Clients

This option is valid for all UNIX and Linux clients except for Mac OS X.

Syntax

>>──VIRTUALFsname =── ──*fsname*──><

Parameters

fsname
>   Specifies the name of the container for the group on which you want to perform the operation.

Examples

Command line:
>   backup group -filelist=/Users/van/Documents/filelist1 -groupname=group1
>   -virtualfsname=/virtfs -mode=full
>
>   backup group -filelist=/home/dir1/filelist1 -groupname=group1
>   -virtualfsname=/virtfs -mode=full

Virtualmountpoint

Authorized User: The virtualmountpoint option defines a virtual mount point for a file system if you want to consider files for backup that begin with a specific directory within that file system.

Using the virtualmountpoint option to identify a directory within a file system provides a direct path to the files you want to back up, saving processing time. It is more efficient to define a virtual mount point within a file system than it is to define that file system using the domain option, and then to use the exclude option in your include-exclude options list to exclude the files that you do not want to back up.

Use the virtualmountpoint option to define virtual mount points for multiple file systems, for local and remote file systems, and to define more than one virtual mount point within the same file system. Virtual mount points cannot be used in a file system handled by automounter.

Note: If the directory that you want to specify as a virtual mount point is a symbolic link, set the followsymbolic option to *Yes*. If that option is set to *no* (the default), you are not permitted to use a symbolic link as a virtual mount point. Also, if you back up a file system, then add a virtual mount point, and then do another incremental on the file system, the files and directories in the virtual mount point directory are expired, because they are logically contained within the virtual mount point directory and not the file system.

After you define a virtual mount point, you can specify the path and directory name with the domain option in either the default client options file or on the **incremental** command to include it for incremental backup services. When you perform a backup or archive using the virtualmountpoint option, the **query filespace** command lists the virtual mount point in its response along with other file systems. Generally, directories that you define as virtual mount points are treated as actual file systems and require that the virtualmountpoint option is specified in the `dsm.sys` file to restore or retrieve the data.

Note: When you specify a virtualmountpoint option, the path that it specifies is added to the default backup domain (domain **all-local**). The virtualmountpoint path is always considered a local "mount point" regardless of the real file system type it points to.

Supported Clients

This option is valid for all UNIX clients except Mac OS X. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza.

Syntax

```
►►─┬─VIRTUALMountpoint──┬──directory─┬──────────────────────►◄
   └──────────────────────────────────┘
```

Parameters

directory
> Specifies the path and directory name for the directory you want to use as the virtual mount point for a file system. You cannot use wildcard characters in either the path or directory names.
>
> Define only one virtual mount point with each virtualmountpoint option that you include in your client system-options file. Use the virtualmountpoint option as many times as necessary to define all of the virtual mount points that you want to use.

Examples

Options file:
```
virtualmountpoint /afs/xyzcorp.com/home/ellen
virtualmountpoint /afs/xyzcorp.com/home/ellen/test/data
```

Command line:
> Does not apply.

Virtualnodename

The virtualnodename option specifies the node name of your workstation when you want to restore or retrieve files to a different workstation.

When you use the virtualnodename option in your client options file, or with a command:

- You must specify the name you specified with the nodename option in your client system-options file (dsm.sys). This name should be different from the name returned by the **hostname** command on your workstation.
- Tivoli Storage Manager prompts for the password assigned to the node you specify, if a password is required. If you enter the correct password, you have access to all backups and archives that originated from the specified node.

When connecting to a server, the client must identity itself to the server. This login identification is determined in the following ways:

- If the nodename and virtualnodename options are not specified, or a virtual node name is not specified on the command line, the default login ID is the name returned by the **hostname** command.
- If the nodename option is specified, the name specified with the nodename option overrides the name returned by the **hostname** command.
- If the virtualnodename option is specified, or a virtual node name is specified on a command line, it cannot be the same name as the name returned by the **hostname** command.

When the virtual node name is accepted by the server, a password is required (assuming authentication is on), even if the passwordaccess option is *generate*. Once a connection to the server is established, then access is permitted to any file backed up using this login ID.

For the Java GUI, the virtualnodename option only works if you are root user. If you need to use the virtualnodename option with a non-root user account, use the command-line client (or native GUI, for platforms that still have one).

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client user-options file (dsm.opt).

Syntax

```
►►──VIRTUALNodename── ──nodename────────────────────────────────────►◄
```

Parameters

nodename
> Specifies a 1- to 64-character name that identifies the node for which you want to request Tivoli Storage Manager services. There is no default.

Examples

Options file:
>   virtualnodename cougar

Command line:
>   -virtualn=banshee

This option is valid only on the initial command line. It is not valid in interactive mode.

Webports

The webports option enables the use of the Web client outside a firewall.

Values for both the client acceptor daemon and the Web Client Agent service are required.

If you do not specify this option, the default value, zero (0), is used for both ports. This causes TCP/IP to randomly assign a free port number for the client acceptor daemon and the Web Client Agent service.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Web Client** tab, **WEB Ports** fields of the Preferences editor.

Syntax

```
►►──WEBPorts── ──cadport── ──agentport────────────────────────────────────►◄
```

Parameters

cadport
> Specifies the *required* Tivoli Storage Manager client acceptor daemon port number. The range of values is 1000 through 32767. If a value is not specified, the default, zero (0), causes TCP/IP to randomly assign a free port number.

agentport
> Specifies the *required* Tivoli Storage Manager Web client agent service port number. The range of values is 1000 through 32767. If a value is not specified, the default, zero (0), causes TCP/IP to randomly assign a free port number.

Examples

Options file:

>   For UNIX and Linux other than Mac OS X: webports 2123 2124
>
>   For Mac OS X: webports 2123

Command line:

>   Does not apply.

Chapter 11. Using commands

Tivoli Storage Manager provides a command-line interface (CLI) that you can use as an alternative to the graphical user interface (GUI). This topic describes how to start or end a client command session and how to enter commands.

The following is a list of tasks related to entering commands.
- "Start and end a client command session" on page 459
- "Enter client command names, options, and parameters" on page 461
- "Wildcard characters" on page 464

The following table provides an alphabetical list of the commands and a brief description.

Table 93. Commands

Command	Description
archive "Archive" on page 465	Archives files from a workstation to Tivoli Storage Manager storage.
backup fastback "Backup FastBack" on page 469	Backs up volumes specified by the fbpolicyname, fbclientname and fbvolumename options for long term retention.
backup group "Backup Group" on page 471	Creates and backs up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.
backup image "Backup Image" on page 473	Creates an image backup of one or more file systems or logical volumes that you specify.
backup nas "Backup NAS" on page 479	Creates an image backup of one or more file systems belonging to a Network Attached Storage (NAS) file server.
cancel process "Cancel Process" on page 481	Displays a list of current NAS (if NDMP support is enabled) image backup and restore processes for which the administrative user has authority.
cancel restore "Cancel Restore" on page 481	Displays a list of restartable restore sessions from which you can select one to cancel.
delete access "Delete Access" on page 482	Deletes authorization rules for files that are stored on the server.
	On those clients that support image backup, this command deletes authorization rules for images that are stored on the server.
delete archive "Delete Archive" on page 483	Deletes archived files from Tivoli Storage Manager server storage.
delete backup "Delete Backup" on page 484	Deletes active and inactive backup files from Tivoli Storage Manager server storage.
delete filespace "Delete Filespace" on page 487	Deletes file spaces in Tivoli Storage Manager server storage.
delete group "Delete Group" on page 488	Deletes a group backup on the Tivoli Storage Manager server.

Table 93. Commands (continued)

Command	Description
expire "Expire" on page 490	Inactivates backup objects that you specify in the file specification or with the filelist option.
help "Help" on page 491	Displays a Table of Contents of help topics for the command-line client.
incremental "Incremental" on page 492	Backs up all new or changed files or directories in the default client domain or from file systems, directories, or files you specify, unless you exclude them from backup services.
loop "Loop" on page 498	Starts an interactive command session.
macro "Macro" on page 500	Executes commands within a macro file that you specify.
monitor process "Monitor Process" on page 500	Displays a list of current NAS image backup and restore processes from which you can select one to cancel.
preview archive "Preview Archive" on page 501	Simulates an archive command without sending data to the server.
preview backup "Preview Backup" on page 502	Simulates a backup command without sending data to the server.
query access "Query Access" on page 503	Displays a list of current authorization rules.
query archive "Query Archive" on page 504	Displays a list of archived files.
query backup "Query Backup" on page 506	Displays a list of backup versions.
query backupset "Query Backupset" on page 508	Queries a backup set from a local file or the Tivoli Storage Manager server. On those clients that support tape devices, this command can query a backup set from a tape device.
query filespace "Query Filespace" on page 512	Displays a list of file spaces in Tivoli Storage Manager storage. You can also specify a single file space name to query.
query group "Query Group" on page 513	Displays information about group backups and their members.
query image "Query Image" on page 515	Displays information about image backups.
query inclexcl "Query Inclexcl" on page 517	Displays a list of include-exclude statements in the order in which they are processed during backup and archive operations.
query mgmtclass "Query Mgmtclass" on page 518	Displays information about available management classes.
query node "Query Node" on page 518	Displays all the nodes for which an administrative user ID has authority to perform operations.
query options "Query Options" on page 519	Displays all or part of your options and their current settings.
query restore "Query Restore" on page 520	Displays a list of your restartable restore sessions in the server database.
query schedule "Query Schedule" on page 521	Displays information about scheduled events for your node.

Command	Description
query session "Query Session" on page 521	Displays information about your session, including the current node name, when the session was established, server information, and server connection information.
query systeminfo "Query Systeminfo" on page 522	Gathers Tivoli Storage Manager system information and outputs this information to a file or the console.
restart restore "Restart Restore" on page 524	Displays a list of restartable restore sessions from which you can one to restart.
restore "Restore" on page 525	Restores copies of backup versions of your files from a Tivoli Storage Manager server.
restore backupset "Restore Backupset" on page 529	Restores a backup set from the Tivoli Storage Manager server or a local file. On those clients that support tape devices, this command can restore a backup set from a tape device.
restore group "Restore Group" on page 537	Restores specific members or all members of a group backup.
restore image "Restore Image" on page 539	Restores a file system or raw volume image backup.
restore nas "Restore NAS" on page 542	Restores the image of a file system belonging to a Network Attached Storage (NAS) file server.
retrieve "Retrieve" on page 544	Retrieves copies of archived files from the Tivoli Storage Manager server.
schedule "Schedule" on page 546	Starts the client scheduler on the workstation.
selective "Selective" on page 548	Backs up selected files.
set access "Set Access" on page 551	Authorizes another user to access your backup versions or archived copies. On those clients that support image backup, this command can set authorization rules for images that are stored on the server.
set event "Set Event" on page 553	Allows you to specify the circumstances for when archived data is deleted.
set password "Set Password" on page 555	Changes the Tivoli Storage Manager password for your workstation.

For proper operation, the was node must be restored to the same location and under the same name.

Important: To avoid problems, restore your data at the Network Deployment Manager node or Application Server node level only.

Related information

Reading syntax diagrams

Start and end a client command session

You can start or end a client command session in either batch mode or interactive mode.

Use batch mode when you want to enter a *single* client command. Tivoli Storage Manager processes the command and returns to the command prompt.

Use interactive mode when you want to enter a *series* of commands. Since Tivoli Storage Manager establishes connection to the server only once for interactive mode, a series of commands can be processed more quickly. Tivoli Storage Manager processes the commands and returns to the **tsm>** prompt.

Process commands in batch mode

Some options are valid *only* on the initial command line and not in interactive mode. These options generally affect the operation of the entire session.

For example, the command **dsmc query session -errorlogname=myerror.log** is accepted and it does name the error log. However, it is accepted simply because it appears in the initial command, even though the option is not valid for the query command.

There are also some options that are always valid on the initial command line as well as on individual commands in interactive mode. Therefore, certain options are accepted on the initial command line even though they have no effect on the command being entered. For example, **dsmc query session -subdir=yes** is a valid command, but in this case the *-subdir* option has no effect on the command that was entered.

When you enter a *single* command in batch mode, precede it with the executable program name, **dsmc**. For example, to process the **incremental** command in batch mode, you would enter:

    dsmc incremental

Tivoli Storage Manager prompts you each time you enter a command if the passwordaccess option is set to *prompt* and authentication on the server is set to *On*. Type your password and press Enter.

You can also enter your password using the password option with a command, but your password appears on the screen. For example, if your password is *secret*, enter:

    dsmc incremental –password=secret

If you set the *passwordaccess* option to *generate* in your dsm.opt file, you do not need to specify the password with the command. Tivoli Storage Manager *only* prompts you for your password if you are registering your workstation with a server or manually changing your password.

Related concepts

Chapter 10, "Processing options," on page 241

Process commands in interactive mode

Use the *interactive* mode (or *loop* mode) to enter a series of commands.

Enter **dsmc** on the command line and press Enter. When the **tsm>** command prompt appears, type the command name and press Enter. *Do not* precede each command with the executable program name, **dsmc**. Alternatively, you can enter **dsmc loop** on the command line to start a client command session in interactive mode. **Loop** is the default command for **dsmc**.

If a password is required, Tivoli Storage Manager prompts you before you enter the first command.

Type your user ID and password and press Enter.

You can also enter your password using the password option with the **loop** command, but your password appears on the screen. For example, if your password is *secret*, enter:

```
dsmc loop –password=secret
```

To end an interactive session, enter **quit** at the prompt.

Note for UNIX and Linux clients:

In loop mode, following a restore operation directly from tape, the mount point is not released in case additional restore requests are made to that volume. If you request a backup operation in the same session and that mount point is the only one available, the backup operation will stop with the following message:

```
Waiting for mount of offline media
```

In this case, the mount point is not released until one of the following conditions is met:
- The device class MOUNTRETENTION limit is satisfied.
- The client idletimeout period is satisfied.
- The dsmc loop session is closed after the restore operation completes, allowing you to start a subsequent loop mode session to perform the backup operation.

Enter client command names, options, and parameters

A client command can include one or more of these components: *Command name*, *options*, and *parameters*. The topics that follow describe each of these components.

Command name

The first part of a command is the command name. The command name consists of a single word, such as **help** or **schedule**, or an action word and an object for that action, such as **query archive**.

Enter the full command name, or its minimum abbreviation.

For example, you can enter any of the following versions of the **query schedule** command:

```
query schedule
q sc
q sched
query sc
```

Options

When you enter options with a command, always precede the option with a dash (–). Do not put a space between the dash and the option name.

Enter more than one option in any order in a command before or after the file specification. Separate multiple options with a blank space.

There are two groups of options that you can use with commands: Client options (set in your options file), or client command options (used on the command line).

- **Client options**: The group of options that are set in your client options file. You can override an option in the client options file when you enter the option with a command on the command line.
- **Client command options**: Use a client command option *only* when you enter the option with a command on the command line. You cannot set these options in an options file.

Related concepts

"Client options reference" on page 267

Options in interactive mode

In interactive mode, options that you enter on the initial command line override the value that you specified in your options file.

This value remains in effect for the entire interactive session unless overridden by a different value on a given interactive command.

For example, if you set the subdir option to *yes* in your dsm.opt or dsm.sys file, and you specify subdir=*no* on the initial command line, the subdir=*no* setting remains in effect for the entire interactive session unless overridden by the subdir=*yes* value on a given interactive command. However, the subdir=*yes* value specified within the interactive session only affects the command on which it is entered. When that command completes, the value reverts back to subdir=*no*, the value at the beginning of the interactive session.

Parameters

Commands can have required parameters, optional parameters, or no parameters at all.

Required parameters provide information to perform a task. The most commonly required parameter is a file specification.

For example, if you want to archive a file named `budget.fin` from the `project` directory, you would enter the following:

```
dsmc archive /project/budget.fin
```

Some commands have optional parameters. If you do not enter a value for an optional parameter, Tivoli Storage Manager uses the default value. For example, the **restore** command includes a required parameter, **sourcefilespec**, that specifies the path and file name in storage that you want to restore. The optional parameter, **destinationfilespec**, specifies the path where you want to place the restored files. If you do not specify the **destinationfilespec**, by default Tivoli Storage Manager restores the files to the original source path. If you want to restore the files to a *different* directory, enter a value for **destinationfilespec**.

Example: Restore the file `/project/budget.fin` to the new path `/newproj/newbudg.fin`

```
dsmc restore /project/budget.fin /newproj/
```

Enter parameters in the order indicated in the command syntax diagram.

File specification syntax

There are some syntax rules that you need to know about when entering file specification parameters, such as **filespec**, **sourcefilespec**, and **destinationfilespec**.

The following are the syntax rules:

- Do not use wildcards as part of the file space name or anywhere in the **destinationfilespec**. The one exception to this rule is the **set access** command where wildcards are permitted in the two lowest levels of the file spec.

  Example: Allow access to all files in all directories in and below the /home directory:

  ```
  set access backup /home/* * *
  set access backup /home/*/* * *
  ```

  With UNIX clients, do not use wildcards for the directory path name, for example:

  ```
  /home/j*asler/file1.c
  ```

- There is a maximum number of file specifications per command:
  - The **Query** commands can accept only one file specification.
  - The **restore** and **retrieve** commands can accept a source file specification and a destination file specification.
  - There is a limit of 20 operands on some commands. This is to prevent excessive sessions that are caused when wildcards are expanded by the UNIX shell command processor.

    You can prevent shell expansion from causing you to go over the 20-operand limit by placing quotation marks around your source filespec expansion characters for restore commands.

    Note: Using quotation marks has the side affect of causing a no-query restore.

    You can use the removeoperandlimit option to specify that Tivoli Storage Manager removes the 20-operand limit. If you specify the removeoperandlimit option with the **incremental**, **selective**, or **archive** commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits.

- The length of a file specification is limited.
  - The maximum number of characters for a file name is 255. The maximum combined length of the file name and path name is 1024 characters. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.

- When entering the **sourcefilespec**, if the directory name ends with \, then * is implied.

  When entering a **destinationfilespec**, if the name ends with /, then it is considered a directory, otherwise it is considered a file.

  The following example illustrates these two rules. Even though `mydir` and `yourdir` are directories, the command will fail because /* is implied after `mydir`, and `yourdir` is considered a file:

  ```
  restore /home/mydir/ /away/yourdir
  ```

  The following example illustrates the second rule. Even though `mydir` and `yourdir` are directories, the command will fail because `mydir` and `yourdir` are considered files:

  ```
  restore /home/mydir /away/yourdir
  ```

- If a file specification does not begin with a file space name (an opening directory delimiter), the file specification is assumed to be a subdirectory of the current working directory. The client appends the destination file specification to the working directory to build the complete path.

  For example, if the current working directory is /home/me and the command is `dsmc res "/fs/dir1/*" mydir/`, the complete restore path is this: /home/me/mydir.

- The only command that accepts a simple file space name is the **incremental** command. The following example is valid:

      dsmc i /Users

  The following example is not valid, because the command is the **selective** command:

      dsmc sel /Users

Related reference

Wildcard characters

Use wildcard characters when you want to specify multiple files with similar names in *one* command. Without wildcard characters, you must repeat the command for each file.

In a command, you can use wildcard characters in the file name or file extension *only*. You cannot use them to specify destination files, file systems, or server names. You cannot specify a directory whose name contains an asterisk (*) or a question mark (?).

Valid wildcard characters that you can use include:

*   Asterisk. Matches zero or more characters.

?   Question mark. Matches any single character at the present position.

The following table shows examples of each wildcard.

Table 94. Wildcard characters

Pattern	Matches	Does not match
Asterisk (*)		
ab*	ab, abb, abxxx	a, b, aa, bb
ab*rs	abrs, abtrs, abrsrs	ars, aabrs, abrss
ab*ef*rs	abefrs, abefghrs	abefr, abers
abcd.*	abcd.c, abcd.txt	abcd, abcdc, abcdtxt
Question Mark (?)		
ab?	abc	ab, abab, abzzz
ab?rs	abfrs	abrs, abllrs
ab?ef?rs	abdefjrs	abefrs, abdefrs, abefjrs
ab??rs	abcdrs, abzzrs	abrs, abjrs, abkkkrs

Important: Use an asterisk (*) instead of a question mark (?) as a wildcard character when trying to match a pattern on a multibyte code page, to avoid unexpected results.

Note: In batch mode, enclose values containing wildcards in quotation marks. Otherwise, UNIX shells expand unquoted wildcards, and it is easy to exceed the 20 operand limit. It is more efficient to let the client process wildcard file specifications because many fewer server interactions are needed to complete the task. For example:

```
dsmc selective "/home/me/*.c"
```

Client commands reference

The following sections contain detailed information about each of the Tivoli Storage Manager commands.

Information for each command includes the following information:
- A description of the command.
- A syntax diagram of the command.
- Detailed descriptions of the command parameters. If the parameter is a constant (a value that does not change), the minimum abbreviation appears in uppercase letters.
- Examples of using the command.

Archive

The **archive** command archives a single file, selected files, or all files in a directory and its subdirectories on a server.

Archive files that you want to preserve in their current condition. To release storage space on your workstation, delete files as you archive them using the deletefiles option. Retrieve the archived files to your workstation whenever you need them again.

Use the snapshotroot option with the **archive** command in conjunction with a independent software vendor application that provides a snapshot of a logical volume, to associate the data the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. The snapshotroot option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot.

AIX only: You can enable snapshot-based file archive by using the option snapshotproviderfs=*JFS2*.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Archive──┬──filespec──┬──────────────────►◄
             └◄─────────┘  └──options──┘
```
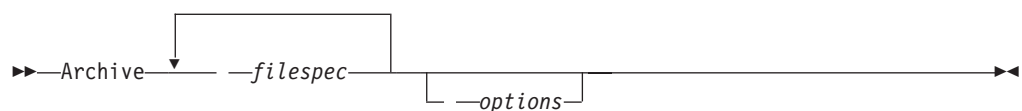
Parameters

filespec

Specifies the path and name of the file you want to archive. You can use wildcards to specify groups of files or all the files in a directory. Separate file specifications with a space. You can also use the filelist option to process a list of files. The Tivoli Storage Manager client opens the file that you specify with this option and processes the list of files within according to the specific command.

There is a limit of 20 operands. This is to prevent excessive sessions that are caused when wildcards are expanded by the UNIX shell command processor. You can prevent shell expansion from causing you to go over the 20-operand limit by placing quotation marks around your source filespec-expansion characters.

Tip: Using quotation marks might cause a no-query restore.

You can use the removeoperandlimit option to specify that Tivoli Storage Manager removes the 20-operand limit. If you specify the removeoperandlimit option, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits.

Table 95. Archive command: Related options

Option	Where to use
archmc	Command line only.
archsymlinkasfile	Client user-options file (dsm.opt) or command line.
autofsrename	Client options file (dsm.opt) only.
changingretries	Client system options file or command line.
compressalways	Client user-options file (dsm.opt) or command line.
compression	dsm.sys file *within* a server stanza or command line.
deletefiles	Command line only.
description	Command line only.
dirsonly	Command line only.
encryptiontype	dsm.sys file *within* a server stanza.
encryptkey	dsm.sys file *within* a server stanza.
filelist	Command line only.
filesonly	Command line only.
preservelastaccessdate	Client user-options file (dsm.opt) or command line.
removeoperandlimit	Command line only.
snapshotroot	Command line only.
subdir	Client options file (dsm.opt) or command line.
tapeprompt	Client options file (dsm.opt) or command line.
v2archive	Command line only.

Examples

Task   Archive a single file named `budget` in the `/home/proj1` directory.

   Command: archive /home/proj1/budget

Task   Archive all files in the `/home/proj1` directory with a file extension of `.txt`.

Command: `archive "/home/proj1/*.txt"`

Task   Archive all files in the directory tree headed by the /home directory.

Command: `archive -subdir=yes "/home/*"`

Task   Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, archive the /usr/dir1/sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name /usr.

Command: `dsmc archive /usr/dir1/sub1/ -subdir=yes -snapshotroot=/snapshot/day1`

Related information

"File system and ACL support" on page 125

"Snapshotproviderfs" on page 423

Archive FastBack

Use the **archive fastback** command to archive Tivoli Storage Manager FastBack volumes specified by the fbpolicyname, fbclientname and fbvolumename options for long-term retention.

Before using this command, configure the client to back up and archive Tivoli Storage Manager FastBack data. Also, before you issue this command, at least one snapshot should exist in the Tivoli Storage Manager FastBack repository for the Tivoli Storage Manager FastBack policy being archived or backed up.

If a policy specification contains both Windows and Linux FastBack clients, only the Windows volumes will be backed up or archived to Tivoli Storage Manager by the Windows Tivoli Storage Manager client.

If a policy specification contains both Windows and Linux FastBack clients, only the Linux volumes will be backed up or archived to Tivoli Storage Manager by the Linux Tivoli Storage Manager client.

You can use Tivoli Storage Manager FastBack options to archive the latest snapshots of the following volumes:
* All clients and volumes associated with a FastBack policy or a list of FastBack policies.
* All volumes associated with a specific FastBack client or a list of FastBack clients for a given FastBack policy.
* A specific volume or volumes associated with a FastBack client for a given FastBack policy.

Supported Clients

This option is valid for Linux x86/x86_64 clients.

Syntax

```
>>──ARCHIVE FASTBack──FBPolicyname──=──▼──name──┴──FBServer──=──name──────────────>
                                    └──,──┘
```
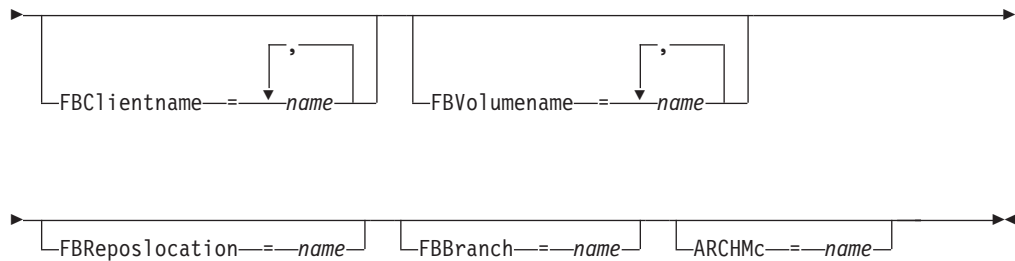
```
                                                                    ┌─────,─────┐                        ┌─────,─────┐
├──────────────────────────────────────┬─────────▼─────────┬────────────────────────▼─────────┬──────────────────────────┤
                 └─FBClientname──=──────name─┘                      └─FBVolumename──=──────name─┘


├───────────────────────────────┬──────────────────────┬────────────────────┬──────────────────────────────────────────────►◄
        └─FBReposlocation──=──name─┘      └─FBBranch──=──name─┘    └─ARCHMc──=──name─┘
```

Important:

1. *At least* one FBpolicyName is always *required*.
2. You can specify up to 10 values for FBPolicyName, if no values are specified for both FBClientName *and* FBVolumeName.
3. When you specify a FBClientName value, there must be *only one* value for FBPolicyName.
4. You can specify up to 10 values for FBClientName if *only one* PolicyName is specified, and *no* values for FBVolumeName are specified.
5. When you specify the FBVolumeName option, you can have *only one* FBPolicy, and *only one* FBCLientName specified.
6. You can specify multiple FBVolumeNames if condition 5 is satisfied.
7. You must always specify the FBReposLocation option for Linux.

Parameters

Table 96. Archive FastBack command: Related options

Option	Where to use
fbpolicyname "Fbpolicyname" on page 328	Command line and scheduler.
fbserver "Fbserver" on page 330	Command line and scheduler.
fbclientname "Fbclientname" on page 326	Command line and scheduler.
fbvolumename "Fbvolumename" on page 332	Command line and scheduler.
fbreposlocation "Fbreposlocation" on page 329	Command line and scheduler.
fbbranch "Fbbranch" on page 326	Command line and scheduler.
archmc "Archmc" on page 267	Command line and scheduler.

Examples

Command line:
> The Tivoli Storage Manager backup-archive client is installed on a Linux

proxy client machine. Use this command to archive all FastBack volumes
for all Linux FastBack clients that are defined for FastBack policy1:

```
dsmc archive fastback -fbpolicyname=Policy1
  -fbserver=myfbserver -fbreposlocation=myfbserver@WORKGROUP
```

The FastBack server name, -myFbDrHub is the short host name of the
FastBack disaster recovery hub server where the repository is located.

The -fbreposlocation parameter specifies the location of the repository. The
repository location is required. If you do not provide the repository
location, the command fails.

FBServer should point to the short host name of the FastBack DR hub in
this case.

Command line:
> The repository, rep_server1, is located on a FastBack disaster hub,
> myFbDrHub

```
dsmc archive fastback -fbpolicyname="Policy 1"
  -fbserver=myFbDrHub -fbreposlocation=\\myFbDrHub\rep_server1
```

The repository location is required. If you do not provide the repository
location, the command fails.

The FastBack server name, -myFbDrHub, is the short host name of the
FastBack disaster recovery hub where the repository is located.

FBServer should point to the short host name of the FastBack DR hub in
this case.

Command line:
> Archive all volumes protected by FastBack policy named policy1 from the
> FastBack server named basil:

```
dsmc archive fastback -Fbpolicyname=policy1
  -FBServer=basil -ARCHMC="my_tsm_mgmt_class"
  -fbreposlocation=basil@WORKGROUP
```

Backup FastBack

Use the **backup fastback** command to back up Tivoli Storage Manager FastBack
volumes specified by the fbpolicyname, fbclientname and fbvolumename options
for long-term retention.

Before using this command, configure the client to back up and archive Tivoli
Storage Manager FastBack data. Also, before you issue this command, at least one
snapshot should exist in the Tivoli Storage Manager FastBack repository for the
Tivoli Storage Manager FastBack policy being archived or backed up.

If a policy specification contains both Windows and Linux FastBack clients, only
the Linux volumes will be backed up or archived to Tivoli Storage Manager by the
Linux Tivoli Storage Manager client.

Tivoli Storage Manager FastBack options are supported for the incremental backup
of the latest snapshots, depending on the option specified:
* All clients and volumes associated with a Tivoli Storage Manager FastBack
  policy or a list of Tivoli Storage Manager FastBack policies.
* All volumes associated with a specific FastBack client or a list of FastBack clients
  for a given Tivoli Storage Manager FastBack policy.

- A specific volume or volumes associated with a FastBack client for a given Tivoli Storage Manager FastBack policy.

Supported Clients

This command is valid for Linux x86/x86_64 clients that are configured as Tivoli Storage Manager FastBack® dedicated proxies.

Syntax
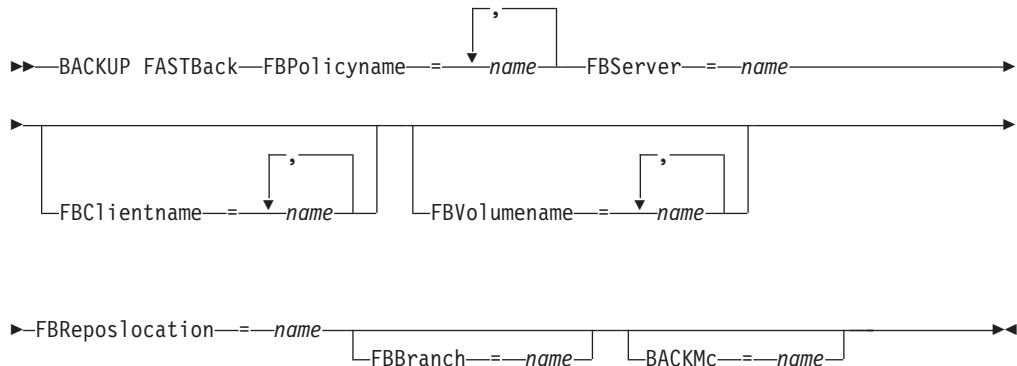
```
                                        ,
                                ┌───◄────┐
►►──BACKUP FASTBack──FBPolicyname──=──▼──name──┴──FBServer──=──name──────────────►

►─────────────────────────────────────────────────────────────────────────────►
                    ,                            ,
             ┌──◄───┐                     ┌───◄──┐
      └─FBClientname──=──▼──name──┘   └─FBVolumename──=──▼──name──┘

►──FBReposlocation──=──name──────────────────────────────────────────────────►◄
                        └─FBBranch──=──name──┘  └─BACKMc──=──name──┘
```

Important:

1. *At least* one FBpolicyName is always *required*.
2. You can specify up to 10 values for FBPolicyName, if no values are specified for both FBClientName *and* FBVolumeName.
3. When you specify a FBClientName value, there must be *only one* value for FBPolicyName.
4. You can specify up to 10 values for FBClientName if *only one* PolicyName is specified, and *no* values for FBVolumeName are specified.
5. When you specify the FBVolumeName option, you can have *only one* FBPolicy, and *only one* FBCLientName specified.
6. You can specify multiple FBVolumeNames if condition 5 is satisfied.
7. You must specify the FBReposLocation option.

Table 97. Backup FastBack command: Related options

Option	Where to use
fbpolicyname "Fbpolicyname" on page 328	Command line and scheduler.
fbserver "Fbserver" on page 330	Command line and scheduler.
fbclientname "Fbclientname" on page 326	Command line and scheduler.
fbvolumename "Fbvolumename" on page 332	Command line and scheduler.

Table 97. Backup FastBack command: Related options (continued)

Option	Where to use
fbreposlocation "Fbreposlocation" on page 329	Command line and scheduler.
fbbranch "Fbbranch" on page 326	Command line and scheduler.
backmc "Backmc" on page 277	Command line and scheduler.

Examples

Command line:

The Tivoli Storage Manager backup-archive client is installed on a Linux proxy client machine. Use this command to back up all FastBack volumes for all Linux FastBack clients that are defined for FastBack policy1:

```
dsmc backup fastback -fbpolicyname=Policy1
  -fbserver=myfbserver
  -fbreposlocation=myfbserver@WORKGROUP
```

The repository location is required. If you do not provide the repository location, the command will fail.

The FastBack server name, -myfbserver, is the short host name of the FastBack server where the repository is located.

Command line:

The repository, rep_server1, is located on a FastBack disaster hub, myFbDrHub

```
dsmc backup fastback -fbpolicyname="Policy 1"
  -fbserver=myFbDrHub -fbreposlocation=\\myFbDrHub\rep_server1
```

The FastBack server name, -myFbDrHub, is the short host name of the FastBack disaster recovery hub server where the repository is located.

The -fbreposlocation option specifies the location of the repository. The repository location is required. If you do not provide the repository location, the command fails.

The FBServer option should point to the short host name of the FastBack DR hub in this case.

Command line:

Back up all volumes protected by FastBack policy named policy1 from the FastBack server named basil:

```
dsmc backup fastback -Fbpolicyname=policy1
  -FBServer=basil -BACKMC="my_tsm_mgmt_class"
  -fbreposlocation=basil@WORKGROUP
```

Backup Group

Use the **backup group** command to create and back up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.

AIX only: You can enable snapshot-based group backup by using the option `snapshotproviderfs=JFS2` .

A group backup allows you to create a consistent point-in-time backup of a group of files that is managed as a single logical entity:

- All objects in the group are assigned to the same management class.
- Existing exclude statements for any files in the group are ignored.
- All objects in the group are exported together.
- All objects in the group are expired together as specified in the management class. No objects in a group are expired until all other objects in the group are expired, even when another group they belong to gets expired.
- If you are performing full and differential group backups to a sequential device, during a restore the data is in no more than two locations. To optimize restore time, perform periodic full backups to back up the data to one location on the sequential media.
- During a full group backup, all objects in the filelist are sent to the server. During a differential group backup, only data that has changed since the last full backup is sent to the server. Objects in the filelist that have not changed since the last full backup are assigned as members of the differential group backup. This data is not resent to the server, reducing backup time.

The **backup group** command requires the following options:

filelist Specifies a list of files to add to a new group.

groupname
> Specifies the fully qualified name of the group containing a list of files.

virtualfsname
> Specifies the name of the virtual file space for the group on which you want to perform the operation. The virtualfsname option cannot be the same as an existing file space name.

mode Specifies whether you want to back up all of the files in the filelist or only files that have changed since the last full backup.

Note:
1. *If any file in the group backup fails, the entire group backup fails.*
2. Use the **query group** command to query members of a group backup on the Tivoli Storage Manager server.
3. Use the **restore group** command to restore specific members or all members of a group backup on the Tivoli Storage Manager server.
4. Unless you are running Mac OS X, use the **delete group** command to delete a specific group backup from the Tivoli Storage Manager server.
5. Use the **query filespace** command to display virtual file space names for your node that are stored on the Tivoli Storage Manager server.
6. A group backup can be added to a backup set.

Supported Clients

This command is valid for all UNIX and Linux clients except Mac OS X.

Syntax

```
►►──Backup GRoup── ──options──────────────────────────────────────►◄
```

Parameters

Table 98. Backup Group command: Related options

Option	Where to use
filelist "Filelist" on page 333	Command line only.
groupname "Groupname" on page 341	Command line only.
mode "Mode" on page 369	Command line only.
snapshotproviderfs "Snapshotproviderfs" on page 423	System-options file (dsm.sys) within a server stanza or with the include.fs option.
virtualfsname "Virtualfsname" on page 452	Command line only.

Examples

Task   Perform a full backup of all the files in the `/home/dir1/filelist1` file to the virtual file space name `accounting` containing the group leader `/home/group1` file.

    Command:

```
backup group -filelist=/home/dir1/filelist1 -groupname=group1
-virtualfsname=/virtfs -mode=full
```

Related information

"Include options" on page 351

"Query Group" on page 513

"Restore Group" on page 537

"Delete Group" on page 488

"Query Filespace" on page 512

Backup Image

The **backup image** command creates an image backup of one or more volumes on your system.

These volumes can be formatted FAT, FAT32, NTFS, or unformatted RAW volumes. If a volume is NTFS-formatted, only those blocks used by the file system are backed up.

If you set the imagegapsize option to 0, all blocks, including unused blocks at the end of the volume, are backed up.

If you specify an AIX JFS2 file system for image backup, only those blocks used by the file system are backed up. If you set the imagegapsize option to zero, all blocks, including blocks at the end of the volume, are backed up.

Note:

1. AIX only: By default, snapshot-based image backup is enabled for JFS2 volumes. You can disable this by using the option `snapshotproviderimage=NONE` .

2. For the Linux clients, image backup is only supported on partitions with id 0x83 or logical volumes created with the Linux Logical Volume Manager. Backing up other partitions, such as extended partitions that contain mounted file systems or database data, can produce inconsistent backup data if the data changes during the image backup operation.

3. Backup image is not supported on any GPFS file system.

4. The Tivoli Storage Manager API must be installed to use the **backup image** command.

5. When you change the attribute of a JFS2 file system to an HSM-managed file system, an image backup is not done for that file system.

Important:

The last incremental backup time refers to the server time and the file modification time refers to the client time. If the client and server time are not synchronized, or the client and server are in different time zones, this affects incremental-by-date backup and image backup with *mode=incremental*.

The client backs up the files whose modification date and time (client time) is later than the date and time of the last incremental backup of the file system on which the file resides (server time).

If the server time is ahead of the client time, an incremental-by-date operation or image backup with *mode=incremental* skips the files, which had been created or modified after the last incremental or image backup with a modification date earlier than the last incremental backup time stamp.

If the client time is ahead of the server time, all files that had been created or modified before the last incremental or image backup and have a modification time stamp later than the last incremental backup time stamp, are backed up again. Typically, these files would not get backed up because they had already been backed up.

The backup date can be checked by the **query filespace** command.

The Tivoli Storage Manager client must support the raw device type on the specific platform to perform an image backup of a raw device. You can only perform an image backup on local devices. Clustered devices or file systems as well as devices or file systems shared between two or more systems are not supported. If you want to perform an image backup for a file system mounted on a raw device, the raw device must be supported.

Use the include.image option to include a file system or logical volume for image backup, or to specify volume-specific options for image backup.

The **backup image** command uses the compression option.

Supported Clients

This command is valid for AIX, HP-UX, all Linux clients, and Solaris.

Syntax

```
►►──Backup Image───────────────────────────────────────────────►◄
                   └─ ─options─┘  └─ ─filespec─┘
```
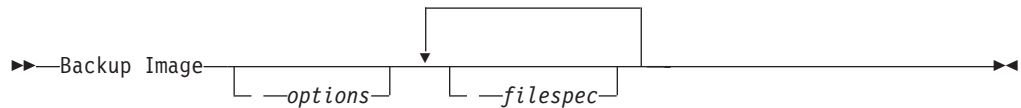
Parameters

filespec

Specifies the name of one or more logical volumes. If you want to back up more than one file system, separate their names with spaces. Do not use pattern matching characters. If you do not specify a volume name, the logical volumes specified with the domain.image option are processed. If you do not use the domain.image option to specify file systems to process, an error message is displayed and no image backup occurs.

Specify the file space over which the logical volume is mounted or the logical volume name. If there is a file system configured in the system for a given volume, you cannot back up the volume with the device name.

For example, if the /dev/lv01 file space is mounted on the /home volume, you can issue backup image /home, but backup image /dev/lv01 file space fails with an error: ANS1063E Invalid path specified.

Note: For Sun systems, specify either a file system name or a raw device name (block device type).

Table 99. Backup Image command: Related options

Option	Where to use
asnodename	Client options file (dsm.opt) or command line.
asnodename "Asnodename" on page 268	Client system options file (dsm.sys) or command line.
compressalways "Compressalways" on page 284	Client system options file (dsm.sys) or command line.
compression "Compression" on page 285	Client options file or command line.
dynamicimage "Dynamicimage" on page 309	Use with the **backup image** command or the include.image option in the options file.
imagegapsize "Imagegapsize" on page 347	Use with the **backup image** command, the include.image option, or in the options file.
mode "Mode" on page 369	Command line only.
postsnapshotcmd "Postsnapshotcmd" on page 386	Use with the **backup image** command, the include.image option , or in the options file.
presnapshotcmd "Presnapshotcmd" on page 392	Use with the **backup image** command, the include.image option , or in the options file.

Table 99. Backup Image command: Related options (continued)

Option	Where to use
snapshotcachesize "Snapshotcachesize" on page 422	Use with the **backup image** command, the include.image option , or in the options file.
snapshotproviderimage "Snapshotproviderimage" on page 424	Client options file or with include.image option.

Examples

Task   Back up the /home/test file space over which the logical volume is mounted and perform an image incremental backup that backs up only new and changed files after the last full image backup.

Command: `dsmc backup image /home/test -mode=incremental`

Task   Perform a static image backup of the logical volume mounted at the /home directory.

Command: `Command: dsmc backup image /home -snapshotproviderimage=none`

Task   Perform a dynamic image backup of the logical volume mounted at the /home directory.

Command: `dsmc backup image /home -dynamicimage=yes`

Task   Perform a snapshot image backup of the /home directory.

Command:

```
AIX client: dsmc backup image /home
    -snapshotproviderimage=JFS2
 LINUX client: dsmc backup image /home
    -snapshotproviderimage=LINUX_LVM
```

Task   Back up the /dev/lv01 raw logical volume.

Command: `dsmc backup image f:`

Related information

"Imagegapsize" on page 347

"Snapshotproviderimage" on page 424

"Snapshotcachesize" on page 422

"Mode" on page 369

"Comparing methods 1 and 2" on page 154 to decide which method is appropriate for your environment.

Static, dynamic, and snapshot image backup

The traditional image backup prevents write access to the volume by other system applications during the operation.

Use the dynamicimage option to back up the volume as is without remounting it read-only. Corruption of the backup can occur if applications write to the volume while the backup is in progress. In this case, run **fsck** after a restore.

The dynamicimage option is not supported for JFS2 volumes.

For Linux x86/x86_64 clients only: By default, Tivoli Storage Manager performs a snapshot image backup of file systems residing on a logical volume created by the Linux Logical Volume Manager during which the volume is available to other system applications. Snapshot image backup requires a Version 5.1 Tivoli Storage Manager server.

For AIX clients only: By default, Tivoli Storage Manager performs a snapshot image backup of JFS2 volumes during which the volume is available to other system applications. AIX allows the creation of a snapshot of a JFS2 volume while it is still online. The snapshot is created inside the same volume group as the source volume. You must ensure that the volume group provides enough free disk space to create the snapshot. The snapshot contains the old data blocks while the modified data is stored in the source volume. Use the snapshotcachesize option with the backup image command, in the dsm.sys file, or with the include.image option to specify an appropriate snapshot size so that all old data blocks can be stored while the image backup occurs.

The Linux Logical Volume Manager allows the creation of a snapshot of a logical volume while the logical volume itself is still online. The snapshot is created inside the same volume group as the source logical volume. You must ensure that the volume group provides enough free disk space to create the snapshot. The snapshot contains the old data blocks while the modified data is stored in the source logical volume. Use the snapshotcachesize option with the **backup image** command, in the dsm.sys file, or with the include.image option to specify an appropriate snapshot size so that all old data blocks can be stored while the image backup occurs. A snapshot size of 100 percent will ensure a valid snapshot.

Utilizing image backup to perform file system incremental backup

There are two methods of utilizing image backups to perform efficient incremental backups of your file system. These backup methods allow you to perform point-in-time restore of your file systems and improve backup and restore performance.

You can perform the backup only on formatted volumes; not on raw logical volumes. You can either use *image backup with file system incremental* or you can use *image backup with image incremental mode* to perform image backups of volumes with mounted file systems.

The following are some examples of using *image backup with file system incremental*.
- To perform a full incremental backup of the file system: `dsmc incremental /myfilesystem`
- To perform an image backup of the same file system: `dsmc backup image /myfilesystem`
- To periodically perform incremental backups: `dsmc incremental /myfilesystem`

You must follow the next steps in the order shown to ensure that the server records additions and deletions accurately.

Use this command to restore the file system to its exact state as of the last incremental backup: `dsmc restore image /myfilesystem -incremental -deletefiles`.

During the restore, the client does the following:
- Restores the most recent image on the server.
- Deletes all of the files restored in the previous step which are inactive on the server. These are files which existed at the time of the image backup, but were subsequently deleted and recorded by a later incremental backup.
- Restores new and changed files from the incremental backups.

If you do not follow the steps exactly, two things can occur:
1. After the original image is restored, all files backed up with the **incremental** command are restored individually.
2. If you perform a **backup image** before performing an **incremental**, files deleted from the original image are *not* deleted from the final restored file system.

The following are some examples of using *image backup with image incremental mode*.
- To perform an image backup of the same file system: `dsmc backup image /myfilesystem`
- To perform an incremental image backup of the file system: `dsmc backup image /myfilesystem -mode=incremental`

  This sends only those files that were added or changed since the last image backup to the server.
- To periodically perform full image backups: `dsmc backup image /myfilesystem`
- To restore the image: `dsmc restore image /myfilesystem -incremental`

  On restore, Tivoli Storage Manager ignores the deletefiles option when the image+image incremental technique of backing up has been used. The restore will include files that were deleted after the last full image backup plus the latest versions of files added or changed after the last image backup.

Note: You should perform full image backups periodically in the following cases. This will improve restore time because fewer changes are applied from incrementals.
- When a file system changes substantially (more than 40%).
- Once each month.
- As appropriate for your environment.

The following restrictions apply when using the image backup with image incremental mode:
- The file system can have no previous full incremental backups produced by the **incremental** command.
- Incremental-by-date image backup does not inactivate files on the server; therefore, when files are restored, none can be deleted.
- If this is the first image backup for the file system, a full image backup is performed.
- Using mode=*incremental* backs up only files with a changed date, not files with changed permissions.
- If file systems are running at or near capacity, an out-of-space condition could result during the restore.

Backup NAS

The **backup nas** command creates an image backup of one or more file systems belonging to a Network Attached Storage (NAS) file server, otherwise known as NDMP Backup.

The NAS file server performs the outboard data movement. A server process starts in order to perform the backup.

Use the nasnodename option to specify the node name for the NAS file server. When using an interactive command-line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID. The NAS node name identifies the NAS file server to the Tivoli Storage Manager server; the NAS node name must be registered at the server. Place the nasnodename option in your client options file (dsm.opt). The value in the client options file is the default, but can be overridden on the command line.

Use the toc option with the **backup nas** command or the include.fs.nas option to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup. If you save TOC information, you can use the **QUERY TOC** server command to determine the contents of a file system backup in conjunction with the **RESTORE NODE** server command to restore individual files or directory trees.

You can also use the Tivoli Storage Manager Web client to examine the entire file system tree and select files and directories to restore. Creation of a TOC requires that you define the **tocdestination** attribute in the backup copy group for the management class to which this backup image is bound. TOC creation requires additional processing, network resources, storage pool space, and possibly a mount point during the backup operation. If you do not save TOC information, you can still restore individual files or directory trees using the RESTORE NODE server command, if you know the fully qualified name of each file or directory and the image in which that object was backed up.

The toc option is only supported for images backed up by Version 5.2 or later client and server.

Specifying mode =*differential* on the **BACKUP NODE** server command or the **backup nas** command where no full image exists, shows that a full backup was started. Using the **QUERY PROCESS** server command shows that a full backup is in process.

Use the mode option to specify whether to perform a full or differential NAS image backup. A full image backup backs up the entire file system. The default is a differential NAS image backup on files that change after the last full image backup. If an eligible full image backup does not exist, a full image backup is performed. If a full image exists, whether it is restorable, or expired and being maintained because of dependent differential images, specifying mode =*differential* sends a differential image backup. If a full image is sent during a differential backup, it is reflected as a full image using the **QUERY NASBACKUP** server command. The **QUERY NASBACKUP** server command also displays NAS images that are restorable and displays full image or differential image as the object type.

Use the monitor option to specify whether you want to monitor a NAS file system image backup and display processing information on your screen.

Use the **monitor process** command to display a list of all processes for which an administrative user ID has authority. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web.

Use the **cancel process** command to stop NAS backup processing.

Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.

Supported Clients

This command is valid for AIX, and Solaris clients only.

Syntax
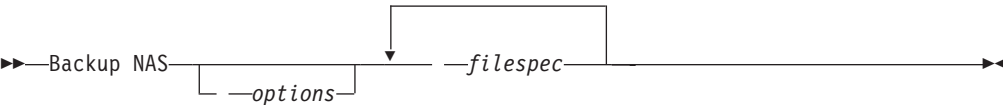
Parameters

filespec

Specifies the name of one or more file systems on the NAS file server. If you do not specify this parameter, Tivoli Storage Manager processes all of the file systems defined by the domain.nas option.

If you do not specify the *filespec* or the domain.nas option, the default **all-nas** value is used for domain.nas and all file systems on the NAS file server are backed up.

Table 100. Backup NAS command: Related options

Option	Where to use
mode "Mode" on page 369	Command line only.
monitor "Monitor" on page 371	Command line only.
nasnodename "Nasnodename" on page 371	Client options file (dsm.sys) or command line.
toc "Toc" on page 444	Command line or with the include.fs.nas option in your client options file (dsm.sys).

Examples

Task   Perform the NAS image backup of the entire file system.

   Command: `backup nas -mode=full -nasnodename=nas1 /vol/vol0 /vol/vol2`

Task   Perform the NAS image backup of the entire file server.

   Command: `backup nas -nasnodename=nas1`

Task   Perform the NAS image backup of the entire file system and save Table of Contents (TOC) information for the file system backup.

Command: `backup nas -mode=full -nasnodename=netappsj /vol/vol0`
`-toc=yes`

Related information

"Nasnodename" on page 371

"Toc" on page 444

"Mode" on page 369

"Monitor" on page 371

"Cancel Process"

"Domain.nas" on page 308

Cancel Process

The **cancel process** command displays a list of current NAS (if NDMP support is enabled) image backup and restore processes for which the administrative user has authority.

From the list, the administrative user can select one process to cancel. Client owner privilege is sufficient authority to cancel the selected NAS image backup or restore processes.

When using an interactive command-line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID.

Supported Clients

This command is valid for AIX, Linux, and Solaris clients only.

Syntax

►►—Cancel Process—————————————————————————————►◄

Parameters

There are no parameters for this command.

Examples

Task   Cancel current NAS image backup or restore processes.

> **Command:** `cancel process`

Cancel Restore

The **cancel restore** command displays a list of your restartable restore sessions in the server database.

You can only cancel one restartable restore session at a time. Run the **cancel restore** command again to cancel additional restores. To restart restartable restore sessions, use the **restart restore** command.

Use the **cancel restore** command under the following circumstances:
- You cannot back up files affected by the restartable restore.
- Restartable restore sessions lock the file space so that files cannot be moved off of the sequential volumes of the server.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Cancel Restore──────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this command.

Examples

Task   Cancel a restore operation.

   Command:
   ```
   cancel restore
   ```

Delete Access

The **delete access** command deletes authorization rules for files that are stored on the server.

When you delete an authorization rule, you revoke user access to any files or images specified by that rule.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Delete── ──ACcess──────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this command.

Examples

Task   Display a list of current authorization rules and select the rules you want to delete.

   Command:
   ```
   delete access
   ```

   See the following screen example:

```
         Index   Type       Node    Owner    Path

         1       Backup     NODE1   USER1    home/dev/proja/list/
         2       Archive    NODE3   LUIE     home/fin/budg/depta/
         3       Backup     NODE4   USER2    home/plan/exp/deptc/
         4       Archive    NODE5   USER2S   home/mfg/invn/parta/
         Enter Index of rule(s) to delete, or quit to cancel:
```

To delete the authorization rules that let **luie** and **user2s** access your files or images, type: **2 4** or **2,4** and press Enter.

Delete Archive

The **delete archive** command deletes archived files from Tivoli Storage Manager server storage. Your administrator must give you authority to delete archived files.

Important: When you delete archived files, you cannot retrieve them. Verify that the files are obsolete before you delete them.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Delete ARchive─┬──────────┬─┬──────────────────────────────┬──►◄
                   └─options──┘ │ ◄────────────────────┐       │
                                │   ┌─filespec───────────────┐ │
                                └─┬─┤                        ├─┘
                                  │ └─{─filespacename─}─filespec─┘
```

Parameters

filespec
   Specifies the path and file name that you want to delete from storage. Use wildcard characters to specify a group of files or all files in a directory. Separate file specifications with a space. You can also use the *filelist* option to process a list of files. The Tivoli Storage Manager client opens the file you specify with this option and processes the list of files within according to the specific command.

   Note: If you indicate *filespacename*, do not include a drive letter in the file specification.

{*filespacename*}
   Specifies the file space (enclosed in braces) on the server that contains the file you want to delete. This is the name on the workstation drive from which the file was archived.

   Use the *filespacename* if the name has changed, or if you are deleting files archived from another node with drive labels that are different from yours.

Table 101. Delete Archive command: Related options

Option	Where to use
dateformat "Dateformat" on page 289	Client options file (dsm.opt) or command line.
description "Description" on page 295	Command line only.

Table 101. Delete Archive command: Related options (continued)

Option	Where to use
filelist "Filelist" on page 333	Command line only.
noprompt "Noprompt" on page 376	Command line only.
numberformat "Numberformat" on page 376	Client options file (dsm.opt) or command line.
pick "Pick" on page 382	Command line only.
subdir "Subdir" on page 432	Client options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 433	Client options file (dsm.opt) or command line.
timeformat "Timeformat" on page 442	Client options file (dsm.opt) or command line.

Examples

Task   Delete a file named budget.

       Command: delete archive /user/home/proj1/budget

Task   Delete all files archived from the /user/home/proj1 directory with a file extension of .txt.

       Command: del arch "/user/home/proj1/*.txt"

Task   Delete files archived from the /user/project directory using the *pick* option to display a list of archive copies that match the file specification. From the list, you can select the versions to process.

       Command: delete archive "/user/project/*" -pick

Task   Delete selected files from the group of files archived with the description "Monthly Budgets 1999" located in /user/projects and its subdirectories.

       Command: delete ar "/user/projects/*" -description="Monthly Budgets 1999" -pick -subdir=yes

Related information

Delete Backup

The **delete backup** command deletes backup files from Tivoli Storage Manager server storage. Your administrator must give you authority to delete backup files.

You can delete your backed up files from the Tivoli Storage Manager client node workstation if the function is enabled on the definition of the Tivoli Storage Manager node. When you do this, the Tivoli Storage Manager server takes all of the backed up files that meet the 'filespec' and 'deltype' specified and deactivate them. The server also assigns a deactivation date of 'infinite-minus' so that the files are no longer available for restore and are purged, immediately on the subsequent run of file expiration. The file is not physically removed until the expiration process runs.

Important: After you delete backup files, you cannot restore them. Verify that the backup files are no longer needed before you delete them. Tivoli Storage Manager prompts whether you want to continue with the delete. If you specify yes, the specified backup files are immediately deleted and removed from Tivoli Storage Manager server storage.

Supported Clients
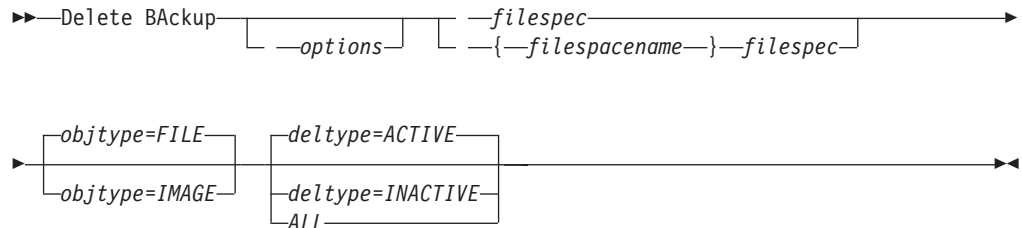
This command is valid for all clients.

Syntax

```
►►──Delete BAckup────────────────┬──────────────────────────────────────────►
                  └─options─┘  └─ ─filespec───────────────────┘
                                  └─{─filespacename─}─filespec─┘


   ┌─objtype=FILE──┐   ┌─deltype=ACTIVE───┐
►──┤               ├───┤                  ├────────────────────────────────►◄
   └─objtype=IMAGE─┘   ├─deltype=INACTIVE─┤
                       └─ALL──────────────┘
```

Parameters

filespec
Specifies the path and file name that you want to delete from storage. Use wildcard characters to specify a group of files or all files in a directory. Separate file specifications with a space. You can also use the filelist option to process a list of files. The Tivoli Storage Manager client opens the file you specify with this option and processes the list of files within according to the specific command.

Note: If you indicate *filespacename*, do not include a drive letter in the file specification.

When using **-deltype=inactive** or **-deltype=active**, use wildcard characters to specify a group of files or all files in a directory.

When using **-deltype=all**, specify a fully-wildcarded directory.

objtype
Specifies the type of backup delete you want to perform. You can specify either of the following values:

FILE
Specifies that you want to delete directories and files. This is the default.

IMAGE
Specifies that you want to delete an image backup. Specifies that you want to delete an image backup. Objtype=image is not supported on Mac OS X.

deltype
Specifies the deletion type. Specify one of the following values:

ACTIVE
Delete only active file objects. Directory objects are not deleted. This is the default.

Note: If there are any inactive objects, then after the active object is deleted, the most current inactive object is changed from inactive to active.

To delete all versions of a file, first issue the **delete backup** command with
`-deltype=inactive`, then issue the command again with `-deltype=active`.

INACTIVE
Delete only inactive file objects. Directory objects are not deleted.

ALL
Delete all active and inactive objects below a given directory, including all
subdirectories and their files.

Note: The parent directory of the deleted files and subdirectories is not
deleted.

Table 102. Delete Backup command: Related options

Option	Where to use
description "Description" on page 295	Command line only.
filelist "Filelist" on page 333	Command line only.
fromdate "Fromdate" on page 338	Command line, and in the GUI find function.
fromtime "Fromtime" on page 340	Command line, and in the GUI find function.
noprompt "Noprompt" on page 376	Command line only.
pick "Pick" on page 382	Command line only.
pitdate "Pitdate" on page 383	Command line, and in the GUI find function.
pittime "Pittime" on page 384	Command line, and in the GUI find function.
subdir "Subdir" on page 432	Client options file (`dsm.opt`) or command line.
tapeprompt "Tapeprompt" on page 433	Client options file (`dsm.opt`) or command line.
timeformat "Timeformat" on page 442	Client options file (`dsm.opt`) or command line.
todate "Todate" on page 445	Command line, and in the GUI find function.
totime "Totime" on page 446	Command line, and in the GUI find function.

Examples

Task   Delete all active and inactive file objects named `budget` in directory
`/data/plan/proj1`.

Commands:

```
delete backup /data/plan/proj1/budget.jan
  -deltype=inactive
delete backup /data/plan/proj1/budget.jan
  -deltype=active
```

Task   Delete all inactive files with names ending in `.txt` backed up from the `/data/plan/proj1` directory and its subdirectories.

        Command: `delete backup "/data/plan/proj1/*.txt" -deltype=inactive -subdir=yes`

Task   Delete selected active files backed up from the `/home/marymb/project` directory. Use the -pick option to display a list of backup copies that match the file specification. From the list, you can select which versions to delete.

        Command: `delete backup "/home/marymb/project/*" -pick`

Task   Delete all active and inactive versions of files and subdirectories in the `/home/storman/myproject` directory. Then delete all active and inactive versions of the `/user/myproject` directory.

        Command:

```
delete backup "/home/storman/myproject*"
 -deltype=all
```

Related information

"Filelist" on page 333

Delete Filespace

The **delete filespace** command deletes file spaces in Tivoli Storage Manager server storage. A file space is a logical space on the server that contains files you backed up or archived.

This command is for an **Authorized User or System Administrator**.

Tivoli Storage Manager assigns a separate file space on the server for each workstation file system from which you back up or archive files. The file space name is the same as the file system name.

When you enter the **delete filespace** command, a list of your file spaces is displayed. From this list, select the file space that you want to delete.

Your Tivoli Storage Manager administrator must give you authority to delete a file space. You need BACKDEL authority if the file space you want to delete contains backup versions, or ARCHDEL authority if the file space contains archive copies. If the file space contains both backup versions and archive copies, you need both types of authority.

Important: When you delete a file space, you delete all backup versions and archive copies within that file space. When you delete a file space, *you cannot restore the files*. Verify that the files are obsolete before you delete them.

You can use the **delete filespace** command to interactively delete NAS file spaces from server storage. Use the nasnodename option to identify the NAS file server. Use the class option to specify the class of the file space to delete.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Delete Filespace───────────────────────────────────────────►◄
                        └──options──┘
```

Parameters

Table 103. Delete Filespace command: Related options

Option	Where to use
class "Class" on page 279	Command line only.
detail "Detail" on page 296	Command line only.
nasnodename "Nasnodename" on page 371	Client system options file or command line.
scrolllines "Scrolllines" on page 410	Client system options file or command line.
scrollprompt "Scrollprompt" on page 411	Client system options file or command line.

Examples

Task   Delete a file space.

    Command: `delete filespace`

Task   Delete NAS file spaces from the **dagordon** NAS file server stored on the server.

    Command: `delete filespace -nasnodename=dagordon -class=nas`

Related information

"Nasnodename" on page 371

"Class" on page 279

Delete Group

Use the **delete group** command to delete a group backup on the Tivoli Storage Manager server.

After deleting a group, the group leader (virtualfsname) remains on the Tivoli Storage Manager server. It contains no members (file or directories) but is reported in a subsequent **query filespace** command. There are no files listed if the showmembers option is added. Deleting a group does not remove the file space that it resides in because there might be other groups in it. Use **delete filespace** if you want to remove the file space and all the data it contains.

Note:

1. Use the inactive option to display both active and inactive group backup versions. By default, Tivoli Storage Manager only displays active versions.
2. Use the pick option to select a specific group to delete from the Tivoli Storage Manager server.
3. Use the noprompt option if you want to suppress the confirmation prompt that normally appears before you delete a group backup version. By default, Tivoli Storage Manager prompts you for confirmation before deleting the group backup. Using this option can speed up the delete procedure. However, it also

increases the danger of accidentally deleting a group backup version that you want to save. Use this option with caution.

4. Use the **query filespace** command to display virtual file space names for your node that are stored on the Tivoli Storage Manager server.

Supported Clients

This command is valid for all UNIX and Linux clients, except for Mac OS X.

Syntax

```
►►──Delete GRoup── ──filespec────────────────────────────────────►◄
                              └─ ──options──┘
```

Parameters

filespec
    Specifies the virtual file space name and the group name that you want to delete from the server storage.

Table 104. Delete Group command: Related options

Option	Where to use
inactive "Inactive" on page 348	Command line only.
noprompt "Noprompt" on page 376	Command line only.
pick "Pick" on page 382	Command line only.
pitdate "Pitdate" on page 383	Command line only.
pittime "Pittime" on page 384	Command line only.

Examples

Task   Delete the current active version of the /virtfs/group1 group.

       Command:
```
delete group /virtfs/group1
```

Task   Delete a backup version of the /virtfs/group1 group from a list of active and inactive versions.

       Command:
```
delete group /virtfs/group1 -inactive -pick
```

Related information

"Inactive" on page 348

"Pick" on page 382

"Noprompt" on page 376

"Query Filespace" on page 512

Expire

The **expire** command inactivates the backup objects you specify in the file specification or with the filelist option.

When working in interactive mode, a prompt notifies you before files are expired.

The **expire** command does not remove workstation files. If you expire a file or directory that still exists on your workstation, the file or directory is backed up again during the next incremental backup unless you exclude the object from backup processing.

If you expire a directory that contains active files, those files will not appear in a subsequent query from the GUI. However, these files are displayed on the command line, if you specify the proper query with a wildcard character for the directory.

Note: Because the **expire** command changes the server picture of the client file system without actually changing the client file system, the **expire** command is not allowed on files located on a file system monitored by the Tivoli Storage Manager journal daemon.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──EXPire──┬──────────────┬──── filespec ────────────────────────►◄
            └── options ──┘
```

Parameters

filespec
> Specifies a path and a file name that you want to expire. You can enter only one file specification on this command. However, you can use wildcards to select a group of files or all the files in a directory. If you specify the filelist option, the `filespec` designation is ignored.

Table 105. Expire command: Related options

Option	Where to use
dateformat "Dateformat" on page 289	Client user-options file (dsm.opt) or command line.
filelist "Filelist" on page 333	Command line only.
noprompt "Noprompt" on page 376	Command line only.
numberformat "Numberformat" on page 376	Client user-options file (dsm.opt) or command line.
pick "Pick" on page 382	Command line only.
timeformat "Timeformat" on page 442	Client user-options file (dsm.opt) or command line.

Examples

Task   Inactivate the `letter1.txt` file in the `home` directory.

   Command: `expire "/home/letter1.txt"`

Task   Inactivate all files in the `/admin/mydir` directory.

   Command: `expire /admin/mydir/*`

Task   Inactivate all files named in the `/home/avi/filelist.txt` file.

   Command: `expire -filelist=/home/avi/filelist.txt`

Help

Use the **help** command to display information about commands, options, and messages.

Tip: If you use the **help** command on the initial command line, no server contact is made and no password is needed.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Help──┬─────────────────────────────────┬──►◄
          ├─command-name [subcommand-name]──┤
          ├─option-name─────────────────────┤
          ├─TOC-section-number──────────────┤
          └─[ANS]message-number─────────────┘
```
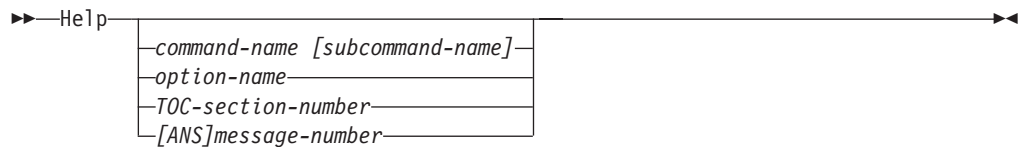
Entering the **help** command with no arguments causes help to display the complete table of contents. Either with the initial command or when HELP displays a prompt, you can enter the following parameters.

Parameters

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

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

TOC-section-number
   Specifies a table of contents section number, for example: 1.5.3. This parameter is optional.

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

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

The requested help text is displayed in one or more sections, depending on the number of display lines that are available in your command window. When enough lines have been displayed to fill the display space, or when the end of the requested help text is displayed, you see a prompt along with instructions for what can be entered at that prompt. To continue displaying text for your current selection, press enter or type the 'd' key to scroll down. To scroll up in the current selection, press the 'u' key and press Enter. Other choices might be presented, so read the prompt instructions.

Proper display of the help text requires a usable display width of 72 characters. A display width less than 72 characters causes sentences that are 72 characters wide to wrap to the next line. This can cause the displayed help text to begin somewhere within the section rather than at the beginning. The undisplayed lines can be viewed by using the scrolling function of the terminal to move up.

Examples

Task   Display the table of contents of the help topics.

       Command: `dsmc help`

Task   Display the information in help topic 2.1.2

       Command: `dsmc help 2.1.2`

Task   Display help information on the **archive** command.

       Command: `dsmc help archive`

Task   Display help information on message ANS1036.

       Command: `dsmc help 1036`

       Command: `dsmc help ANS1036`

Incremental

The **incremental** command backs up all new or changed data in the locations that you specify, unless you exclude them from backup services.

You can back up all new or changed files or directories in the default client domain or from file systems, directories, or files.

To incrementally back up selected files or directories, enter a file specification in the command. If you do not enter a file specification, the default is to back up files or directories in the default domain.

AIX only: You can enable snapshot-based incremental backup by using the option snapshotproviderfs =*JFS2* .

The following attributes in the management class assigned to the file or directory affect whether the data is actually backed up:

Frequency

The number of days that must elapse between successive backups for the file. The **frequency** attribute applies only to a full incremental backup.

This management class attribute is ignored during a journal-based backup.

Mode Permits you to back up only files that changed since the last backup (*modified*) or back up the files whether they changed or not (*absolute*).

Serialization

Permits or denies backup of files or directories according to the following values:

- **static**: In order to be backed up, data must not be modified during backup or archive.
- **shared static**: If data in the file or directory changes during each of the allowed attempts to back up or archive it, it is not backed up or archived. The value of the changingretries option determines how many attempts are made. The default is 4.
- **dynamic**: The object is backed up or archived on the first attempt whether or not data changes during the process.
- **shared dynamic**: The object is backed up or archived on the last attempt, even if data changes during the process.

Using the include option in an include-exclude list, you can override the default management class for a file or group of files.

You can perform either a *full incremental* backup or an *incremental by date* backup. The default is a full incremental backup.

If you are journaling a file system and the journal is valid, the full incremental backup performs a journal-based backup. More than one journal-based backup session can be started, but only one journal-based backup session can proceed. All other journal-based backup sessions that need access to the same filespace must wait until the current journal-based backup session has completed before the next session can proceed. You can perform a full incremental backup without the journal by using the nojournal option.

You can also use the **selective** command to perform a *selective* backup that backs up only the files, directories or empty directories that you specify regardless of whether they have changed.

A full incremental backs up all files or directories that are new or have changed since the last incremental backup. During a full incremental backup, the client queries the server. Tivoli Storage Manager uses this information to:
- Back up new files or directories.
- Back up files or directories whose contents have changed.
- Mark inactive backup versions on the server for files or directories that are deleted from the workstation.
- Rebind backup versions to management classes if the management class assignments change.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Incremental─────────────────────┬─────────────┬─────────────────►◄
                  └──options──┘      ▼             │
                                     ├──filespec───┤
                                     └─"filespec"──┘
```
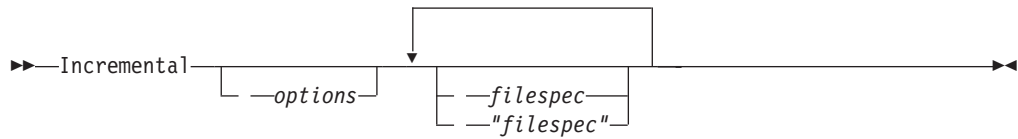
Parameters

filespec

Specifies the path and file name that you want to back up. Use wildcards to select a group of files or all the files in a directory. If you do not specify a file specification, the default domain or the domain specified as an option is backed up.

If you specify a file system, all new and changed files are backed up. In addition, the last incremental date for the file space is updated on the server. If you specify a file or directory, the last incremental date is not updated. This means the file or directory might be backed up again if a later backup is performed using the incrbydate option. If you specify a file system, specify the file system without a trailing slash.

Table 106. Incremental command: Related options

Option	Where to use
changingretries "Changingretries" on page 278	dsm.sys file or command line.
compressalways "Compressalways" on page 284	Client user-options file (dsm.opt) or command line.
compression "Compression" on page 285	dsm.sys file *within* a server stanza or command line.
diffsnapshot "Diffsnapshot" on page 297	Command line only.
dirsonly "Dirsonly" on page 300	Command line only.
domain "Domain" on page 303	dsm.sys file or the client user-options file (dsm.opt) or command line only.
encryptiontype "Encryptiontype" on page 315	System-options file (dsm.sys) within a server stanza.
encryptkey "Encryptkey" on page 316	System-options file (dsm.sys) within a server stanza.
filelist "Filelist" on page 333	Command line only.
filesonly "Filesonly" on page 336	Command line only.
incrbydate "Incrbydate" on page 357	Command line only.
memoryefficientbackup "Memoryefficientbackup" on page 368	Client user-options file (dsm.opt), client system-options file (dsm.sys), server, or command line.
nojournal "Nojournal" on page 375	Command line only.
preservelastaccessdate "Preservelastaccessdate" on page 389	Client user-options file (dsm.opt) or command line.
removeoperandlimit "Removeoperandlimit" on page 396	Command line only.
snapdiff "Snapdiff" on page 418	Command line only.

Table 106. Incremental command: Related options (continued)

Option	Where to use
snapshotproviderfs "Snapshotproviderfs" on page 423	System-options file (dsm.sys) within a server stanza or with the include.fs option.
snapshotroot "Snapshotroot" on page 425	Command line only.
subdir "Subdir" on page 432	Client user-options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 433	Client user-options file (dsm.opt) or command line.

Examples

Task   Run an incremental backup of the default client domain specified in your client user-options file (dsm.opt).

     Command: `Incremental`

Task   Run an incremental backup for the /home, /usr, and /proj file systems.

     Command: `Incremental /home /usr /proj`

Task   Run an incremental backup for the /proj/test directory.

     Command: `Incremental /proj/test/`

Task   Run an incremental-by-date backup for the /home file system.

     Command: `Incremental -incrbydate /home`

Task   Run an incremental backup of the abc file in the /fs/dir1 directory.

     Command: `Incremental -subdir=yes /fs/dir1/abc`

Task   Run an incremental backup of the directory object /fs/dir1, but not any of the files in the /fs/dir1 directory.

     Command: `Incremental /fs/dir1`

Task   Run an incremental backup of the directory object /fs/dir1, all of the files in the fs/dir1 directory, and all files and subdirectories under /fs/dir1.

     Command: `Incremental -subdir=yes /fs/dir1/`

Task   Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, run an incremental backup of all files and directories under the local snapshot and manage them on the Tivoli Storage Manager server under the file space name /usr.

     Command: `dsmc inc /usr -snapshotroot=/snapshot/day1`

Task   Run an incremental backup for the /home file system using the snapdiff option. Tell Tivoli Storage Manager to create the difference snapshot. Here /home is the NFS mount point for a NAS/N-Series file server volume.

     Command: `incremental /home -snapdiff -diffsnapshot=create`

Task   Run an incremental backup of the /proj file system using the snapdiff option. Tell Tivoli Storage Manager to use the latest snapshot on the file server as the difference snapshot. Here /proj is the NFS mount point for a NAS/N-Series file server volume.

     Command: `incremental /proj -snapdiff -diffsnapshot=latest`

Related information

"Journal-based backup"

"Selective" on page 548

"Include options" on page 351

Journal-based backup

A backup for a particular file system is journal-based when the Tivoli Storage Manager journal daemon is installed and configured to journal the file system, and a valid journal has been established.

Journal-based backup is supported on the AIX Backup-Archive client, on JFS, JFS2, and VxFs file systems.

If the *journal daemon* is installed and running, then by default the **incremental** command performs a journal-based backup on file systems which are being monitored by the journal engine daemon. The following conditions must be met in order to successfully perform a journal-based backup:

- The journal daemon must be set up to monitor the file system that contains the files and directories being backed up.
- A full incremental backup must have been run successfully at least once on the file system being backed up.
- The file space image of the file system at the server cannot have been modified by an administrative command since the last full incremental backup.
- The storage management policy for the files being backed up cannot have been updated since the last full incremental backup.

The *journal daemon* records changes to an object or its attributes in a journal database. During a journal-based backup, the client obtains a list of files that are eligible for backup from the journal database. Journal-based backup can increase backup performance because the client does not scan the local file system or contact the server to determine which files to process. Journal-based backup also reduces network traffic between the client and server

Tivoli Storage Manager filters the list based on the current include-exclude list and processes, expires, and updates the resulting files according to policy constraints, such as serialization. However, the client ignores the server frequency attribute during a journal-based backup. The reason for this is because a journal-based backup eliminates the backup version query to the server; therefore, the client does not know how many days have transpired since the last backup of the file.

The journal daemon does not record changes in UNIX special files.

The journal daemon excludes specific system files from having changes recorded in the journal. Because changes to these files are not journaled, Tivoli Storage Manager does not back up these files. See the journal daemon configuration file `tsmjbbd.ini` located in the Tivoli Storage Manager installation directory for specific system files that are excluded.

Note:
1. When using antivirus software, there are limitations to journal-based backup. Some antivirus software can incorrectly generate change notifications to the

TSM journal service, causing files that have not changed to be incorrectly backed up during journal based backup. To avoid these problems, use Norton Anti-Virus Corporate Edition 8.0 and higher.

2. A journal-based backup might not fall back to the traditional incremental backup if the policy domain of your node is changed on the server. This depends on when the policy set within the domain was last updated and the date of the last incremental backup. In this case, you must force a full traditional incremental backup to rebind the files to the new domain. Use the nojournal option with the **incremental** command to specify that you want to perform a traditional full incremental backup, instead of the default journal-based backup. Add an exclude snapshot statement to the `tsmjbbd.ini` file for AIX 6.1 or later to avoid JFS2 internal snapshot directories to be monitored by the journal-based backup daemon.

Under the following conditions, the journal database is considered invalid and the client reverts to the traditional full incremental backup:
- A journaled file space name has changed.
- The client node name has changed.
- The client contacts a different server to do the backup.
- Policy changes have occurred (new policy set activation).
- The journal is corrupt (out of space conditions, disk error).
- The journal is not running.

Journal-based backup differs from the traditional full incremental backup in the following ways:
- Tivoli Storage Manager does not enforce non-default copy frequencies (other than 0).
- Changes to UNIX special files are not detected.

You can use the nojournal option with the **incremental** command to perform a traditional full incremental backup instead of the default journal-based backup.

Incremental-by-Date

An incremental-by-date backup backs up new and changed files with a modification date later than the date of the last incremental backup stored at the server, unless the files are excluded from backup by an **exclude** statement.

If an incremental-by-date is performed on only part of a file system, the date of the last full incremental is not updated, and the next incremental-by-date will back up these files again. Changes to the access control lists (ACL) or Extended Attributes do not cause the files to be backed up during an incremental-by-date. Use the **query filespace** command to determine the date and time of the last incremental backup of the entire file system.

To perform an incremental-by-date backup, use the incrbydate option with the **incremental** command.

Unlike a full incremental, an incremental-by-date does not maintain current server storage of *all* your workstation files for the following reasons:
- It does not expire backup versions of files that are deleted from the workstation.
- It does not rebind backup versions to a new management class if the management class has changed.
- It does not back up files with attributes that have changed, unless the modification dates and times have also changed.

- It ignores the copy group frequency attribute of management classes.

For these reasons, if you have limited time during the week to perform backups, but extra time on the weekends, you can perform an incremental-by-date backup on weekdays and a full incremental backup on weekends to maintain current server storage of your workstation files.

If the **incremental** command is retried because of a communication failure or session loss, the transfer statistics will display the number of bytes Tivoli Storage Manager attempted to transfer during *all* command attempts. Therefore, the statistics for bytes transferred might not match the file statistics, such as those for file size.

Associate a local snapshot with a server file space

Use the snapshotroot option with the **incremental** command in conjunction with a vendor-supplied application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server.

The snapshotroot option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot.

Loop

The **loop** command starts an interactive command-line session that is maintained until you enter **quit**.

If you are required to enter a password, you are prompted for it before the loop mode prompt appears.

Note: It is not possible to enter loop mode without a valid server contact. One of the consequences is that certain commands, such as `restore backupset -location=file`, are only accepted on the initial command line when a valid server is not available.

In an interactive command-line session, it is unnecessary to precede each command name with **dsmc** and your password, if one is required.

In interactive mode, options you enter on the initial command line overrides the value that you specified in your client user-options file (dsm.opt) or dsm.sys file. This value remains in effect for the entire interactive session unless overridden by a different value on a given interactive command. For example, if you set the subdir option to *yes* in your client user-options file (dsm.opt), and you specify subdir=*no* on the initial command line, the subdir=*no* setting remains in effect for the entire interactive session unless overridden by the subdir=*yes* value on a given interactive command. However, the subdir=*yes* value only affects the command it is entered on. When that command completes, the value reverts back to subdir=*no*, the value at the beginning of the interactive session.

You can enter all valid commands in interactive mode *except* the **schedule** and **loop** commands.

There are some options that you cannot use in the interactive session created by the **loop** command and are identified in the option description by this statement: *This option is valid only on the initial command line. It is not valid in interactive mode.*

Note:

1. In loop mode, following a restore operation directly from tape, the mount point is not released in case additional restore requests are made to that volume. If you request a backup operation in the same session and that mount point is the only one available, the backup operation stops with the following message:

   ```
   Waiting for mount of offline media
   ```

   In this case, the mount point is not released until one of the following conditions is met:
   - The device class MOUNTRETENTION limit is satisfied.
   - The client idletimeout period is satisfied.
   - The dsmc loop session is closed after the restore operation completes, allowing you to start a subsequent loop mode session to perform the backup operation.

2. In interactive mode, you cannot enter a file specification that contains national language characters. If a command contains national characters, process the command in batch mode by preceding the command with the executable program name, **dsmc**.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──LOOP──────────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this command.

Examples

Task   Start an interactive command-line session.

Command: dsmc

At the **tsm>** prompt, enter a command.

There are two methods for ending an interactive session:
- Enter quit
- If you set editor=*yes*, you can do the following:
  1. Press the Escape key (Esc).
  2. Type Q and press the Enter key.

  Note: The default setting is editor=*yes*.

  Note: To interrupt a **dsmc** command before Tivoli Storage Manager has finished processing, enter *QQ* on the Tivoli Storage Manager console. In many cases, but not all, this interrupts the command.

Related information

Chapter 10, "Processing options," on page 241 for options that you cannot use in interactive mode.

Macro

The **macro** command executes a series of commands that you specify in a macro file.

By including the **macro** command within a macro file, you can nest as many as ten levels of commands.

Comment lines are not supported within the macro file that you specify for the **macro** command.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──MAcro─── ──macroname────────────────────────────────────────────────►◄
```

Parameters

macroname
    Specifies the fully qualified name of the file containing the commands.

Examples

The following is an example of how to use the **macro** command.

Task   Selectively back up files in the following directories:
- /devel/project/proja
- /devel/project/projb
- /devel/project/projc

Command: macro backabc.mac

where backabc.mac contains the following statements:

```
Selective /devel/project/proja/
Selective /devel/project/projb/
Selective /devel/project/projc/
```

Monitor Process

The **monitor process** command displays a list of current NAS (if NDMP support is enabled) image backup and restore processes for which the administrative user has authority.

The administrative user can then select one process to monitor. Client owner privilege is sufficient authority to monitor the selected NAS image backup or restore processes.

When using an interactive command-line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID.

Supported Clients

This command is valid for AIX, Linux, and Solaris clients only.

Syntax

```
►►──MONitor Process──────────────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this command.

Examples

Task   Monitor current NAS image backup or restore processes.

Command: `monitor process`

Preview Archive

The **preview archive** command simulates an archive command without sending data to the server.

The **preview archive** command generates a tab-delineated text file that can be imported into a spreadsheet program. The preview contains information such as whether the file is excluded or included. If the file is excluded, the pattern, or reason, that the file is excluded is listed, along with the source for the pattern.

Supported Clients

This command is valid for all clients.

Syntax

```
                                         ┌─-filter=ALL──┐
►►──PREview──Archive──filespec────┼─-filter=INCL─┤──────────────────────────────►
                                         └─-filter=EXCL─┘      └─-FILEName= filename─┘

                           ┌─-TRAverse=Yes─┐
►──────────────────────────┼───────────────┼───────────────────────────────────►◄
        └─-CONsole─┘      └─-TRAverse=No──┘
```

Parameters

filespec
> Specifies the path and file name that you want to archive. Use wildcard characters to select a group of files or all the files in a directory.

-filter   Specifies the output to display – included objects, excluded objects, or both.

> **ALL**   Display output for included and excluded objects. This is the default.

> **INCLuded**
> > Display output for included objects only.

> **EXCLuded**
> > Display output for excluded objects only.
>
> **-FILEName=**
> > Specifies the filename in which to write the tab-delineated output. The default is dsmprev.txt.
>
> **-CONsole**
> > Output is written to the console, and the file.
>
> **-TRAverse**
> > Preview the current directory and subdirectories.
> >
> > **Yes** Preview the current directories and subdirectories. This is the default.
> >
> > **No** Preview only the current directory, not subdirectories.
> >
> > **Important:** Specifying **-traverse** does not preview directories excluded using the exclude.dir option.

Preview Backup

The **preview backup** command simulates a backup command without sending data to the server.

The **preview backup** command generates a tab-delineated text file that can be imported into a spreadsheet program. The preview contains information such as whether the file is excluded or included. If the file is excluded, the pattern, or reason, that the file is excluded is listed, along with the source for the pattern.

Supported Clients

This command is valid for all clients.

Syntax

```
>>--PREview--backup--filespec--+-filter=ALL--+------------------------------->
                               +-filter=INCL-+
                               '-filter=EXCL-'   '--FILEName= filename-'

                  --TRAverse=Yes--
>--+----------+--+----------------+-------------------------------------------><
   '-CONsole--'  '--TRAverse=No---'
```

Parameters

filespec
> > Specifies the path and file name that you want to back up. Use wildcard characters to select a group of files or all the files in a directory.

-filter Specifies the output to display – included objects, excluded objects, or both.

> > **ALL** Display output for included and excluded objects. This is the default.
> >
> > **INCLuded**
> > > Display output for included objects only.

> **EXCLuded**
> > Display output for excluded objects only.
>
> **-FILEName=**
> > Specifies the filename in which to write the tab-delineated output. The default is dsmprev.txt.
>
> **-CONsole**
> > Output is written to the console, and the file.
>
> **-TRAverse**
> > Preview the current directory and subdirectories.
> >
> > **Yes** Preview the current directories and subdirectories. This is the default.
> >
> > **No** Preview only the current directory, not subdirectories.
> >
> > **Important:** Specifying **-traverse** does not preview directories excluded using the exclude.dir option.

Query Access

The **query access** command shows who you have given access to backup versions or archive copies of specific files.

Tivoli Storage Manager displays a list of authorization rules that you defined with the **set access** command or with *Node Access List* on the graphical user interface (GUI) Utilities menu.

The following information is included.
- Authority you gave a user to restore backup versions or retrieve archive copies.
- The node name of the user to whom you gave authorization.
- The ID of the user at that node to whom you gave authorization.
- The files to which the user has access.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Query ACcess──────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this command.

Examples

Task Display a list of users who have access to your files.

> **Command:** `query access`

Query Archive

The **query archive** command displays a list of your archived files and the following information about each file: file size, archive date, file specification, expiration date, and archive description.

If you use the detail option with the **query archive** command, the client displays the following additional information:
- Last modification date
- Last access date
- Compression
- Encryption type
- Client-side data deduplication
- Retention initiation
- Whether or not the file is on hold

```
Size Archive Date - Time File - Expires on - Description
---- ------------------ -------------------------------
219 B 09/03/2009 09:32:13 /Volumes/Data/info.txt 09/03/2010
Archive Date: 09/03/2009
RetInit:STARTED Obj
Held:NO
Modified: 09/02/2009 19:43:00 Created: 09/01/2009 15:31:23
Compressed: NO Encryption Type: None
Client-deduplicated: NO
```

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Query ARchive──┬──────────┬──┬──────────────┬──────────────────────────►◄
                   └─options─┘   ├──filespec───┤
                                 └──"filespec"─┘
```

Parameters

filespec
> Specifies the path and file name that you want to query. Use wildcard characters to specify a group of files or all the files in a directory. If you use wildcard characters, enclose the file specification in double quotation marks. Specify an asterisk (*) to query all archived files in the current directory.

Table 107. Query Archive command: Related options

Option	Where to use
dateformat "Dateformat" on page 289	Client user-options file (dsm.opt) or command line.
description "Description" on page 295	Command line only.
detail "Detail" on page 296	Command line only.
dirsonly "Dirsonly" on page 300	Command line only.
filelist "Filelist" on page 333	Command line only.

Table 107. Query Archive command: Related options (continued)

Option	Where to use
filesonly "Filesonly" on page 336	Command line only.
fromdate "Fromdate" on page 338	Command line only.
fromnode "Fromnode" on page 339	Command line only.
fromowner "Fromowner" on page 339	Command line only.
fromtime "Fromtime" on page 340	Command line only.
numberformat "Numberformat" on page 376	Client user-options file (dsm.opt) or command line.
querysummary "Querysummary" on page 394	Command line only.
scrolllines "Scrolllines" on page 410	Client user-options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 411	Client user-options file (dsm.opt) or command line.
subdir "Subdir" on page 432	Client user-options file (dsm.opt) or command line.
timeformat "Timeformat" on page 442	Client user-options file (dsm.opt) or command line.
todate "Todate" on page 445	Command line only.
totime "Totime" on page 446	Command line only.

Examples

Task Display a list of all your archived files in the current working directory.

  Command: `q archive "*"`

Task Display a list of all your archived files in the /devel directory and all of its subdirectories.

  Command: `query archive "/devel/*" -subdir=yes`

Task Display a list of all your archived files in the current directory. Use the dateformat and timeformat options to reformat the dates and times.

  Command: `q ar —date=5 —time=1 "*"`

Task Display a list of all your archived files in the current directory. Use the detail option to display the last modification date and the last access date of each file.

  Command: `q ar -detail "*"`

Task Display a list of archived files in the /home/proj directory whose first four characters of the file name begin with proj.

Query Backup

The **query backup** command displays a list of backup versions of your files that are stored on the Tivoli Storage Manager, or that are inside a backup set from the server when the backupsetname option is specified.

The command displays the following file information:
- File specification
- File size
- Backup date
- Whether the file is active or inactive
- The management class that is assigned to the file. Only the first ten characters of the management class name appear.

If you use the **detail** option with the **query backup** command, the client displays the following additional information:
- Last modification date
- Last access date
- Compression
- Encryption type
- Client-side data deduplication

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Query Backup──────────────────────────┬──filespec────┬─────────────────────────►◄
                  └──options──┘           └──"filespec"──┘
```

Parameters

filespec
    Specifies the path and file name that you want to query. Use wildcard characters to specify a group of files or all the files in a directory. If you use wildcard characters, enclose the file specification in double quotation marks. Specify an asterisk (*) to display information about backup versions for all of your files in the current directory. Do not use wild cards when you query NAS file system images with `-class=nas` option setting.

Table 108. Query Backup command: Related options

Option	Where to use
backupsetname "Backupsetname" on page 277	Command line only.
class "Class" on page 279	Command line only.
dateformat "Dateformat" on page 289	Client system options file (dsm.sys) or command line.

Table 108. Query Backup command: Related options (continued)

Option	Where to use
detail "Detail" on page 296	Command line only.
dirsonly "Dirsonly" on page 300	Command line only.
filelist "Filelist" on page 333	Command line only.
filesonly "Filesonly" on page 336	Command line only.
fromdate "Fromdate" on page 338	Command line only.
fromowner "Fromnode" on page 339	Command line only.
fromowner "Fromowner" on page 339	Command line only.
fromtime "Fromtime" on page 340	Command line only.
inactive "Inactive" on page 348	Command line only.
nasnodename "Nasnodename" on page 371	Client system options file (dsm.sys) or command line.
numberformat "Numberformat" on page 376	Client user-options file (dsm.opt) or command line.
pitdate "Pitdate" on page 383	Command line only.
pittime "Pittime" on page 384	Command line only.
querysummary "Querysummary" on page 394	Command line only.
scrolllines "Scrolllines" on page 410	Client user-options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 411	Client user-options file (dsm.opt) or command line.
subdir "Subdir" on page 432	Client user-options file (dsm.opt) or command line.
timeformat "Timeformat" on page 442	Client user-options file (dsm.opt) or command line.
todate "Todate" on page 445	Command line only.
totime "Totime" on page 446	Command line only.

Examples

Task Display a list of all active and inactive backup versions of your files in the current directory.

> **Command:** `query backup –inactive "*"`

Task Display a list of all your backups in the current directory. Use the detail option to display the last modification date and the last access date of each file.

> **Command:** `q backup -detail "*"`

Task Display a list of files that were backed up from the `/home/proj` directory with file names that begin with proj.

> **Command:** `q b "/home/proj/proj*"`

Task Display a list of active and inactive backup file versions in the `/home` file system.

> **Command:** `q b –ina –su=yes /home/`

Task Query file system images from the **nas2** NAS file server.

> **Command:** `query backup -nasnodename=nas2 –class=nas`

Related information

"Restore data from a backup set" on page 179

Query NAS file system images

You can use the **query backup** command to display information about file system images backed up for a NAS file server.

Where supported, use the nasnodename option to identify the NAS file server to query. When using an interactive command-line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID. Place the nasnodename option in your client system-options file (dsm.sys). The value in the client system-options file is the default, but this value can be overridden on the command line.

Use the class option to specify the class of the file space to query. To display a list of images belonging to a NAS node, use the -class=*nas* option.

Related reference

"Class" on page 279

"Nasnodename" on page 371

Query Backupset

The **query backupset** command queries a backup set from a local file, tape device (if applicable) or the Tivoli Storage Manager server.

This command displays the backup set name, generation date, retention (for a backup set on a Tivoli Storage Manager server), and user-supplied description.

Supported Clients

This command is valid for all clients.

Tape support is only available on AIX, Solaris, and HP-UX.

Syntax

```
►►──Query BACKUPSET──────────────BACKUPSETName=──────backupsetname─────────────►
                    └─options─┘                    ├─localfilename─┤
                                                   └─tapedevice────┘

►──────────────────────────────────────────────────────────────────────────◄
   └─LOCation=──┬─server─┬──┘
               ├─file───┤
               └─tape───┘
```
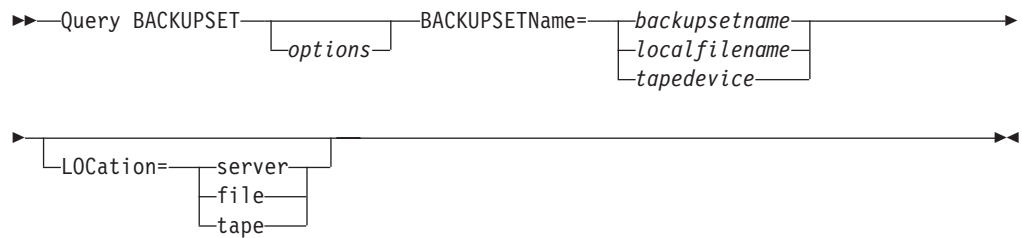
Parameters

BACKUPSETName=
> Specifies the name of a backup set you want to query. You can use wildcards to specify the backup set name. If you use wildcards or do not specify a backup set name, all backup sets that you own are displayed. This parameter is required.
>
> When a backup set is created, the server assigns root as the owner of the backup set. When querying a backup set on the server, a non-root user will not see the backup set listed, even if they know the backup set name and use it in the query.
>
> The value of **backupsetname** depends on the location of the backup set, and corresponds to one of the three choices below.
>
> **backupsetname**
>> Specifies the name of the backup set from the Tivoli Storage Manager server. If the **location** parameter is specified, you must set `-location=server`.
>
> **localfilename**
>> Specifies the file name of the first backup set volume. You must set `-location=file`.
>
> **tapedevice**
>> Specifies the name of the tape device containing the backup set volume. You must use a Windows native device driver, not the device driver provided by Tivoli. You must set `-location=tape`.

LOCation=
> Specifies where Tivoli Storage Manager searches for the backup set. If you do not specify the location parameter, the client searches for backup sets on the Tivoli Storage Manager server.
>
> **server** Specifies that Tivoli Storage Manager searches for the backup set from the server. This is the default.
>
> **file** Specifies that Tivoli Storage Manager searches for the backup set from a local file.
>
> **tape** Specifies that Tivoli Storage Manager searches for the backup set from a local tape device.

Table 109. Query Backupset command: Related options

Option	Where to use
description "Description" on page 295	Command line only.

Table 109. Query Backupset command: Related options (continued)

Option	Where to use
scrolllines "Scrolllines" on page 410	Client user-options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 411	Client user-options file (dsm.opt) or command line.

Examples

Task   Query all backup sets from the Tivoli Storage Manager server.

   Command: `query backupset -backupsetname=*`

Task   Query a backup set called `monthy_financial_data` from the Tivoli Storage Manager server.

   Command: `query backupset -backupsetname=monthly_financial_data.12345678`

Task   Query the backup set in the file /home/budget/weekly_budget_data.ost.

   Command: `dsmc query backupset -backupsetname="/home/budget/ weekly_budget_data.ost" -loc=file`

Task   Query the backup set from the /dev/rmt0 tape device.

   Command: `dsmc query backupset -backupsetname=/dev/rmt0 -loc=tape`

Related information
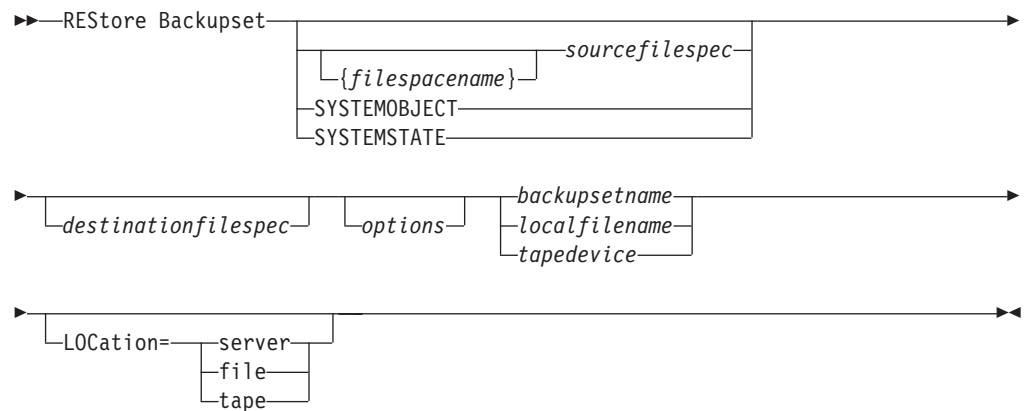
Query Backupset without the backupsetname parameter

The **query backupset** command can be used without the **backupsetname** parameter.

The preferred syntax for **query backupset** command requires the **backupsetname** parameter. Prior to the introduction of the **backupsetname** parameter, the Tivoli Storage Manager client queried backup sets with a different syntax.

While you can use syntax from previous releases for this command, do not do so unless you have a specific need and cannot replace the old syntax with the syntax in Tivoli Storage Manager Version 6.1. For best results, use the **backupsetname** parameter.

Supported Clients

This command is valid for all clients. Tape support is only available on AIX, Solaris, and HP-UX.

Syntax

```
►►─Query BACKUPSET──────────────┬─backupsetname─┬───────────────────►◄
                 └─options─┘     ├─localfilename─┤  └─LOCation=─┬─server─┬─┘
                                 └─tapedevice────┘             ├─file───┤
                                                               └─tape───┘
```

Parameters

backupsetname

Specifies the name of the backup set from the Tivoli Storage Manager server. If the **location** parameter is specified, you must set `-location=server`.

localfilename

Specifies the file name of the first backup set volume. You must set `-location=file`.

tapedevice

Specifies the name of the tape device containing the backup set volume. You must use a Windows native device driver, not the device driver provided by Tivoli. You must set `-location=tape`.

LOCation=

Specifies where Tivoli Storage Manager searches for the backup set. If you do not specify the location parameter, the client searches for backup sets on the Tivoli Storage Manager server.

server Specifies that Tivoli Storage Manager searches for the backup set from the server. This is the default.

file Specifies that Tivoli Storage Manager searches for the backup set from a local file.

tape Specifies that Tivoli Storage Manager searches for the backup set from a local tape device.

Table 110. Query Backupset command: Related options

Option	Where to use
description "Description" on page 295	Command line only.
scrolllines "Scrolllines" on page 410	Client user-options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 411	Client user-options file (dsm.opt) or command line.

Examples

Task Query all backup sets from the Tivoli Storage Manager server.

Command: `query backupset`

Task Query a backup set called `monthy_financial_data` from the Tivoli Storage Manager server.

Command: `query backupset monthly_financial_data.12345678`

Task Query the backup set in the file /home/budget/weekly_budget_data.ost.

Command: `dsmc query backupset /home/budget/weekly_budget_data.ost -loc=file`

Task Query the backup set from the /dev/rmt0 tape device.

Command: `dsmc query backupset /dev/rmt0 -loc=tape`

Related information

"Restore data from a backup set" on page 179

Query Filespace

The **query filespace** command displays a list of file spaces for a node that are stored on the Tivoli Storage Manager server, or that are inside a backup set from the Tivoli Storage Manager server, when the backupsetname option is specified. You can also specify a single file space name to query.

A *file space* is a logical space on the server that contains files you backed up or archived. Tivoli Storage Manager assigns a separate file space on the server for each node at your workstation from which you back up or archive files.

Tivoli Storage Manager assigns a separate file space on the server for each file system at your workstation from which you back up or archive files. The file space name is the same as the file system name.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Query Filespace─────────────────────────────────────────────────►◄
                    └─ filespacename ─┘  └─ options ─┘
```

Parameters

filespacename
> Specifies an optional character string which can include wildcards. Use this argument to specify a subset of file spaces. The default is to display all file spaces.

Table 111. Query Filespace command: Related options

Option	Where to use
backupsetname "Backupsetname" on page 277	Command line only.
class "Class" on page 279	Command line only.
dateformat "Dateformat" on page 289	Client user-options file (dsm.opt) or command line.
detail "Detail" on page 296	Command line only.
fromnode "Fromnode" on page 339	Command line only.
fromowner "Fromowner" on page 339	Command line only.
nasnodename "Nasnodename" on page 371	Client system options file (dsm.sys) or command line.
scrolllines "Scrolllines" on page 410	Client user-options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 411	Client user-options file (dsm.opt) or command line.

Table 111. Query Filespace command: Related options (continued)

Option	Where to use
timeformat "Timeformat" on page 442	Client user-options file (dsm.opt) or command line.

Examples

Task   Display your file spaces. Use the dateformat and timeformat options to reformat the dates and times.

> **Command:** `query filespace –date=5 –time=4`

Task   Display the /home file space.

> **Command:** `query filespace /home`

Task   Display file space names that include the pattern `smith`.

> **Command:** `query filespace "*smith*"`

Task   Query a file space from the **nas2** NAS file server.

> **Command:** `query filespace -nasnodename=nas2 -class=nas`

Related information

"Restore data from a backup set" on page 179

"Nasnodename" on page 371

"Class" on page 279

Query NAS file spaces

Use the nasnodename option to identify the NAS file server to query. When using an interactive command-line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID.

Place the nasnodename option in your client system-options file (dsm.sys). The value in the client system-options file is the default, but this value can be overridden on the command line. If the nasnodename option is not specified in the client system-options file, it must be specified on the command line when processing NAS file systems.

Use the class option to specify the class of the object to query. To display a list of file spaces belonging to a NAS node, use the -class=*nas* option.

Query Group

Use the **query group** command to display information about a group backup and its members.

Note:

1. Use the showmembers option to display and select individual group members that you want to query. The showmembers option is not valid with the inactive option. If you want to display members of a group that are not currently active, use the pitdate and pittime options to specify the backup date and time of the member you want to query.

2. Use the **query filespace** command to display virtual file space names for your node that are stored on the Tivoli Storage Manager server.
3. If you perform a full and differential group backup, a query of this group using the -inactive option displays two active backups of the same name, one of type FULL and one of type DIFF.

```
tsm> q group {/fs}/v1 -inactive

Size Backup Date       Mgmt Class A/I Group
____ _____       _____ ___ _____
 978 B 06/02/2007 11:57:04  DEFAULT   A  FULL /fs/v1
  32 B 06/05/2007 13:52:04  DEFAULT   A  DIFF /fs/v1
```

If you query a group backup without the -inactive option, the query displays only the latest group backup, whether it is type FULL or type DIFF:

```
tsm> q group {/fs}/v1

Size Backup Date       Mgmt Class A/I Group
____ _____       _____ ___ _____
  32 B 06/05/2007 13:52:04  DEFAULT   A  DIFF /fs/v1
```

Supported Clients

This option is valid for all clients, except for Mac OS X.

Syntax

```
>>──Query GRoup── ─filespec─────────────────────────────────><
                            └─ ─options─┘
```

Parameters

filespec
    Specifies the virtual file space name and the group name on the server that you want to query.

Table 112. Query Group command: Related options

Option	Where to use
fromnode "Fromnode" on page 339	Command line only.
fromowner "Fromowner" on page 339	Command line only.
inactive "Inactive" on page 348	Command line only.
pitdate "Pitdate" on page 383	Command line only.
pittime "Pittime" on page 384	Command line only.
showmembers "Showmembers" on page 416 (does not apply to Mac OS X)	Command line only.

Examples

Task   Display all the groups in the /virtfs file space.

> **Command:**
>
> ```
> query group /virtfs/*
> ```

Task Display active and inactive versions of the /virtfs/group1 file space.

> **Command:**
>
> ```
> query group /virtfs/group1 -inactive
> ```

Task Display the /virtfs/group1 file space. Use the showmembers option to display a list of group members from which you can select one or more to query.

> **Command:**
>
> ```
> query backup /virtfs/group1 -showmembers
> ```

Related information

"Query Filespace" on page 512

Query Image

The **query image** command displays information about file system images that are stored on the Tivoli Storage Manager server, or that are inside a backup set from the Tivoli Storage Manager server, when the backupsetname option is specified.

The following information about file system images is displayed:
- Image Size - This is the volume size which was backed up.
- Stored Size - This is the actual image size stored on the server. The stored image on the Tivoli Storage Manager server is the same size as the volume capacity. For online snapshot-based image backups, the stored image can be larger than the file system based on the size of the cache files. The stored image on the Tivoli Storage Manager server is the same size as the volume capacity.
- File system type
- Backup date and time
- Management class assigned to image backup
- Whether the image backup is an active or inactive copy
- The image name

Note: The Tivoli Storage Manager API must be installed to use the **query image** command.

Supported Clients

This command is valid for AIX, HP-UX, all Linux clients, and Solaris.

Syntax

```
►►──Query IMage──┬──────────┬──┬──logicalvolumename──┬──────────────────►◄
                 └─options──┘  └──filespacename──────┘
```

Parameters

logicalvolumename
    The name of a logical volume you want to query. You must specify the exact

name of the image. You cannot use wildcards. The default is all active images (unless restricted by one or more options).

filespacename
> Specifies the file system name that you want to query.

Omitting *logicalvolumename* and *filespacename* causes all images to be displayed.

Table 113. Query Image command: Related options

Option	Where to use
backupsetname "Backupsetname" on page 277	Command line only.
dateformat "Dateformat" on page 289	Client user option file (dsm.opt) or command line.
fromnode "Fromnode" on page 339	Command line only.
fromowner "Fromowner" on page 339	Command line only.
inactive "Inactive" on page 348	Command line only.
numberformat "Numberformat" on page 376	Client user option file (dsm.opt) or command line.
pitdate "Pitdate" on page 383	Command line only.
pittime "Pittime" on page 384	Command line only.
scrolllines "Scrolllines" on page 410	Client user options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 411	Client user options file (dsm.opt) or command line.
timeformat "Timeformat" on page 442	Client user option file (dsm.opt) or command line.

Examples

Task   Display all backed up images.

> **Command:** `q image`

Task   Display all backed up images owned by **kutras** at node avalon.

> **Command:** `query image -fromnode=avalon -fromowner=kutras`

Task   Display active and inactive version of the /usr image.

> **Command:** `q i /usr -inactive`

Task   Display all images that are contained within the backup set weekly_backup_data.32145678.

> **Command:** `query image -backupsetname=weekly_backup_data.32145678`

Related information

"Restore data from a backup set" on page 179

Query Inclexcl

The **query inclexcl** command displays a list of include-exclude statements in the order in which they are processed during backup and archive operations. The list displays the type of option, the scope of the option (archive, all, etc.), and the name of the source file.

You can test the validity of patterns you want to use in your include-exclude list before you actually insert them in your options file. See the *test pattern* explanation below.

Use the detail option to display the management class that is associated with an include-exclude statement.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Query INCLexcl─┬───────────────┬─┬─────────┬──────────►◄
                   └─test pattern──┘ └─-DETail─┘
```

Parameters

test pattern
> Use for testing the validity of patterns you want to use in your include-exclude list. When you use a test pattern with this command, the following occurs:
>
> - The internal include-exclude list is *not* displayed
> - The pattern is processed as if it had come from an include-exclude statement, including all the usual error checking
> - The pattern is displayed as it would appear in the include-exclude list
>
> If the test pattern has no errors, the compiled pattern result is the same as the test pattern.

-DETail
> Display the management class that is associated with the include-exclude statement.

Examples

Here is an example include-exclude list for client-side data deduplication:

```
Exclude Dedup /.../file2 dsm.opt
IEOBJTYPE: FILE
Exclude Dedup /home/*/* dsm.opt
IEOBJTYPE: IMAGE
```

Task Display a list of include-exclude statements.

   Command: `query inclexcl`

Task Display a list of include-exclude statements. Display the management class that is associated with each statement.

   Command: `query inclexcl -detail`

Task Test the validity of this pattern: /.../?x?/*.log

> **Command:** `query inclexcl /.../?x?/*.log`

> **Note:** Some system files are excluded explicitly by Tivoli Storage Manager. You can use the **query inclexcl** command to display a list of these files.

Query Mgmtclass

The **query mgmtclass** command displays information about the management classes available in your active policy set.

Your administrator defines management classes that contain attributes controlling whether a file is eligible for backup or archive services. Management classes also determine how backups and archives are managed on the server.

Your active policy set contains a default management class; it can contain any number of additional management classes. You can assign specific management classes to files using include options that are located in the client user-options file (dsm.opt). If you do not assign a management class to a file, Tivoli Storage Manager uses the default management class.

When you archive files, you can override the assigned management class by using the archmc option.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Query Mgmtclass────────────────────────────────────────►◄
                     └──options──┘
```

Parameters

Table 114. Query Mgmtclass command: Related options

Option	Where to use
detail "Detail" on page 296	Command line only.
fromnode "Fromnode" on page 339	Command line only.

Examples

Task   Display default and available management classes.

> **Command:** `query mgmtclass`

Query Node

The **query node** command displays all the nodes for which an administrative user ID has authority to perform operations.

The authorized administrative user ID should have at least client owner authority over the client workstation node they are using either from the command line or from the web.

When using an interactive command-line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID.

Use the type option to specify the type of node to filter for. The following are the valid values.

- nas
- client
- server
- any

The default is *any*.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Query Node──────────────────────────────────────────────────────►◄
                  └──options──┘
```

Parameters

Table 115. Query Node command: Related options

Option	Where to use
type "Type" on page 448	Command line only.
scrolllines "Scrolllines" on page 410	Client user options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 411	Client user options file (dsm.opt) or command line.

Examples

Task   Display all NAS nodes.

   Command: `query node -type=nas`

Task   Display all client nodes that are backup-archive clients.

   Command: `query node -type=client`

Related information

"Type" on page 448

Query Options

Use the **query options** command to display all or part of your options and their current settings relevant to the command-line client.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Query Options─────────────── ──pattern──────────────────────────────►◄
                  └──options──┘
```

Parameters

pattern
> An optional character string which can include wildcards. Use this argument to specify a subset of options. The default is to display all options.

Table 116. Query Options command: Related options

Option	Where to use
scrolllines "Scrolllines" on page 410	Client user options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 411	Client user options file (dsm.opt) or command line.

Examples

Task Display all options and their values.

> **Command:** `query options`

Task Display only options beginning with *comm*.

> **Command:** `query options comm*`

Task Display the value of the replace option.

> **Command:** `query options replace`

Query Restore

The **query restore** command displays a list of your restartable restore sessions in the server database. The list contains these fields: owner, replace, subdir, preservepath, source, and destination.

A restartable restore session is created when a wildcard restore command fails because of network outage, client failure, server outage, or a similar problem. When such a failure occurs, the file space is locked on the server and its files cannot be moved off the sequential volumes of the server. To unlock the file space, either restart the restore and allow it to complete (**query restore** command), or cancel the restore (**cancel restore** command). Use **query restore** to determine if you have any restartable restore sessions and which file spaces are affected.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Query Restore──────────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this command.

Examples

Task Display your restartable restore session in the server database.

  Command: `query restore`

Query Schedule

The **query schedule** command displays the events scheduled for your node. Your administrator can set up schedules to perform automatic backups and archives for you. To plan your work, use this command to determine when the next scheduled events occur.

The enhanced **query schedule** command on a Tivoli Storage Manager Version 5.3 and above client reports new parameters. The **query schedule** command prior to Tivoli Storage Manager Version 5.3 is referred to as classic. Refer to the Administrator's Guide and Administrator's Reference for your operating system for more information on the classic and enhanced commands.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Query SChedule──────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this command.

Examples

Task Display your scheduled events.

  Command: `query schedule`

Query Session

The **query session** command displays information about your session, including the current node name, when the session was established, server information, and server connection information.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Query SEssion───────────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this command.

Examples

Task   Display your session information.

Command: `query session`

A sample **query session** display follows:

```
Server Name.............: HALLEY_SERVER1
Server Type.............: Windows
Archive Retain Protect..: "No"
Server Version..........: Ver. 6, Rel. 2, Lev. 0.0
Last Access Date........: 09/03/2009 09:08:13
Delete Backup Files.....: "No"
Delete Archive Files....: "Yes"
Deduplication...........: "Server Only"

Node Name...............: HALLEY
User Name...............:
```

The possible client-side deduplication values are:
- None
  - This is displayed when connected to pre-V6.1 Tivoli Storage Manager server
- Server Only
- Client Or Server

Task

A sample **query session** display with LAN-free enabled follows:

```
TSM Server Connection Information

Server Name.............: TEMPLAR
Server Type.............: AIX
Archive Retain Protect..: "No"
Server Version..........: Ver. 6, Rel. 1, Lev. 4.0
Last Access Date........: 08/12/10   22:10:15
Delete Backup Files.....: "No"
Delete Archive Files....: "Yes"

Node Name...............: LAN2
User Name...............: root

Storage Agent Name......: TEMPLAR_STA
Storage Agent Type......: AIX
Storage Agent Version...: Ver. 6, Rel. 1, Lev. 3.3
```

Query Systeminfo

Use the **query systeminfo** command to gather information and output this information to a file or the console.

This command is intended primarily as an aid for IBM support to assist in diagnosing problems, although users who are familiar with the concepts addressed by this information might also find it useful.

If you use the *console* option, no special formatting of the output is performed to accommodate screen height or width. Therefore, the console output can be difficult to read due to length and line-wrapping. If the console output is difficult to read,

use the filename option with the **query systeminfo** command to allow the output to be written to a file that can subsequently be submitted to IBM support.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──Query SYSTEMInfo──┬──── item ──┬──────────────────────────►◄
                      └◄───────────┘
                              └── options ──┘
```
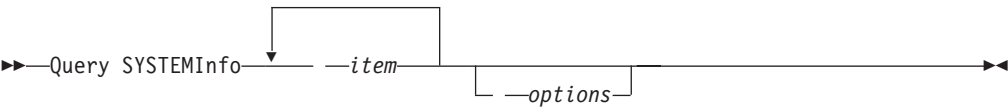
Parameters

item
  Specifies one or more items from which you want to gather information and output the information to the file name that you specify with the filename option or to the console. The default is to gather all items.

You can gather information on one or more of the following items:
- DSMOPTFILE - The contents of dsm.opt file.
- DSMSYSFILE - The contents of the dsm.sys file.
- ENV - Environment variables.
- ERRORLOG - The Tivoli Storage Manager error log file.
- FILE - Attributes for the file name that you specify.
- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- OPTIONS - Compiled options.
- OSINFO - Name and version of the client operating system (includes ULIMIT information for UNIX).
- POLICY - Policy set dump.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually dsmsched.log).
- CLUSTER - AIX cluster information.
- ENCRYPT - Available encryption methods.

Note:
1. Use the filename option to specify a file name in which to store the information gathered from the items you specify. If you do not specify a file name, by default the information is stored in the /Library/Application Support/tivoli/tsm/client/ba/bin/dsminfo.txt file (for Mac OS X) or the dsminfo.txt file (for other UNIX and Linux).
2. Use the console option if you want to output the information to the console.

Table 117. Query Systeminfo command: Related options

Option	Where to use
console "Console" on page 286	Command line only.

Table 117. Query Systeminfo command: Related options (continued)

Option	Where to use
filename "Filename" on page 335	Command line only.

Examples

Task   Gather and store the contents of the dsm.opt file and the Tivoli Storage Manager error log file in the `tsminfo.txt` file.

Command: `query systeminfo dsmoptfile errorlog -filename=tsminfo.txt`

Related information

"Filename" on page 335

"Console" on page 286

Restart Restore

The **restart restore** command displays a list of your restartable restore sessions in the server database.

You can only restart one restartable restore session at a time. Run the **restart restore** command again to restart additional restores.

The restarted restore uses the same options you used in the failed restore. The restarted restore continues from the point at which the restore previously failed.

To cancel restartable restore sessions, use the **cancel restore** command. Use the **restart restore** command when:
- Restartable restore sessions lock the file space at the server so that files cannot be moved off the sequential volumes of the server.
- You cannot back up files affected by the restartable restore.

Options from the failed session supersede new or changed options for the restarted session.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──RESTArt Restore────────────────────────────────────────────►◄
```

Parameters

There are no parameters for this command.

Examples

Task   Restart a restore.

Restore

The **restore** command obtains copies of backup versions of your files from a Tivoli Storage Manager server, or inside a backup set.

To restore files, specify the directories or selected files, or select the files from a list. Restore files to the directory from which you backed them up or to a different directory. Tivoli Storage Manager uses the preservepath option with the *subtree* value as the default for restoring files.

Note:
1. On UNIX and Linux systems when a symbolic link is created its modification time is set to the current system time and cannot be changed. So, when restoring a symbolic link its modification date and time is set to the date and time of the restore, not to the date and time the link had when it was backed up. As a result, Tivoli Storage Manager backs up the symbolic link during the next incremental backup because its modification time changed since the last backup.
2. When restoring a directory, its modification date and time is set to the date and time of the restore, not to the date and time the directory had when it was backed up. This is because Tivoli Storage Manager restores the directories first, then adds the files to the directories.

If you set the subdir option to *yes* when restoring a specific path and file, Tivoli Storage Manager recursively restores *all* subdirectories under that path, and any instances of the specified file that exist under *any* of those subdirectories.

When you restore an entire directory or directory tree, and you do not specify the inactive, latest, pick, todate, and fromdate options on the **restore** command, Tivoli Storage Manager tracks which objects are restored. If the restore process is interrupted for any reason, you can restart the restore at the point of interruption by entering the **restart restore** command. It is possible to create more than one restartable restore session. Restores are only restartable if the *filespec* is fully wildcarded. For example, for a restore which is restartable, enter:

```
dsmc rest /home/* -sub=yes
```

For a restore which is not restartable, enter:

```
dsmc rest "/Users/user1/file?.c" -sub=yes
```

Use the **query restore** command to display a list of your restartable restore sessions in the server database. Further backups of the file system cannot be performed unless the restartable restore completes using the **restart restore** command or is cancelled using the **cancel restore** command.

```
dsmc rest "/Users/user1/file?.c" -sub=yes
```

Supported Clients

This command is valid for all clients.

Syntax

```
                              ┌─FILE─┐
►►──REStore─────┴──────┴──────┬─ sourcefilespec ─┬──────────►
                 └─ options ─┘  └─ "sourcefilespec" ─┘

►──┬─────────────────────┬──┬─BACKUPSETName=─┬─ backupsetname ─┬──►
    └─ destinationfilespec ─┘                  ├─ localfilename ─┤
                                               └─ tapedevice ────┘

►──┬──────────────────────────┬─────────────────────────────►◄
    └─LOCation=─┬─ server ─┬──┘
                ├─ file ───┤
                └─ tape ───┘
```
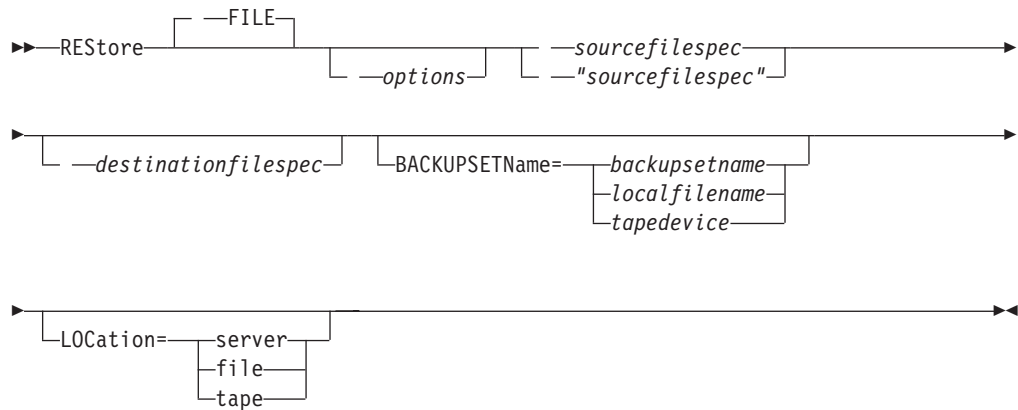
Parameters

file

This parameter specifies that the source file specification is an explicit file name. This parameter is required when you restore a file name from the current path, when you do not specify a relative or absolute path, and when the file name conflicts with one of the reserved **restore** command keywords, such as **restore backupset**.

sourcefilespec

Specifies the path and file name in storage that you want to restore. Use wildcard characters to specify a group of files or all the files in a directory.

{_filespacename_**}**

Specifies the file space (enclosed in braces) on the server that contains the files you want to restore. This is the name on the workstation drive from which the files were backed up.

Specify the file space name if the drive label name has changed or if you are restoring files backed up from another node with drive labels that are different from yours.

destinationfilespec

Specifies the path and file name where you want to place the restored files. If you do not specify a destination, Tivoli Storage Manager restores the files to the original source path.

- If the *sourcefilespec* names a single file, the *destinationfilespec* can be a file or a directory. If you are restoring a single file, you can optionally end the specification with a file name if you want to give the restored file a new name.
- If the *sourcefilespec* is wildcarded or **subdir**=*yes* is specified, the *destinationfilespec* must be a directory and end with a directory delimiter (\).

Note: If the destination path or any part of it does not exist, Tivoli Storage Manager creates it.

Note: If you do not specify a destination, Tivoli Storage Manager determines whether the original file system can be reached. If the original file system cannot be reached, Tivoli Storage Manager will not restore the file. In this case, you can specify a different destination and retry the command.

BACKUPSETName=
> Specifies the name of a backup set. This parameter is optional. If you specify the **backupsetname** parameter with the **restore** command, you cannot use the pick option.
>
> The value of **backupsetname** depends on the location of the backup set, and corresponds to one of the three choices below.
>
> **backupsetname**
>> Specifies the name of the backup set from the Tivoli Storage Manager server. If the **location** parameter is specified, you must set `-location=server`. If the backup set resides in Tivoli Storage Manager server storage, the backup set must have a TOC.
>
> **localfilename**
>> Specifies the file name of the first backup set volume. You must set `-location=file`.
>
> **tapedevice**
>> Specifies the name of the tape device containing the backup set volume. You must use a Windows native device driver, not the device driver provided by Tivoli. You must set `-location=tape`.

LOCation=
> Specifies where Tivoli Storage Manager searches for the backup set. If you do not specify the location parameter, the client searches for backup sets on the Tivoli Storage Manager server.
>
> **server** Specifies that Tivoli Storage Manager searches for the backup set from the server. This is the default.
>
> **file** Specifies that Tivoli Storage Manager searches for the backup set from a local file.
>
> **tape** Specifies that Tivoli Storage Manager searches for the backup set from a local tape device.

Table 118. Restore command: Related options

Option	Where to use
dateformat "Dateformat" on page 289	Client user options file (dsm.opt) or command line.
dirsonly "Dirsonly" on page 300	Command line only.
filelist "Filelist" on page 333	Command line only.
filesonly "Filesonly" on page 336	Command line only.
followsymbolic "Followsymbolic" on page 336	Client user options file (dsm.opt) or command line.
fromdate "Fromdate" on page 338	Command line only.
fromnode "Fromnode" on page 339	Command line only.
fromowner "Fromowner" on page 339	Command line only.
fromtime "Fromtime" on page 340	Command line only.

Table 118. Restore command: Related options (continued)

Option	Where to use
ifnewer "Ifnewer" on page 346	Command line only.
inactive "Inactive" on page 348	Command line only.
latest "Latest" on page 362	Command line only.
numberformat "Numberformat" on page 376	Client user options file (dsm.opt) or command line.
pick **Note:** If you specify the **backupsetname** parameter with the **restore** command, you cannot use the pick option."Pick" on page 382	Command line only.
pitdate "Pitdate" on page 383	Command line only.
pittime "Pittime" on page 384	Command line only.
preservepath "Preservepath" on page 390	Command line only.
replace "Replace" on page 397	Client user options file (dsm.opt) or command line.
subdir "Subdir" on page 432	Client user options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 433	Client user options file (dsm.opt) or command line.
timeformat "Timeformat" on page 442	Client user options file (dsm.opt) or command line.
todate "Todate" on page 445	Command line only.
totime "Totime" on page 446	Command line only.

Examples

Task   Restore a single file named `budget` in the `/Users/user1/Documents` directory.

Command: `restore /home/devel/projecta/budget`

Task   Restore a single file named `budget` which resides in the current directory.

Command: `restore file budget`

Task   Restore all files with a file extension of `.c` from the `/home/devel/projecta` directory.

Command: `restore "/home/devel/projecta/*.c"`

Task   Restore files in the `/user/project` directory. Use the pick and inactive options to select active and inactive backup versions.

Command: `restore "/user/project/*" -pick -inactive`

Task Restore all files from the /home/devel/projecta directory that end with the character .c to the /home/newdevel/projectn/projecta directory. If the projectn or the projectn/projecta directory does not exist, it is created.

Command: `restore "/home/devel/projecta/*.c" /home/newdevel/projectn/`

Task Restore all files in the /home/mydir directory to their state as of 1:00 PM on August 17, 2002.

Command: `restore -pitd=8/17/2002 -pitt=13:00:00 /home/mydir/`

Task Restore all objects in the /home/myid/ directory. Since this restore is fully wildcarded, if the restore process is interrupted, a restartable restore session is created.

Command: `res /home/myid/*`

Task Restore all files in the /home/mydir directory to their state as of 1:00 PM on August 17, 2002.

Command: `restore -pitd=8/17/2002 -pitt=13:00:00 /home/mydir/`

Related information

"Restore data from a backup set" on page 179

"Preservepath" on page 390

"File system and ACL support" on page 125

Restore from file spaces that are not Unicode-enabled

If you want to restore from file spaces that are not Unicode-enabled, you must specify the source on the server and a destination on the client, prior to installing the Unicode-enabled client.

Note: This Unicode section applies only to Mac OS X.

For example, assume that Jaguar is the name of your startup disk and you back up all of the .log files in the /Users/user5/Documents directory. Before the backup takes place, the server renames the file space to Jaguar_OLD. The backup places the data specified in the current operation into the Unicode-enabled file space named /. The new Unicode-enabled file space now contains only the /Users/user5/ Documents directory and the *.log files specified in the operation.

If you want to restore a file from the *renamed* (old) file space to its original location, you must enter both the source and destination as follows:

```
restore Jaguar_OLD/Users/user5/Documents
  /mylog.log /Users/user5/Documents/
```

Restore Backupset

The **restore backupset** command restores a backup set from the Tivoli Storage Manager server, a local file, or a local tape device. You can restore the entire backup set, or, in some cases, specific files within the backup set.

Supported Clients

This command is valid for all clients.

Syntax

```
>>--REStore Backupset----------------------------------------------------->
                      |                              |
                      |  |-{filespacename}-|  sourcefilespec |
                      |--SYSTEMOBJECT------------------|
                      |--SYSTEMSTATE-------------------|

>--------------------------------------------- -BACKUPSETName= ---backupsetname----->
   |-destinationfilespec-|  |-options-|                    |-localfilename-|
                                                           |-tapedevice----|

>------------------------------------------------------------------------><
   |- -LOCation= ---server---|
                 |-file---|
                 |-tape---|
```

Parameters

{*filespacename*}
Specifies the file space (enclosed in braces) on the server that contains the files you want to restore. This is the name on the workstation drive from which the files were backed up, or the virtual file space name for a group.

Specify a file space name when you restore a backup set containing a group.

Specify a file space name when the **sourcefilespec** does not exist on the target computer. This can occur if the drive label name has changed or if you are restoring files backed up from another node with drive labels that are different from yours.

sourcefilespec
Specifies the source path of a portion of the backup set. The default is to restore the entire backup set.

SYSTEMOBJECT
Specifies that you want to restore all system objects from the backup set. You cannot restore individual system objects from the backup set. This parameter is for the Windows 2003 and Windows XP clients only. To restore individual system objects from a backup set, use the **restore systemobject** command.

SYSTEMSTATE
Specifies that you want to restore the entire backup of the system state from the backup set. You cannot restore individual system state components from the backup set. This parameter is valid only for systems that support system state. To restore individual system state components from a backup set, use the **restore systemstate** command.

destinationfilespec
Specifies the destination path for the restored files. If you do not specify a *sourcefilespec*, you cannot specify a *destinationfilespec*. If you do not specify a destination, Tivoli Storage Manager restores the files to the original source path. If you are restoring more than one file, you must end the file specification with a directory delimiter (/), otherwise, Tivoli Storage Manager assumes the last name is a file name and reports an error. If you are restoring a

single file, you can optionally end the destination file specification with a file name if you want to give the restored file a new name. When the **sourcefilespec** does not exist on the target workstation, you must specify **destinationfilespec**.

-BACKUPSETName=
Specifies the name of the backup set from which to perform a restore operation. You cannot use wildcard characters to specify the backup set name. The value of backupsetname depends on the location of the backup set, and corresponds to one of the three choices below.

backupsetname
Specifies the name of the backup set on the server from which to perform a restore operation. If **location** option is specified, you must set `-location=server`.

localfilename
Specifies the file name of the first backup set volume. You must set `-location=file`.

tapedevice
Specifies the name of the tape device containing the backup set volume. You must use a Windows native device driver, not the device driver provided by Tivoli. You must set `-location=tape`.

-LOCation=
Specifies the location of the backup set. If you do not specify the location parameter, the Tivoli Storage Manager client searches for backup sets on the Tivoli Storage Manager server. If you specify the location parameter, the value must be one of the three choices below.

server Specifies that the backup set is on the Tivoli Storage Manager server. This is the default.

file Specifies that the backup set is on an available file system.

tape Specifies that the backup set is on an available tape device.

Table 119. Restore Backupset command: Related options

Option	Where to use
dirsonly "Dirsonly" on page 300	Command line only.
filesonly "Filesonly" on page 336	Command line only.
ifnewer "Ifnewer" on page 346	Command line only.
preservepath "Preservepath" on page 390	Command line only.
quiet "Quiet" on page 396	Client user options file (dsm.opt) or command line.
replace "Replace" on page 397	Client user options file (dsm.opt) or command line.
subdir "Subdir" on page 432	Client user options file (dsm.opt) or command line.

Examples

Task Restore the entire backup set called `monthly_financial_data.87654321` from the server.

Command:

```
dsmc restore backupset
  -backupsetname=monthly_financial_data.87654321
   -loc=server
```

Task Restore the entire backup set contained in the file: /home/budget/weekly_budget_data.ost.

Command:

```
dsmc restore backupset
  -backupsetname="/home/budget/weekly_budget_data.ost"
   -loc=file
```

Task Restore the entire backup set from the /dev/rmt0 device.

Command:

```
dsmc restore backupset
  "-backupsetname=/dev/rmt0" -loc=tape
```

Task Restore a single file named /home/jones/budget.dev from the /dev/rmt0 tape device, to the original source path.

Command:

```
dsmc restore backupset
  -backupsetname=/dev/rmt0 "/home/jones/budget.dev"
   -loc=tape
```

Task Restore all files in the budget directory that contain a file extension of .txt from the tape(s) on the /dev/rmt0 device, to the original source path.

Command:

```
dsmc restore backupset "/home/budget/*.txt"
  -backupsetname=/dev/rmt0 -loc=tape
```

Task Restore the entire backup set contained in local file "/home/jones/bset01.file"

Command:

```
dsmc restore backupset
  -backupsetname="/home/jones/bset01.file"
   -loc=file
```

Task Restore groups from the backup set `mybackupset.12345678` on the Tivoli Storage Manager server to the /home/devel/projectb directory. The groups' virtual file space is `accounting`.

Command:

```
dsmc restore backupset {/accounting}/*
   /home/devel/projectb/
  -backupsetname=mybackupset.12345678 -loc=server
   -subdir=yes
```

Task Restore groups from the local backup set `mybackupset.ost` to the
`/home/devel/projectb/` directory. The groups' virtual file space is
`accounting`.

Command:

```
dsmc restore backupset {/accounting}/*
    /home/devel/projectb/
    -backupsetname=mybackupset.ost
    -loc=server -subdir=yes
```

Related information

"Restore data from a backup set" on page 179

Restore backup sets: considerations and restrictions

This topic lists some considerations and restrictions that you must be aware of
when restoring backup sets.

Backup set restore considerations

Consider the following when restoring backup sets:
- If the object you want to restore was generated from a Tivoli Storage Manager
  node whose name is different from your current node, specify the original node
  name with the **filespacename** parameter on any of the restore commands.
- If you are unable to restore a backup set from portable media, check with your
  Tivoli Storage Manager administrator to ensure that the portable media was
  created on a device using a compatible format.
- If you use the **restore backupset** command on the initial command line with the
  parameter `-location=tape` or `-location=file`, the client does not attempt to
  contact the Tivoli Storage Manager server.
- When restoring a group from a backup set:
  - The entire group, or all groups, in the virtual file space are restored. You
    cannot restore a single group by specifying the group name, if there are
    several groups in the same virtual file space. You cannot restore a part of a
    group by specifying a file path.
  - Specify a group by using the following values:
    - Specify the virtual file space name with the **filespacename** parameter.
    - Use the subdir option to include subdirectories.
- Limited support is provided for restoring backup sets from tape devices attached
  to the client system. A native device driver provided by the device manufacturer
  must always be used. The device driver provided by Tivoli to be used with the
  Tivoli Storage Manager server cannot be used on the client system for restoring
  local backup sets.
- If a backup set contains files from several owners, the backup set itself is owned
  by the root user ID, and non-root user IDs cannot see the backup set. In this
  case, non-root user IDs can restore their files by obtaining the backup set name
  from the Tivoli Storage Manager administrator. Non-root users can restore only
  their own files.
- To enable the client GUI to restore a backup set from a local device, without
  requiring a server connection, use the localbackupset option.

Backup set restore restrictions

Be aware of the following restrictions when restoring backup sets:

- A backup set data that was backed up with the API cannot be restored or used.
- You cannot restore image data from a backup set using the **restore backupset** command. You can restore image data from a backup set only with the **restore image** command.
- You cannot restore image data from a local backup set (`location=tape` or `location=file`). You can restore image data from a backup set only from the Tivoli Storage Manager server.

Related information

"Localbackupset" on page 363

"Restore" on page 525

"Restore Image" on page 539

"Restore Backupset" on page 529

Restore backup sets in a SAN environment

You can restore backup sets in a storage area network (SAN) in the following ways:

- If the backup set is on a SAN-attached storage device, specify the device using the *filename* parameter and use the location=*tape* option, where applicable. Tivoli Storage Manager restores the backup set directly from the SAN-attached storage device, gaining high-speed restore performance.

    Note: You must ensure that the correct tape is mounted in the SAN-attached tape drive prior to issuing the **restore** command. The backup-archive client will not initiate a SCSI autochanger to mount the tape automatically.

- If the backup set is not on local media or a SAN-attached storage device, you can specify the backup set using the backupsetname option. Use the location=*server* option to restore the backup set directly from the server using the LAN.

Restore Backupset without the backupsetname parameter

The **restore backupset** command can be used without the **backupsetname** parameter.

The preferred syntax for **restore backupset** command requires the **backupsetname** parameter. Prior to the introduction of the **backupsetname** parameter, the Tivoli Storage Manager client restored backup sets with a different syntax. The previous syntax is currently supported, but whenever possible, follow the syntax that requires the **backupsetname** parameter. The previous syntax is documented for those cases when it cannot be replaced by the preferred syntax.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──REStore Backupset─────────────────────────────────────────────────────────►
                     ┌─────────────────────┐  ─sourcefilespec─
                     └─{filespacename}─┘
                     ├─SYSTEMOBJECT──────────┤
                     └─SYSTEMSTATE───────────┘

►──────────────────────────────────────────────────────────────────────────────►
   └─destinationfilespec─┘  └─options─┘    ─backupsetname─
                                           ├─localfilename─┤
                                           └─tapedevice────┘

►──────────────────────────────────────────────────────────────────────────────►◄
   └─LOCation=──┬─server─┬──┘
               ├─file───┤
               └─tape───┘
```

Parameters

options
>   All options that are valid with the preferred syntax of **restore backupset** are valid with the legacy syntax of **restore backupset**.

{*filespacename*}
>   Specifies the file space (enclosed in braces) on the server that contains the files you want to restore. This is the name on the workstation drive from which the files were backed up, or the virtual file space name for a group.
>
>   Specify a file space name when you restore a backup set containing a group.
>
>   Specify a file space name when the **sourcefilespec** does not exist on the target computer. This can occur if the drive label name has changed or if you are restoring files backed up from another node with drive labels that are different from yours.

sourcefilespec
>   Specifies the source path of a portion of the backup set. The default is to restore the entire backup set.

SYSTEMOBJECT
>   Specifies that you want to restore all system objects from the backup set. You cannot restore individual system objects from the backup set. This parameter is for the Windows 2003 and Windows XP clients only. To restore individual system objects from a backup set, use the **restore systemobject** command.

SYSTEMSTATE
>   Specifies that you want to restore the entire backup of the system state from the backup set. You cannot restore individual system state components from the backup set. This parameter is valid only for systems that support system state. To restore individual system state components from a backup set, use the **restore systemstate** command.

destinationfilespec
>   Specifies the destination path for the restored files. If you do not specify a *sourcefilespec*, you cannot specify a *destinationfilespec*. If you do not specify a destination, Tivoli Storage Manager restores the files to the original source path. If you are restoring more than one file, you must end the file specification with a directory delimiter (/), otherwise, Tivoli Storage Manager assumes the last name is a file name and reports an error. If you are restoring a single file, you can optionally end the destination file specification with a file

name if you want to give the restored file a new name. When the
sourcefilespec does not exist on the target workstation, you must specify
destinationfilespec.

backupsetname
Specifies the name of the backup set from the Tivoli Storage Manager server. If
the **location** parameter is specified, you must set `-location=server`.

localfilename
Specifies the file name of the first backup set volume. You must set
`-location=file`.

tapedevice
Specifies the name of the tape device containing the backup set volume. You
must use a Windows native device driver, not the device driver provided by
Tivoli. You must set `-location=tape`.

LOCation=
Specifies the location of the backup set. If you do not specify the location
parameter, the Tivoli Storage Manager client searches for backup sets on the
Tivoli Storage Manager server. If you specify the location parameter, the value
must be one of the three choices below.

server Specifies that the backup set is on the Tivoli Storage Manager server.
This is the default.

file Specifies that the backup set is on an available file system.

tape Specifies that the backup set is on an available tape device.

Examples

Task Restore the entire backup set called `monthly_financial_data.87654321`
from the server.

Command: `dsmc restore backupset monthly_financial_data.87654321
-loc=server`

Task Restore the entire backup set contained in the file: `/home/budget/
weekly_budget_data.ost`.

Command: `dsmc restore backupset "/home/budget/
weekly_budget_data.ost" -loc=file`

Task Restore the entire backup set from the `/dev/rmt0` device.

Command: `dsmc restore backupset "/dev/rmt0" -loc=tape`

Task Restore a single file named `/home/jones/budget.dev` from the `/dev/rmt0`
tape device, to the original source path.

Command: `dsmc restore backupset /dev/rmt0 "/home/jones/budget.dev"
-loc=tape`

Task Restore all files in the budget directory that contain a file extension of `.txt`
from the tape(s) on the `/dev/rmt0` device, to the original source path.

Command: `dsmc restore backupset /dev/rmt0 "/home/budget/*.txt"
-loc=tape`

Task Restore the entire backup set contained in local file `"/home/jones/
bset01.file"`

Command: `dsmc restore backupset "/home/jones/bset01.file"
-loc=file`

Task Restore groups from the backup set `mybackupset.12345678` on the Tivoli Storage Manager server to the `/home/devel/projectb` directory. The groups' virtual file space is `accounting`.

Command: `dsmc restore backupset mybackupset.12345678 {/accounting}/* /home/devel/projectb/ -loc=server -subdir=yes`

Task Restore groups from the local backup set `mybackupset.ost` to the `/home/devel/projectb/` directory. The groups' virtual file space is `accounting`.

Command: `dsmc restore backupset mybackupset.ost {/accounting}/* /home/devel/projectb/ -loc=server -subdir=yes`

Related information

Restore Group

Use the **restore group** command to restore specific members or all members of a group backup.

Note:
1. Use the pick option to display a list of groups from which you can select one group to restore.
2. Use the showmembers option with the pick option to display and restore one or more members of a group. In this case, you first select the group from which you want to restore specific members, then you select one or more group members to restore.
3. You can restore a group from a backup set.

Supported Clients

This command is valid for all clients, except Mac OS X.

Syntax

```
►►──REStore GRoup─────────────────sourcefilespec────────────────────►◄
                    └options─┘                 └destinationfilespec─┘
```

Parameters

sourcefilespec
Specifies the virtual file space name and the group name on the server that you want to restore.

destinationfilespec
Specifies the path where you want to place the group or one or more group members. If you do not specify a destination, the client restores the files to their original location.

Table 120. Restore Group command: Related options

Option	Where to use
backupsetname "Backupsetname" on page 277	Command line only.
followsymbolic "Followsymbolic" on page 336	Client user options file (dsm.opt) or command line.
fromdate "Fromdate" on page 338	Command line only.
fromnode "Fromnode" on page 339	Command line only.
fromowner "Fromowner" on page 339	Command line only.
fromtime "Fromtime" on page 340	Command line only.
ifnewer "Ifnewer" on page 346	Command line only.
inactive "Inactive" on page 348	Command line only.
latest "Latest" on page 362	Command line only.
pick "Pick" on page 382	Command line only.
pitdate "Pitdate" on page 383	Command line only.
pittime "Pittime" on page 384	Command line only.
preservepath "Preservepath" on page 390	Command line only.
replace "Replace" on page 397	Client user options file (dsm.opt) or command line.
showmembers "Showmembers" on page 416 (does not apply to Mac OS X)	Command line only.
subdir "Subdir" on page 432	Client user options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 433	Client user options file (dsm.opt) or command line.
todate "Todate" on page 445	Command line only.
totime "Totime" on page 446	Command line only.

Examples

Task   Restore all members in the /virtfs/group1 group backup to their original location on the client system.

Command:

```
restore group /virtfs/group1
```

Task Display all groups within the /virtfs virtual file space. Use the showmembers option to display a list of group members from which you can select one or more to restore.

> **Command:**
>
> ```
> restore group /virtfs/
>    * -pick -showmembers
> ```

Task Display a list of groups within the /virtfs virtual file space from which you can select one or more groups to restore.

> **Command:**
>
> ```
> restore group /virtfs/* -pick
> ```

Related information

"Restore Backupset" on page 529

Restore Image

The **restore image** command restores a file system or raw volume image that was backed up using the **backup image** command.

The restore obtains the backup image from a Tivoli Storage Manager server, or inside a backup set from the Tivoli Storage Manager server, when the backupsetname option is specified. This command can restore an active base image, or a point-in-time base image, with associated incremental updates.

Note:
1. Using the incremental option with the **restore image** command to perform a dynamic image backup is not supported.
2. If you use HSM for Windows or HSM for UNIX, and you restore a file system image backup and plan to run reconciliation, you must restore the files that have been backed up after the image backup. Otherwise, migrated files that have been created after the image backup expire from the HSM archive storage on the Tivoli Storage Manager server.

You can use the verifyimage option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log.

If bad sectors are present on the target volume, you can use the imagetofile option with the **restore image** command to specify that you want to restore the source image to a file. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume.

Considerations:
- The API must be installed to use the **restore image** command.
- Image restore of the Sun QFS file system is not supported.
- Image restore is not supported for GPFS file systems on Linux x86/x86_64, Linux on POWER and Linux on System z.
- If you use the pick option, the following information is displayed for file system images backed up by the client:
  – Image Size

- Stored Size - This is the actual image size stored on the server. The stored image on the Tivoli Storage Manager server is the same size as the volume capacity.
  - File system type
  - Backup date and time
  - Management class assigned to image backup
  - Whether the image backup is an active or inactive copy
  - The image name
- If for some reason a restored image is corrupted, you can use the **fsck** tool to attempt to repair the image.

Supported Clients

This command is valid for AIX, HP-UX, all Linux clients, and Solaris.

Syntax

```
►►──REStore Image──┬─────────┬──┬──sourcefilespec───┬──────────────►
                   └─options─┘  └──"sourcefilespec"─┘

►──┬──────────────────────┬──────────────────────────────────────►◄
   └──destinationfilespec──┘
```
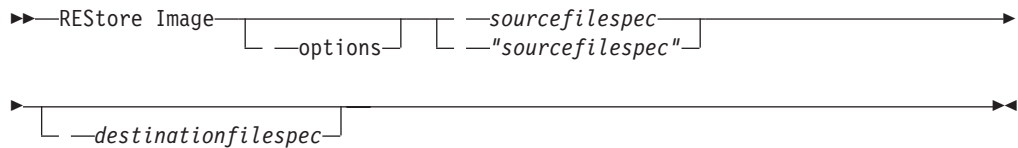
Parameters

sourcefilespec
Specifies the name of a source image file system to be restored. Only a single source image can be specified; you cannot use wildcard characters.

destinationfilespec
Specifies the name of an existing mounted file system or the path and file name to which the source file system is restored. The default is the original location of the file system.

Table 121. Restore Image command: Related options

Option	Where to use
backupsetname "Backupsetname" on page 277	Command line only.
dateformat "Dateformat" on page 289	Client user option file (dsm.opt) or command line.
deletefiles "Deletefiles" on page 295	Command line only.
fromnode "Fromnode" on page 339	Command line only.
fromowner "Fromowner" on page 339	Command line only.
imagetofile "Imagetofile" on page 348	Command line only.
inactive "Inactive" on page 348	Command line only.
incremental "Incremental" on page 358	Command line only.

Table 121. Restore Image command: Related options (continued)

Option	Where to use
noprompt "Noprompt" on page 376	Command line only.
pick "Pick" on page 382	Command line only.
pitdate "Pitdate" on page 383	Command line only.
pittime "Pittime" on page 384	Command line only.
timeformat "Timeformat" on page 442	Client user option file (dsm.opt) or command line.
verifyimage "Verifyimage" on page 451	Command line only.

The **restore image** command does not define or mount the destination file space. The destination volume must exist, must be large enough to hold the source, and if it contains a file system, must be mounted. If an image backup contains a file system, and you restore them to a different location, be aware of the following points:

- If the destination volume is smaller than the source volume, the operation fails.
- If the destination volume is larger than the source, after the restore operation you *lose* the difference between the sizes. The lost space can be recovered by increasing the size of the volume. This also increases the size of the restored volume.

Examples

Task   Restore the /home/test directory over which the logical volume is mounted, to its original location.

   Command: `dsmc rest image /home/test`

Task   Restore the /home/proj directory over which the logical volume is mounted, to its original location and apply the changes from the last incremental backup of the original image recorded at the server. The changes include deletion of files.

   Command: `dsmc restore image /home/proj -incremental -deletefiles`

Task   Restore the /usr file system to its original location. Use the verifyimage option to enable detection of bad sectors on the target volume.

   Command: `dsmc restore image /usr -verifyimage`

Task   If bad sectors present on the target volume, use the imagetofile option to restore the /usr file system to the /home/usr.img file to avoid data corruption.

   Command: `dsmc restore image /usr /home/usr.img -imagetofile`

Related information

"Verifyimage" on page 451

"Imagetofile" on page 348

Restore NAS

The **restore nas** command restores the image of a file system belonging to a Network Attached Storage (NAS) file server.

The NAS file server performs the outboard data movement. A server process performs the restore.

If you used the toc option with the **backup nas** command or the include.fs.nas option to save Table of Contents (TOC) information for each file system backup, you can use the **QUERY TOC** server command to determine the contents of a file system backup in conjunction with the **RESTORE NODE** server command to restore individual files or directory trees. You can also use the Web client to examine the entire file system tree and select files and directories to restore. If you do not save TOC information, you can still restore individual files or directory trees using the **RESTORE NODE** server command, provided that you know the fully qualified name of each file or directory and the image in which that object was backed up.

Use the nasnodename option to specify the node name for the NAS file server. When using an interactive command-line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID. The NAS node name identifies the NAS file server to the Tivoli Storage Manager server. You must register the NAS node name at the server. Place the nasnodename option in your client system-options file (dsm.sys). The value in the client system-options file is the default, but this value can be overridden on the command line.

You can use the pick option to display a list of NAS images owned by the NAS node you specify. From this list you can select one or more images to restore. If you select multiple images to restore using the pick option, do not use the monitor option or you will serialize the restores. To start multiple restore processes simultaneously when restoring multiple images, do not specify monitor=*yes*.

Use the monitor option to specify whether you want to monitor a NAS file system image restore and display processing information on your screen.

Use the **monitor process** command to display a list of current restore processes for all NAS nodes for which your administrative user ID has authority. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web.

Use the **cancel process** command to stop NAS restore processing.

Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.

Supported Clients

This command is valid for AIX, and Solaris clients only.

Syntax

```
►►──REStore NAS──┬──────────────┬──── sourcefilespec ──┬─────────────────────────┬──►◄
                 └── —options ──┘                       └── —destinationfilespec ──┘
```

Parameters

sourcefilespec

Specifies the name of the NAS file system image you want to restore. This parameter is required unless you use the pick option to display a list of NAS images from which to choose. You cannot use wildcard characters when specifying the *sourcefilespec*.

destinationfilespec

Specifies the name of an existing mounted file system on the NAS device over which you want to restore the image. This parameter is optional. The default is the original location of the file system on the NAS device.

Table 122. Restore NAS command: Related options

Option	Where to use
dateformat "Dateformat" on page 289	Client user option file (dsm.opt) or command line.
inactive "Inactive" on page 348	Command line only.
mode "Mode" on page 369	Command line only.
monitor "Monitor" on page 371	Command line only.
nasnodename "Nasnodename" on page 371	Client system options file (dsm.sys) or command line.
numberformat "Numberformat" on page 376	Client user option file (dsm.opt) or command line.
pick "Pick" on page 382	Command line only.
pitdate "Pitdate" on page 383	Command line only.
pittime "Pittime" on page 384	Command line only.
timeformat "Timeformat" on page 442	Client user option file (dsm.opt) or command line.

Examples

Task   Restore the NAS file system image /vol/vol1 to the /vol/vol2 file system on the NAS file server called **nas1**.

Command: `restore nas -nasnodename=nas1 /vol/vol1 /vol/vol2`

Task   Restore inactive NAS images.

Command: `restore nas -nasnodename=nas2 -pick -inactive`

Related information

"Nasnodename" on page 371

"Monitor" on page 371

"Cancel Process" on page 481

Retrieve

The **retrieve** command obtains copies of archived files from the Tivoli Storage Manager server. You can retrieve specific files or entire directories.

Use the description option to specify the descriptions assigned to the files you want to retrieve.

Use the pick option to display a list of your archives from which you can select an archive to retrieve.

Retrieve the files to the same directory from which they were archived, or to a different directory. Tivoli Storage Manager uses the preservepath option with the *subtree* value as the default for restoring files.

Note: When retrieving a directory, its modification date and time is set to the date and time of the retrieval, not to the date and time the directory had when it was archived. This is because Tivoli Storage Manager retrieves the directories first, then adds the files to the directories.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──RETrieve─────────────────┬─sourcefilespec─────┬───────────────────────►
                  └─options─┘  └─"sourcefilespec"─┘

►──────────────────────────────────────────────────────────────────────►◄
   └─destinationfilespec─┘
```
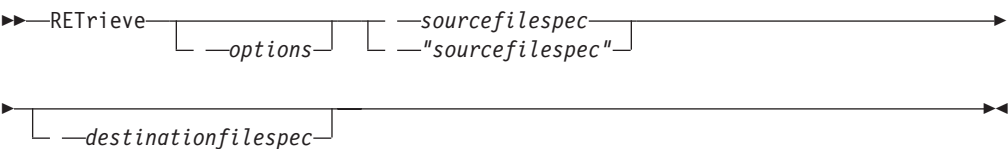
Parameters

sourcefilespec
> Specifies the path and file name that you want to retrieve. Use wildcard characters to specify a group of files or all the files in a directory.

destinationfilespec
> Specifies the path and file name where you want the files to reside. If you do not specify a destination, Tivoli Storage Manager restores the files to the original source path.
>
> **Note:** If you do not specify a destination, Tivoli Storage Manager determines whether the original file system can be reached. If the original file system cannot be reached, Tivoli Storage Manager will not restore the file.
>
> This failure can also occur if you remove the virtualmountpoint option from the dsm.sys file. In this case, you can specify a different destination or restore the original virtualmountpoint option to the dsm.sys file, restart the client, and retry the command.

Table 123. Retrieve command: Related options

Option	Where to use
dateformat "Dateformat" on page 289	Client user options file (dsm.opt) or command line.

Table 123. Retrieve command: Related options (continued)

Option	Where to use
description "Description" on page 295	Command line only.
dirsonly "Dirsonly" on page 300	Command line only.
filelist "Filelist" on page 333	Command line only.
filesonly "Filesonly" on page 336	Command line only
followsymbolic "Followsymbolic" on page 336	Client user options file (dsm.opt) or command line.
fromdate "Fromdate" on page 338	Command line only
fromnode "Fromnode" on page 339	Command line only.
fromowner "Fromowner" on page 339	Command line only
fromtime "Fromtime" on page 340	Command line only
ifnewer "Ifnewer" on page 346	Command line only
pick "Pick" on page 382	Command line only.
preservepath "Preservepath" on page 390	Command line only.
replace "Replace" on page 397	Client user options file (dsm.opt) or command line.
subdir "Subdir" on page 432	Client user options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 433	Client user options file (dsm.opt) or command line.
timeformat "Timeformat" on page 442	Client user options file (dsm.opt) or command line.
todate "Todate" on page 445	Command line only.
totime "Totime" on page 446	Command line only.

Examples

Task   Retrieve a single file named `budget`.

        Command: `retrieve /home/devel/projecta/budget`

Task   Retrieve all files with an extension of `.c` from the `/home/devel/projecta` directory.

        Command: `retrieve "/home/devel/projecta/*.c"`

Task   Retrieve all files in the `/home` directory.

Command: `retrieve /home/`

Task   Retrieve all files with a file extension of `.c` from the `/home/devel/projecta` directory to the `/home/newdevel/projectn/projecta` directory. If the `/projectn` or the `/projectn/projecta` directory does not exist, it is created.

Command: `retrieve "/home/devel/projecta/*.c" /home/newdevel/projectn/`

Task   Retrieve files in the `/user/project` directory. Use the pick option.

Command: `ret "/user/project/*" -pick`

Task   Retrieve all files archived from the `/proj` directory with the description "2002 survey results."

Command: `retrieve "/proj/*" -desc="2002 survey results"`

Task   Retrieve archived file `/home/devel/budget` with description "my budget" to the `/dev/rmt1` tape drive.

Command:

```
mkfifo fifo
dd if=fifo of=/dev/rmt1&
dsmc retrieve -replace=yes -description="mybudget"
/home/devel/budget fifo
```

Task   Retrieve a file from the renamed file space `Jaguar_OLD` to its original location. Enter both the source and destination as follows:

Command: `ret Jaguar_OLD/user5/Documents/myresume.doc /Users/user5/Documents/`

Related information

"Client options reference" on page 267

Retrieve archives from file spaces that are not Unicode-enabled

If you want to retrieve archives from file spaces that were renamed by the Unicode-enabled client, you must specify the source on the server and a destination on the client.

This section applies to Mac OS X only. For example, assume that `Jaguar` is the name of your startup disk and you archive all of the .log files in the /Users/user5/Documents directory. Before the archive takes place, the server renames the file space to `Jaguar_OLD`. The archive places the data specified in the current operation into the Unicode-enabled file space named /. The new Unicode-enabled file space now contains only the Users/user5/Documents directory and the *.log files specified in the operation.

If you want to retrieve a file from the *renamed* (old) file space to its original location, you must enter both the source and destination as follows:

```
retrieve Jaguar_OLD/Users/user5/Documents/mylog.log /Users/user5/Documents/
```

Schedule

The **schedule** command starts the client scheduler on your workstation. The client scheduler must be running before scheduled work can start.

Authorized User: The **schedule** command starts the client scheduler on your workstation. The client scheduler must be running before scheduled work can start.

Note:

1. The **schedule** command cannot be used if the managedservices option is set to *schedule* (does not apply to Mac OS X).
2. For Mac OSX only, to use the **schedule** command, specify managedservices *none* in the dsm.sys file.
3. This command is valid only on the initial command line. It is not valid in interactive mode or in a macro file.

If the schedmode option is set to polling, the client scheduler contacts the server for scheduled events at the hourly interval you specified with the queryschedperiod option in your client user-options file (dsm.opt). If your administrator sets the queryschedperiod option for all nodes, that setting overrides the client setting.

If you are using TCP/IP communications, the server can prompt your workstation when it is time to run a scheduled event. To do so, set the schedmode option to *prompted* in the client user-options file (dsm.opt) or on the **schedule** command.

You can use the sessioninitiation option with the **schedule** command to control whether the server or client initiates sessions through a firewall.

After you start the client scheduler, it continues to run and to start scheduled events until you press Ctrl+C, stop the scheduler process with the UNIX **kill** command, start the workstation again, or turn off the workstation to end it.

After you start the client scheduler, it continues to run and to start scheduled events until you press Ctrl+C, press the Q key twice, start the workstation again, or turn off the workstation to end it.

Note: You *cannot* enter this command in interactive mode.

Supported Clients

This command is valid for all clients.

Syntax
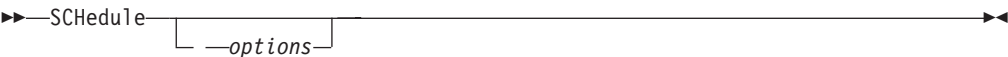
▶▶──SCHedule──┬────────────┬──────────────────────────────────▶◀
              └──*options*──┘

Parameters

Table 124. Schedule command: Related options

Option	Where to use
maxcmdretries "Maxcmdretries" on page 367	Client system options file (dsm.sys) or command line.
password "Password" on page 379	client user options file (dsm.opt)

Table 124. Schedule command: Related options (continued)

Option	Where to use
queryschedperiod "Queryschedperiod" on page 394	Client system options file (dsm.sys) or command line.
retryperiod "Retryperiod" on page 401	Client system options file (dsm.sys) or command line.
schedlogname "Schedlogname" on page 405	Client system options file (dsm.sys) or command line.
schedmode "Schedmode" on page 408	Client system options file (dsm.sys) or command line.
sessioninitiation "Sessioninitiation" on page 414	Client system options file (dsm.sys) or command line.
tcpclientport "Tcpclientport" on page 438	Client system options file (dsm.sys) or command line.

Examples

Task   Start the client scheduler.

  Command: `dsmc sch —password=notell`

Task   For AIX: Start the scheduler at system bootup time by entering this command in the /etc/inittab file. Ensure the ***passwordaccess*** option is set to *generate*.

  Command: `tsm::once:/usr/bin/dsmc sched > /dev/null 2>&1 #TSM`

Task   Interactively start the scheduler and keep it running in the background.

  Command: `nohup dsmc sched 2> /dev/null &`

When you run the **schedule** command, all messages regarding scheduled work are sent to the `dsmsched.log` file or to the file you specify with the schedlogname option in your client system-options file (dsm.sys). If you do not specify a directory path with the file name in the schedlogname option, the `dsmsched.log` resides in the current working directory, except for Mac OS X. For Mac OS X, the `dsmsched.log` resides in the `/Library/Logs/tivoli/tsm/` directory.

Important: To prevent log write failures and process termination in certain cases, set the DSM_LOG environment variable to name a directory where default permissions allow the required access.

Related information

"Sessioninitiation" on page 414

Selective

The **selective** command backs up files that you specify. If you damage or lose these files, you can replace them with backup versions from the server.

When you run a selective backup, all the files are candidates for backup unless you exclude them from backup, or they do not meet management class requirements for serialization.

During a selective backup, copies of the files are sent to the server even if they did not change since the last backup. This might result in more than one copy of the same file on the server. If this occurs, you might not have as many different down-level versions of the file on the server as you intended. Your version limit might consist of identical files. To avoid this, use the **incremental** command to back up only new and changed files.

You can selectively back up single files *or* directories. You can also use wildcard characters to back up groups of related files.

If you set the subdir option to *yes* when backing up a specific path and file, Tivoli Storage Manager recursively backs up *all* subdirectories under that path, and any instances of the specified file that exist under *any* of those subdirectories.

During a selective backup, a directory path might be backed up, even if the specific file that was targeted for backup is not found. For example the following command still backs up dir1 and dir2 even if the file bogus.txt does not exist.

```
selective /Users/user1/Documents/dir1/bogus.txt
selective "/dir1/dir2/bogus.txt"
```

If the **selective** command is retried because of a communication failure or session loss, the transfer statistics displays the number of bytes Tivoli Storage Manager attempts to transfer during *all* command attempts. Therefore, the statistics for bytes transferred might not match the file statistics, such as those for file size.

You can use the removeoperandlimit option to specify that Tivoli Storage Manager removes the 20-operand limit. If you specify the removeoperandlimit option with the **selective** command, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits.

Supported Clients
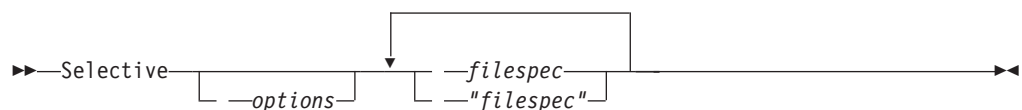
This command is valid for all clients.

Syntax

```
►►──Selective──┬──────────┬──▼──┬──filespec───┬──┬──────────►◄
               └──options─┘     └──"filespec"─┘
```

Parameters

filespec
    Specifies the path and file name that you want to back up. Use wildcard characters to select a group of files or all the files in a directory. When backing up a file system, specify the file system with a trailing slash; for example: /home/.

Table 125. Selective command: Related options

Option	Where to use
changingretries "Changingretries" on page 278	Client system options file (dsm.sys) or command line.
compressalways "Compressalways" on page 284	Client user options file (dsm.opt) or command line.
compression "Compression" on page 285	Client user options file (dsm.opt) or command line.
dirsonly "Dirsonly" on page 300	Command line only.
filelist "Filelist" on page 333	Command line only.
filesonly "Filesonly" on page 336	Command line only.
preservelastaccessdate "Preservelastaccessdate" on page 389	Client user options file (dsm.opt) or command line.
removeoperandlimit "Removeoperandlimit" on page 396	Command line only.
snapshotproviderfs "Snapshotproviderfs" on page 423	System-options file (dsm.sys) within a server stanza or with the include.fs option.
snapshotroot "Snapshotroot" on page 425	Command line only.
subdir "Subdir" on page 432	Client user options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 433	Client user options file (dsm.opt) or command line.

Examples

Task Back up the `proja` file in the `/home/devel` directory.

  Command: `selective /home/devel/proja`

Task Back up all files in the `/home/devel` directory whose file names begin with `proj`.

  Command: `selective "/home/devel/proj*"`

Task Back up all files in the `/home/devel` directory whose file names begin with `proj`. Back up the single file named budget in the `/user/home` directory.

  Command: `selective "/home/devel/proj*" /user/home/budget`

Task Back up the `/home` file system.

  Command: `selective /home/ -subdir=yes`

Task Assuming that you initiated a snapshot of the `/usr` file system and mounted the snapshot as `/snapshot/day1`, run a selective backup of the

/usr/dir1/sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name /usr.

Command: `dsmc sel "/usr/dir1/sub1/*" -subdir=yes -snapshotroot=/snapshot/day1`

Related information

Associate a local snapshot with a server file space

Use the snapshotroot option with the **selective** command in conjunction with an independent software vendor application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. The snapshotroot option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot.

AIX only: You can perform a snapshot-based selective backup by specifying the option snapshotproviderfs=*JFS2*.

Set Access

The **set access** command gives users at other nodes access to your backup versions or archived copies.

You can also use the **set access** command to give users at other nodes access to your backup images.

You can give another user access to a specific file or image, multiple files or images, or all files in a directory. When you give access to another user, that user can restore or retrieve your objects. Specify in the command whether you are giving access to archives or backups.

When a node is exported to another Tivoli Storage Manager server, the access rules can change on the importing server. If an access rule is applied to all file spaces on the exporting server, the access rule on the importing server is restricted to only those file spaces that are imported. The file spaces are restricted in the access rule on the importing server for security reasons. Additionally, the access rules do not recognize the first occurrence of a wildcard character in the file specification when you restore or retrieve. This means that if you restore or retrieve with a wildcard character in the file specification, subdirectories are ignored.

Tip: If you might export a node to another Tivoli Storage Manager server, do not use a single wildcard character as the file specification in the access rule. Instead, create an access rule for each filespace.

Note: You cannot give access to both archives and backups using a single command.

When an existing file space is renamed during Unicode conversion, any access rules that are defined for the file space remain applicable to the original file space. However, new access rules must be defined to apply to the new Unicode file space.

Supported Clients

This command is valid for all clients.

Syntax

```
►►─SET Access─┬─ Archive ─┬─┬─ filespec ─┬─ node ─┬───────────┬─►◄
              └─ Backup ──┘ └─ image-fs ─┘         └─ user ─┘
```
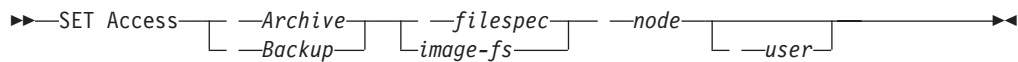
Parameters

Archive
> Permits access to archived files or images.

Backup
> Permits access to backup versions of files or images.

filespec
> Specifies the path, file, image, or directory to which your are giving access to another node or user. Use wildcard characters to specify a group of files or images, or all files in a directory; all objects in a directory branch; or all objects in a file system. Use a single asterisk "*" for the file spec to give access to all files or images owned by you and backed up on the server. When the command **set access backup "*" node** is entered, no check is made with the server; it is assumed you have at least one object backed up.
>
> If you give access to a branch of the current working directory, you only need to specify the branch. If you give access to objects that are not in a branch of the current working directory, you must specify the complete path. The file spec to which you gave access must have at least one backup version or archive copy object (file or directory) on the server.
>
> To specify all files in a named directory, enter /home/mine/proj1/* on the command line.
>
> To give access to all objects *below* a certain level, use an asterisk, directory delimiter, and an asterisk at the end of your file spec. For example, to give access to all objects below home/test, use file spec home/test/*/*.
>
> **Important:** Use of the form /*/* alone will not give access to objects in the named directory; only those in directories below the named directory are accessible.
>
> The rules are essentially the same when considering the root directory. Enter /* on one set access command and /*/* on another if you want another user to have access to all files and directories *in* and *below* the root directory. The first /* gives access to all directories and all files in the root directory. The second /* allows access to all directories and files below the root directory.
>
> For example:
> - Your directory structure is multilevel: /home/sub1/subsub1.
> - The /home directory contains the h1.txt and h2.txt files.
> - The /home/sub1 directory contains file s1.htm.
> - The /home/sub1/sub2 directory contains the ss1.cpp file.
>
> To allow access to all files in the /home/sub1/sub2 directory, enter:
> ```
> set access backup /home/sub1/sub2/* * *
> ```
>
> To allow access to only those files in the /home directory, enter:
> ```
> set access backup /home/* * *
> ```
>
> To allow access to all files in all directories *in* and *below* the /home directory, enter:

```
                  set access backup /home/* * *
                  set access backup /home/*/* * *
```

image-fs

> The name of the image file system to be shared. This can be specified as an asterisk (*) to allow access to all images owned by the user granting access.

node

> Specifies the client node of the user to whom you are giving access. Use wildcards to give access to more than one node with similar node names. Use an asterisk (*) to give access to all nodes.

user

> This is an optional parameter that restricts access to the named user at the specified node. This parameter is applicable for UNIX clients only. Windows clients can omit this parameter. If this parameter is specified on a Windows client, it is ignored.

Examples

Task   Give the user at `node_2` authority to restore the `budget` file from the `/home/user` directory.

> **Command:** `set access backup /home/user/budget node_2`

Task   Give `node_3` the authority to retrieve all files in the `/home/devel/proja` directory.

> **Command:** `set ac archive /home/devel/proja/ node_3`

Task   Give all nodes whose names end with `bldgb` the authority to restore all backup versions from directories with a file space name of `project`.

> **Command:** `set ac b "{project}/*" "*bldgb"`

Task   Give user serena at `node_5` authority to restore all images of the file space mounted on directory `/home/devel/proja`.

> **Command:** `set acc backup "home/devel/proja/*/*" node_5 serena`

Set Event

The **set event** command allows you to specify the circumstances for when archived data is deleted.

You can use the **set event** command to:
- Prevent the deletion of data at the end of its assigned retention period (Deletion hold)
- Allow the expiration to take place, as defined by the archive copy group (Release a deletion hold)
- Start the expiration clock running when a particular event occurs (Notify the server that an event has occurred)

Objects affected can be specified with a standard Tivoli Storage Manager filespec (including wildcards), a list of files whose names are in the file specified using the filelist option, or a group of archived files with the description specified with the description option.

Note: When only a <filespec> is used, *all* archived copies of files or folders matching the filespec are affected. If you only want to affect certain versions of a file, use the *-pick* option and select from the displayed list.

Interaction with down-level servers:

If the **set event** command is issued when the client is connected to a server that does not support event-based policy (previous to Tivoli Storage Manager 5.2.2), the command is rejected with an error message indicating that the current server does not support event-based policy.

Supported Clients

This command is valid for all clients.

Syntax

```
►►──SET Event──-TYPE=──┬──Hold──────────┬──  ─<filespec>────────────────────►
                       ├──Release───────┤
                       └──Activateretention─┘

►──  ──-filelist=<filespec>──  ──-description=──-pick──────────────────────►◄
```

Parameters

TYPE=

    Specifies the event type setting. This parameter must be specified.

hold

    Prevents the object from being deleted regardless of expiration policy.

release

    Allows normal event-controlled expiration to take place.

activateretention

    Signals the server that the controlling event has occurred and starts the expiration clock running.

-pick

    Provides a list of objects from which the user can select to apply the event.

The following options can also be used and serve their usual purpose:
- Dateformat
- Numberformat
- Noprompt
- Subdir
- Timeformat

Examples

Task   The following example displays the verbose and statistics output from the **set event** command `set event type=hold /home/accounting/ledgers/ *05.books`, with objects rebound (as opposed to archived or some other notation).

```
Rebinding--> 274 /home/accounting/ledgers/
  jan05.books
Rebinding--> 290 /home/accounting/ledgers/
  feb05.books
```

```
Total number of objects archived:        0
Total number of objects failed:          0
Total number of objects rebound:         2
Total number of bytes transferred:       0  B
Data transfer time:                   0.00 sec
Network data transfer rate:           0.00 KB/sec
Aggregate data transfer rate:         0.00 KB/sec
Objects compressed by:                   0%
Elapsed processing time:          00:00:02
```

Task The -pick option used with the set event command `set event type=activate /user/tsm521/common/unix` shows the event type instead of the command name:

```
TSM Scrollable PICK Window - Retention Event :  ACTIVATE

 #     Archive Date/Time      File Size   File
-------------------------------------------------------
 1. | 08/05/2003 08:47:46       766  B   /user/tsm521
                                          /common/unix
 2. | 08/01/2003 10:38:11       766  B   /user/tsm521
                                          /common/unix
 3. | 08/05/2003 08:47:46      5.79 KB   /user/tsm521
                                          /common/unix
 4. | 08/01/2003 10:38:11      5.79 KB   /user/tsm521
                                          /common/unix
 5. | 08/05/2003 08:47:46     10.18 KB   /user/tsm521
                                          /common/unix
```

Related information

"Dateformat" on page 289

"Numberformat" on page 376

"Noprompt" on page 376

"Subdir" on page 432

"Timeformat" on page 442

Set Password

The **set password** command changes the Tivoli Storage Manager password for your workstation.

If you omit the old and new passwords when you enter the **set password** command, you are prompted once for the old password and twice for the new password.

You must be an authorized user or system administrator to use this command.

A password is not case-sensitive, and it can be as many as 63 characters. Valid characters are:

a–z   Any letter, a through z, upper or lower-case

0–9   Any number, 0 through 9

+   Plus

.   Period

_	Underscore
-	Hyphen
&	Ampersand

Supported Clients

This command is valid for all clients.

Syntax

```
►►──SET Password──────────────────────────────────────────────────────►
                  └──oldpw── ──newpw──┘
```

```
►──────────────────────────────────────────────┬──TYPE=TSM──────┬──►◄
   └──otherserver── ──otheruserid── ──otherpassword──┘   ├──TYPE=FILER────┤
                                                         └──TYPE=FASTBack──┘
```
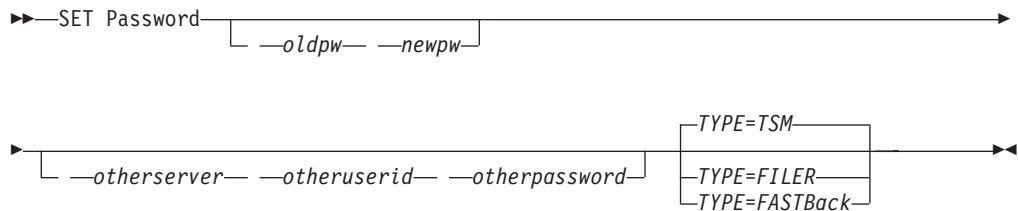
Parameters

oldpw
> Specifies the current password for your workstation.

newpw
> Specifies the new password for your workstation.

TYPE
> Specifies whether this password is for the Tivoli Storage Manager client or file server. Use *TYPE=TSM* to specify the Tivoli Storage Manager password for your client. *TYPE=FILER* (valid for AIX 64 bit and Linux x86/86_64 clients) to specify that this password is for snapshot difference operations on a file server. The default is *TSM*.
>
> For *TYPE=FILER* you must specify the file server name, the user ID, and the password. For example: dsmc set password -type=filer myfiler filerid filerpasswd.
>
> If you are using the **dsmc** command in batch mode, and your password contains any special characters (for example, pa$sword), you must enclose the password in single quotation marks to avoid truncation (for example, pa):
> ```
> dsmc
> set password -type=filer myfiler "pa$sword"
> ```
>
> Use the *TYPE=FastBack* parameter to store the Tivoli Storage Manager FastBack credentials that are required for mounting and dismounting the FastBack volumes on the Windows FastBack server or DR Hub server.
>
> The password file on the Linux proxy machine must have either the Windows administrator ID and password, or the FastBack administration ID and password. If the password file entry is for a FastBack DR Hub, specify the Windows administrator ID and password.
>
> Here are some examples:
> ```
> dsmc set password 9.47.98.169 admin admin 123 -type=fastback
> dsmc set password 9.47.98.169 WORKGROUP:admin admin 123 -type=fastback
> dsmc set password windserv administrator windpass4 -type=fastback
> ```

Important: You must define the user credentials that are required to mount and unmount FastBack volumes from a repository to the Tivoli Storage Manager client before you enter the Tivoli Storage Manager backup-archive FastBack subcommand. Use the fbserver option to define the credentials.

Here is a brief description of the various configurations and credentials that you need:

- The Tivoli Storage Manager backup-archive client is installed on a dedicated proxy workstation. The Tivoli Storage Manager dedicated backup-archive client proxy workstation must connect to multiple network share repositories.

Follow these steps for each of the network share repositories where the Tivoli Storage Manager client is connected:

1. Configure the repository for remote network access from FastBack Manager. Refer to the *Tivoli Storage Manager FastBack Installation and User's Guide* for details.

   This step establishes a domain name, a network share user ID, and a network share password to connect remotely to the repository.

2. On the Tivoli Storage Manager client workstation, manually enter the following command:

   ```
   dsmc set password type=fastback FBServer domain:networkaccessuserid
      networkaccesspassword
   ```

   The fbserver option specifies the short host name of the FastBack server workstation. For a FastBack DR Hub, the fbserver option specifies the short name of the workstation where the DR Hub is installed.

   Networkaccessuserid is either the Windows administrator ID or the FastBack administration password.

   Domain is the domain name of the user ID.

   Networkaccesspassword is either the Windows administrator ID or the FastBack administration password.

These credentials are retrieved based on the short host name that you specify with the fbserver option.

otherserver
   Specifies the file server host name if you specify *TYPE=FILER*. The file server host name must be the fully qualified DNS name and cannot be the IP address in dotted-decimal format or the short local-host name. This parameter only applies if you specify *TYPE=FILER*.

otheruserid
   Specifies the user ID to be used for the entity named by the *otherserver* parameter. This parameter only applies if you specify *TYPE=FILER*.

otherpassword
   Specifies the password to be used for the entity named by the *otherserver* parameter. This parameter only applies if you specify *TYPE=FILER*.

When you specify *TYPE=FILER* (valid for AIX 64 bit and Linux x86/86_64 clients), the password is stored in the `TSM.PWD` file without validating that the password is valid. Passwords that are stored with *TYPE=FILER* can be shared between Tivoli Storage Manager nodes. For example, a password stored by NODEA can be used by NODEB. Only one set of credentials is stored per file server.

Examples

The following examples use the **set password** command.

Task Change your password from osecret to nsecret.

Command: `set password osecret nsecret`

Task Set up a user ID and password for the root user on the file server
myFiler.example.com.

Command: `dsmc set password -type=filer myFiler.example.com root`

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

Task Set up a user ID and password for the root user on the file server
myFiler.example.com.

Command: `dsmc set password -type=filer myFiler.example.com root
secret`

Task Set up a user ID and password for the FastBack server myFastBackServer.
Use the -fbserver option in the **archive fastback** and **backup fastback**
commands for the server name.

Command: `dsmc set password -type=FASTBack myFastBackServer
myUserId "pa$sword"`

Important:
1. The `dsmc set password -type=fastback` command must be repeated on
   a dedicated client proxy workstation once for each FastBack repository
   where the Tivoli Storage Manager backup-archive client is expected to
   connect.
2. For network share repositories the `dsmc set password -type=fastback`
   command should be in this format: `dsmc set password -type=fastback
   myFBServer domainName:userId password`.
3. For a FastBack server or a Tivoli Storage Manager FastBack Hub, the
   user ID and password specified must have Tivoli Storage Manager
   FastBack administrator privileges. You must issue the `dsmc set
   password -type=fastback` command once for each FastBack Server
   branch repository on the FastBack DR Hub that the Tivoli Storage
   Manager backup-archive client is expected to connect to.
4. The server name specified, `myFBServer`, in the above example, must
   match the name to be specified on the -fbserver option specified on the
   backup fastback or archive fastback command.

Task The Tivoli Storage Manager backup-archive client is connecting to a
FastBack server repository whose short hostname is myFBServer. user ID is
the FastBack network user ID that has read/write access to the repository
share. DOMAIN is the domain to which the user ID belongs. myNetworkPass is
the corresponding password for the user ID.

Command: `dsmc set password -type=fastback myFbServer DOMAIN:USERID
myNetworkPass`

Task The Tivoli Storage Manager client is connecting to a repository on a DR
Hub machine whose short hostname is myFbDrHub. The user ID is the

Windows administrator ID. `DOMAIN` is the domain to which the DR Hub machine belongs. `myNetworkPass` is the corresponding password for the administrator ID.

Command: `dsmc set password -type=fastback myFbDrHub DOMAIN:administrator adminPasswd`

Related reference

"Snapdiff" on page 418

Appendix. Accessibility features for Tivoli Storage Manager

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

Accessibility features

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

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices
- User documentation provided in HTML and PDF format. Descriptive text is provided for all documentation images.

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

Keyboard navigation

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

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

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

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

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

Vendor software

Tivoli Storage Manager includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

Related accessibility information

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

IBM and accessibility

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

Notices

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

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

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

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

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

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

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

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

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs. Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Trademarks

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

 Java and all Java-based trademarks are trademarks are trademarks or registered trademarks of Oracle or its affiliates.

Linux is a registered trademark of Linus Torvalds 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.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Manager.

To view glossaries for other IBM products, go to http://www.ibm.com/software/globalization/terminology/.

The following cross-references are used in this glossary:

- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

A

absolute mode
In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup even if the file has not changed since the last backup. See also *modified mode*.

access control list (ACL)
In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights. For example, an access control list is associated with a file that identifies the users who can access that file and their access rights.

access mode
An attribute of a storage pool or a storage volume that specifies whether the server can write to or read from the storage pool or storage volume. The access mode can be read/write, read-only, or unavailable. Volumes in primary storage pools can also have an access mode of destroyed. Volumes in copy storage pools can also have an access mode of offsite.

acknowledgment
The transmission of acknowledgment characters as a positive response to a data transmission.

ACL　See *access control list*.

activate
To validate the contents of a policy set and then make it the active policy set.

active-data pool
A named set of storage pool volumes that contain only active versions of client backup data.

active file system
A file system to which space management has been added. With space management, tasks for an active file system include automatic migration, reconciliation, selective migration, and recall. Contrast with *inactive file system*.

active policy set
The activated policy set that contains the policy rules in use by all client nodes that are assigned to the policy domain. See also *policy domain* and *policy set*.

active version
The most recent backup copy of a file stored. The active version of a file cannot be deleted until a backup process detects that the user has either replaced the file with a newer version or has deleted the file from the file server or workstation. Contrast with *inactive version*.

activity log
A log that records normal activity messages that are generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors.

adaptive subfile backup
A type of backup that sends only changed portions of a file to the server, instead of sending the entire file. Adaptive subfile backup reduces network traffic and increases the speed of the backup.

administrative client
A program that runs on a file server, workstation, or mainframe that administrators use to control and monitor the Tivoli Storage Manager server. Contrast with *backup-archive client*.

administrative command schedule
A database record that describes the planned processing of an administrative

　　　567

command during a specific time period.
See also *client schedule*.

administrative privilege class
See *privilege class*.

administrative session
A period of time during which an
administrator user ID communicates with
a server to perform administrative tasks.
Contrast with *client node session*.

administrator
A user who is registered to the server as
an administrator, and who is authorized
to perform tasks and issue commands
through the assignment of an
administrative privilege class.

**Advanced Program-to-Program Communication
(APPC)**
An implementation of the SNA LU 6.2
protocol that allows interconnected
systems to communicate and share the
processing of programs.

agent node
A client node that has been granted proxy
authority to perform operations on behalf
of another client node, which is the target
node.

aggregate
An object, stored in one or more storage
pools, consisting of a group of logical files
that are packaged together. See also *logical
file* and *physical file*.

aggregate data transfer rate
A performance statistic that indicates the
average number of bytes that were
transferred per second while processing a
given operation.

APPC See *Advanced Program-to-Program
Communication*.

application client
A program that is installed on a system to
protect an application. The Tivoli Storage
Manager server provides backup services
to an application client.

archive
To copy programs, data, or files to
another storage media, usually for
long-term storage or security. Contrast
with *retrieve*.

archive copy
A file or group of files that was archived
to server storage.

archive copy group
A policy object containing attributes that
control the generation, destination, and
expiration of archived files.

archive-retention grace period
The number of days that the storage
manager retains an archived file when the
server is unable to rebind the file to an
appropriate management class. See also
bind.

association
(1) The defined relationship between a
client node and a client schedule. An
association identifies the name of a
schedule, the name of the policy domain
to which the schedule belongs, and the
name of a client node that performs
scheduled operations.

(2) On a configuration manager, the
defined relationship between a profile and
an object such as a policy domain. Profile
associations define the configuration
information that is distributed to a
managed server when it subscribes to the
profile.

audit To check for logical inconsistencies
between information that the server has
and the actual condition of the system.
The storage manager can audit
information about items such as volumes,
libraries, and licenses. For example, when
a storage manager audits a volume, the
server checks for inconsistencies between
information about backed-up or archived
files that are stored in the database and
the actual data that are associated with
each backup version or archive copy in
server storage.

authentication
The process of checking a user's password
before permitting user access to the Tivoli
Storage Manager server. Authentication
can be turned on or off by an
administrator with system privilege.

authentication rule
A specification that another user can use
to either restore or retrieve files from
storage.

authority

The right to access objects, resources, or functions. See also *privilege class*.

authorization rule

A specification that permits another user to either restore or retrieve a user's files from storage.

authorized user

A user who has administrative authority for the Tivoli Storage Manager client on a workstation. This user changes passwords, performs open registrations, and deletes file spaces.

AutoFS

See *automounted file system*.

automatic detection

A feature that detects, reports, and updates the serial number of a drive or library in the database when the path from the local server is defined.

automatic migration

The process that is used to automatically move files from a local file system to storage, based on options and settings that are chosen by a root user on a workstation. See also *threshold migration* and *demand migration*.

automatic reconciliation

The process that is used to reconcile file systems at regular intervals. The intervals are set by a user with root user authority. See also *reconciliation*.

automounted file system (AutoFS)

A file system that is managed by an automounter daemon. The automounter daemon monitors a specified directory path, and automatically mounts the file system to access data.

B

backup-archive client

A program that runs on a workstation or file server and provides a means for users to back up, archive, restore, and retrieve files. Contrast with *administrative client*.

backup copy group

A policy object containing attributes that control the generation, destination, and expiration of backup versions of files. A backup copy group belongs to a management class.

backup-retention grace period

The number of days the storage manager retains a backup version after the server is unable to rebind the file to an appropriate management class.

backup set

A portable, consolidated group of active versions of backup files that are generated for a backup-archive client.

backup set collection

A group of backup sets that are created at the same time and which have the same backup set name, volume names, description, and device classes. The server identifies each backup set in the collection by its node name, backup set name, and file type.

backup version

A file or directory that a client node backed up to server storage. More than one backup version can exist in server storage, but only one backup version is the active version. See also *active version* and *inactive version*.

bind   To associate a file with a management class name. See *rebind*.

bindery

A database that consists of three system files for a NetWare server. The files contain user IDs and user restrictions.

C

cache   To place a duplicate copy of a file on random access media when the server migrates a file to another storage pool in the hierarchy.

cache file

A snapshot of a logical volume created by Logical Volume Snapshot Agent. Blocks are saved immediately before they are modified during the image backup and their logical extents are saved in the cache files.

CAD   See *client acceptor*.

central scheduler

A function that permits an administrator to schedule client operations and administrative commands. The operations can be scheduled to occur periodically or on a specific date. See *client schedule* and *administrative command schedule*.

client A software program or computer that requests services from a server.

client acceptor
An HTTP service that serves the Java applet for the Web client to Web browsers. On Windows systems, the client acceptor is installed and run as a service. On AIX, UNIX, and Linux systems, the client acceptor is run as a daemon, and is also called the *client acceptor daemon* (CAD).

client acceptor daemon (CAD)
See *client acceptor*.

client domain
The set of drives, file systems, or volumes that the user selects to back up or archive data, using the backup-archive client.

client node
A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

client node session
A session in which a client node communicates with a server to perform backup, restore, archive, retrieve, migrate, or recall requests. Contrast with *administrative session*.

client options file
An editable file that identifies the server and communication method, and provides the configuration for backup, archive, hierarchical storage management, and scheduling.

client option set
A group of options that are defined on the server and used on client nodes in conjunction with client options files.

client-polling scheduling mode
A method of operation in which the client queries the server for work. Contrast with *server-prompted scheduling mode*.

client schedule
A database record that describes the planned processing of a client operation during a specific time period. The client operation can be a backup, archive, restore, or retrieve operation, a client operating system command, or a macro. See also *administrative command schedule*.

client/server
Pertaining to the model of interaction in distributed data processing in which a program on one computer sends a request to a program on another computer and awaits a response. The requesting program is called a client; the answering program is called a server.

client system-options file
A file, used on AIX, UNIX, or Linux system clients, containing a set of processing options that identify the servers to be contacted for services. This file also specifies communication methods and options for backup, archive, hierarchical storage management, and scheduling. This file is also called the dsm.sys file. See also *client user-options file*.

client user-options file
A file that contains the set of processing options that the clients on the system use. The set can include options that determine the server that the client contacts, and options that affect backup operations, archive operations, hierarchical storage management operations, and scheduled operations. This file is also called the dsm.opt file. For AIX, UNIX, or Linux systems, see also *client system-options file*.

closed registration
A registration process in which only an administrator can register workstations as client nodes with the server. Contrast with *open registration*.

collocation
The process of keeping all data belonging to a single-client file space, a single client node, or a group of client nodes on a minimal number of sequential-access volumes within a storage pool. Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored.

collocation group
A user-defined group of client nodes whose data is stored on a minimal number of volumes through the process of collocation.

commit point
A point in time when data is considered consistent.

Common Programming Interface for Communications (CPI-C)
A call-level interface that provides a consistent application programming interface (API) for applications that use program-to-program communications. CPI-C uses LU 6.2 architecture to create a set of interprogram services that can establish and end a conversation, send and receive data, exchange control information, and notify a partner program of errors.

communication method
The method by which a client and server exchange information. See also *Transmission Control Protocol/Internet Protocol*.

communication protocol
A set of defined interfaces that permit computers to communicate with each other.

compression
A function that removes repetitive characters, spaces, or strings of characters from the data being processed and replaces the repetitive characters with control characters. Compression reduces the amount of storage space that is required for the data.

configuration manager
A server that distributes configuration information, such as policies and schedules, to managed servers according to their profiles. Configuration information can include policy and schedules. See also *managed server* and *profile*.

conversation
A connection between two programs over a session that allows them to communicate with each other while processing a transaction.

copy backup
A full backup in which the transaction log files are not deleted so that backup procedures that use incremental or differential backups are not disrupted

copy group
A policy object containing attributes that control how backup versions or archive copies are generated, where backup versions or archive copies are initially located, and when backup versions or archive copies expire. A copy group belongs to a management class. See also *archive copy group*, *backup copy group*, *backup version*, and *management class*.

copy storage pool
A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files). See also *primary storage pool* and *destination*.

CPI-C See *Common Programming Interface for Communications*.

D

daemon
A program that runs unattended to perform continuous or periodic functions, such as network control.

damaged file
A physical file in which Tivoli Storage Manager has detected read errors.

data access control mode
A mode that controls whether a command can access a migrated file, see a migrated file as zero-length, or receive an input/output error if it attempts to access a migrated file. See also *execution mode*.

database backup series
One full backup of the database, plus up to 32 incremental backups made since that full backup. Each full backup that is run starts a new database backup series. A number identifies each backup series.

database snapshot
A complete backup of the entire database to media that can be taken off-site. When a database snapshot is created, the current database backup series is not interrupted. A database snapshot cannot have incremental database backups associated with it. See also *database backup series*. Contrast with *full backup*.

data deduplication
A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage

media. Other instances of the same data are replaced with a pointer to the retained instance.

data manager server
A server that collects metadata information for client inventory and manages transactions for the storage agent over the local area network. The data manager server informs the storage agent with applicable library attributes and the target volume identifier.

data mover
A device that moves data on behalf of the server. A network-attached storage (NAS) file server is a data mover.

data storage-management application-programming interface (DSMAPI)
A set of functions and semantics that can monitor events on files, and manage and maintain the data in a file. In an HSM environment, a DSMAPI uses events to notify data management applications about operations on files, stores arbitrary attribute information with a file, supports managed regions in a file, and uses DSMAPI access rights to control access to a file object.

default management class
A management class that is assigned to a policy set. This class is used to govern backed up or archived files when a file is not explicitly associated with a specific management class through the include-exclude list.

deduplication
See *data deduplication.*

demand migration
The process that is used to respond to an out-of-space condition on a file system for which hierarchical storage management (HSM) is active. Files are migrated to server storage until space usage drops to the low threshold that was set for the file system. If the high threshold and low threshold are the same, one file is migrated.

desktop client
The group of backup-archive clients that includes clients on Microsoft Windows, Apple, and Novell NetWare operating systems.

destination
A copy group or management class attribute that specifies the primary storage pool to which a client file will be backed up, archived, or migrated.

device class
A named set of characteristics that are applied to a group of storage devices. Each device class has a unique name and represents a device type of disk, file, optical disk, or tape.

device configuration file
(1) For a server, a file that contains information about defined device classes, and, on some servers, defined libraries and drives. The information is a copy of the device configuration information in the database.

(2) For a storage agent, a file that contains the name and password of the storage agent, and information about the server that is managing the SAN-attached libraries and drives that the storage agent uses.

device driver
A program that provides an interface between a specific device and the application program that uses the device.

disaster recovery manager (DRM)
A function that assists in preparing and using a disaster recovery plan file for the server.

disaster recovery plan
A file that is created by the disaster recovery manager (DRM) that contains information about how to recover computer systems if a disaster occurs and scripts that can be run to perform some recovery tasks. The file includes information about the software and hardware that is used by the server, and the location of recovery media.

domain
A grouping of client nodes with one or more policy sets, which manage data or storage resources for the client nodes. See *policy domain* or *client domain*.

DRM See *disaster recovery manager*.

DSMAPI
See *data storage-management application-programming interface*.

dynamic serialization
A type of copy serialization in which a file or folder is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.

E

EA See *extended attribute*.

EB See *exabyte*.

EFS See *Encrypted File System*.

Encrypted File System (EFS)
A file system that uses file system-level encryption.

enterprise configuration
A method of setting up servers so that the administrator can distribute the configuration of one of the servers to the other servers, using server-to-server communication. See also *configuration manager*, *managed server*, *profile*, and *subscription*.

enterprise logging
The process of sending events from a Tivoli Storage Manager server to a designated event server. The event server routes the events to designated receivers, such as to a user exit. See also *event*.

error log
A data set or file that is used to record error information about a product or system.

estimated capacity
The available space, in megabytes, of a storage pool.

event (1) An administrative command or a client operation that is scheduled to be run using Tivoli Storage Manager scheduling.

(2) A message that an Tivoli Storage Manager server or client issues. Messages can be logged using Tivoli Storage Manager event logging.

event record
A database record that describes actual status and results for events.

event server
A server to which other servers can send events for logging. The event server routes the events to any receivers that are enabled for the sending server's events.

exabyte (EB)
For processor storage, real and virtual storage, and channel volume, 1 152 921 504 606 846 976 bytes. For disk storage capacity and communications volume, 1 000 000 000 000 000 000 bytes.

exclude
The process of identifying files in an include-exclude list. This process prevents the files from being backed up or migrated whenever a user or schedule enters an incremental or selective backup operation. A file can be excluded from backup and space management, backup only, or space management only.

exclude-include list
See *include-exclude list*.

execution mode
A mode that controls the space-management related behavior of commands that run under the **dsmmode** command.

expiration
The process by which files, data sets, or objects are identified for deletion because their expiration date or retention period has passed.

expiring file
A migrated or premigrated file that has been marked for expiration and removal from storage. If a stub file or an original copy of a premigrated file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run.

extend
To increase the portion of available space that can be used to store database or recovery log information.

extended attribute (EA)
Names or value pairs that are associated with files or directories. There are three classes of extended attributes: user attributes, system attributes, and trusted attributes.

extent The part of a file that is created during the data-deduplication process. Extents are compared with other file extents to identify duplicates.

external library
A type of library that is provided by Tivoli Storage Manager that permits LAN-free data movement for StorageTek libraries that are managed by Automated Cartridge System Library Software (ACSLS). To activate this function, the Tivoli Storage Manager library type must be EXTERNAL.

F

file access time
On AIX, UNIX, or Linux systems, the time when the file was last accessed.

file age
For migration prioritization purposes, the number of days since a file was last accessed.

file device type
A device type that specifies the use of sequential access files on disk storage as volumes.

file server
A dedicated computer and its peripheral storage devices that are connected to a local area network that stores programs and files that are shared by users on the network.

file space
A logical space in server storage that contains a group of files that have been backed up or archived by a client node, from a single logical partition, file system, or virtual mount point. Client nodes can restore, retrieve, or delete their file spaces from server storage. In server storage, files belonging to a single file space are not necessarily stored together.

file space ID (FSID)
A unique numeric identifier that the server assigns to a file space when it is stored in server storage.

file state
The space management mode of a file that resides in a file system to which space management has been added. A file can be in one of three states: resident, premigrated, or migrated. See also *resident file*, *premigrated file*, and *migrated file*.

file system migrator (FSM)
A kernel extension that intercepts all file system operations and provides any space management support that is required. If no space management support is required, the operation is passed to the operating system, which performs its normal functions. The file system migrator is mounted over a file system when space management is added to the file system.

file system state
The storage management mode of a file system that resides on a workstation on which the hierarchical storage management (HSM) client is installed. A file system can be in one of these states: native, active, inactive, or global inactive.

frequency
A copy group attribute that specifies the minimum interval, in days, between incremental backups.

FSID See *file space ID*.

FSM See *file system migrator*.

full backup
The process of backing up the entire server database. A full backup begins a new database backup series. See also *database backup series* and *incremental backup*. Contrast with *database snapshot*.

fuzzy backup
A backup version of a file that might not accurately reflect what is currently in the file because the file was backed up at the same time as it was being modified.

fuzzy copy
A backup version or archive copy of a file that might not accurately reflect the original contents of the file because it was backed up or archived the file while the file was being modified. See also *backup version* and *archive copy*.

G

General Parallel File System
A high-performance shared-disk file system that can provide data access from nodes in a cluster environment.

gigabyte (GB)
In decimal notation, 1 073 741 824 when referring to memory capacity; in all other cases, it is defined as 1 000 000 000.

global inactive state
The state of all file systems to which

space management has been added when space management is globally deactivated for a client node. When space management is globally deactivated, hierarchical storage management (HSM) cannot perform migration, recall, or reconciliation. However, a root user can update space management settings and add space management to additional file systems. Users can access resident and premigrated files.

Globally Unique Identifier (GUID)
An algorithmically determined number that uniquely identifies an entity within a system.

GPFS See *General Parallel File System*.

GPFS node set
A mounted, defined group of GPFS file systems.

group backup
The backup of a group containing a list of files from one or more file space origins.

GUID See *Globally Unique Identifier*.

H

hierarchical storage management (HSM)
A function that automatically distributes and manages data on disk, tape, or both by regarding devices of these types and potentially others as levels in a storage hierarchy that range from fast, expensive devices to slower, cheaper, and possibly removable devices. The objectives are to minimize access time to data and maximize available media capacity.

hierarchical storage management (HSM) client
A client program that works with the Tivoli Storage Manager server to provide hierarchical storage management (HSM) for a system. See also *hierarchical storage management* and *space manager client*.

HSM See *hierarchical storage management*.

HSM client
See *hierarchical storage management client*.

I

ILM See *information lifecycle management*.

image A file system or raw logical volume that is backed up as a single object.

image backup
A backup of a full file system or raw logical volume as a single object.

inactive file system
A file system for which space management has been deactivated. Contrast with *active file system*.

inactive version
A backup version of a file that is either not the most recent backup version, or that is a backup version of a file that no longer exists on the client system. Inactive backup versions are eligible for expiration processing according to the management class assigned to the file. Contrast with *active version*.

include-exclude file
A file containing statements to determine the files to back up and the associated management classes to use for backup or archive. See also *include-exclude list*.

include-exclude list
A list of options that include or exclude selected files for backup. An exclude option identifies files that should not be backed up. An include option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services.

incremental backup
(1) A copy of all database data that has changed since the most recent successful full backup operation. An incremental backup is also known as a *cumulative backup image* because each incremental backup includes the contents of the previous incremental backup.

(2) The process of backing up information in the database that is new or changed since the last full backup. Contrast with *full backup*. See also *database backup series*.

(3) For Data Protection for Microsoft Exchange Server, a backup in which the transaction logs are backed up and then cleared.

individual mailbox restore
See *mailbox restore*.

information lifecycle management (ILM)
GPFS policy-based file management for storage pools and file sets.

i-node The internal structure that describes the individual files on AIX, UNIX, or Linux systems. An i-node contains the node, type, owner, and location of a file.

i-node number
A number specifying a particular i-node file in the file system.

IP address
A unique address for a device or logical unit on a network that uses the IP standard.

J

job file
A generated file that contains configuration information for a migration job. The file is XML format and can be created and edited in the hierarchical storage management (HSM) client for Windows client graphical user interface.

journal-based backup
A method for backing up Windows clients and AIX clients that exploits the change notification mechanism in a file to improve incremental backup performance by reducing the need to fully scan the file system.

journal daemon
On AIX, UNIX, or Linux systems, a program that tracks change activity for files residing in file systems.

journal service
In Microsoft Windows, a program that tracks change activity for files residing in file systems.

K

kilobyte (KB)
For processor storage, real and virtual storage, and channel volume, 210 or 1 024 bytes. For disk storage capacity and communications volume, 1 000 bytes.

L

LAN See *local area network*.

LAN-free data movement
The movement of client data between a client system and a storage device on a storage area network (SAN), bypassing the local area network. This process is also referred to as *LAN-free data transfer*.

LAN-free data transfer
See *LAN-free data movement*.

leader data
Bytes of data, from the beginning of a migrated file, that are stored in the file's corresponding stub file on the local file system. The amount of leader data that is stored in a stub file depends on the stub size that is specified.

library
(1) A repository for demountable recorded media, such as magnetic disks and magnetic tapes.

(2) A collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

library client
A server that uses server-to-server communication to access a library that is managed by another storage management server. See also *library manager*.

library manager
A server that controls device operations when multiple storage management servers share a storage device. See also *library client*.

local Pertaining to a device, file, or system that is accessed directly from a user's system, without the use of a communication line.

local area network (LAN)
A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

local shadow volumes
Data that is stored on shadow volumes localized to a disk storage subsystem.

LOFS See *loopback virtual file system*.

logical file
A file that is stored in one or more server storage pools, either by itself or as part of an aggregate. See also *aggregate* and *physical file*.

logical occupancy
The space that is used by logical files in a storage pool. This space does not include the unused space created when logical

files are deleted from aggregate files, so it might be less than the physical occupancy.

logical unit (LU)
An access point through which a user or application program accesses the Systems Network Architecture (SNA) network to communicate with another user or application program.

logical unit number (LUN)
In the Small Computer System Interface (SCSI) standard, a unique identifier that is used to differentiate devices, each of which is a logical unit (LU).

logical volume
A portion of a physical volume that contains a file system.

logical volume backup
A back up of a file system or logical volume as a single object.

Logical Volume Snapshot Agent (LVSA)
Software that can act as the snapshot provider for creating a snapshot of a logical volume during an online image backup.

loopback virtual file system (LOFS)
A file system that is created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

LU   See *logical unit*.

LUN   See *logical unit number*.

LVSA See *Logical Volume Snapshot Agent*.

M

macro file
A file that contains one or more storage manager administrative commands, which can be run only from an administrative client using the MACRO command. Contrast with *Tivoli Storage Manager command script*.

mailbox restore
A function that restores Microsoft Exchange Server data (from IBM Data Protection for Exchange backups) at the mailbox level or mailbox-item level.

managed object
In Tivoli Storage Manager, a definition in the database of a managed server that was distributed to the managed server by a configuration manager. When a managed server subscribes to a profile, all objects that are associated with that profile become managed objects in the database of the managed server. In general, a managed object cannot be modified locally on the managed server. Objects can include policy, schedules, client option sets, server scripts, administrator registrations, and server and server group definitions.

managed server
A Tivoli Storage Manager server that receives configuration information from a configuration manager using a subscription to one or more profiles. Configuration information can include definitions of objects such as policy and schedules. See also *configuration manager*, *subscription*, and *profile*.

management class
A policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. See also *copy group*, *space manager client*, *bind*, and *rebind*.

maximum transmission unit
The largest possible unit of data that can be sent on a given physical medium in a single frame. For example, the maximum transmission unit for Ethernet is 1500 bytes.

MB   See *megabyte*.

megabyte (MB)
(1) 1 048 576 bytes (two to the twentieth power) when used in this publication.

(2) For processor storage, real and virtual storage, and channel volume, 2 to the power of 20 or 1 048 576 bits. For disk storage capacity and communications volume, 1 000 000 bits.

metadata
Data that describes the characteristics of data; descriptive data.

migrate
To move data from one storage location to another. In Tivoli Storage Manager products, migrating can mean moving

data from a client node to server storage, or moving data from one storage pool to the next storage pool defined in the server storage hierarchy. In both cases the movement is controlled by policy, such as thresholds that are set. See also *migration threshold*.

migrated file
A file that has been copied from a local file system to Tivoli Storage Manager storage. For HSM clients on UNIX or Linux systems, the file is replaced with a stub file on the local file system. On Windows systems, creation of the stub file is optional. See also *stub file* and *resident file*. For HSM clients on UNIX or Linux systems, contrast with *premigrated file*.

migrate-on-close recall mode
A mode that causes a migrated file to be recalled back to its originating file system temporarily. Contrast with *normal recall mode* and *read-without-recall recall mode*.

migration job
A specification of files to migrate, and actions to perform on the original files after migration. See also *job file*.

migration threshold
High and low capacities for storage pools or file systems, expressed as percentages, at which migration is set to start and stop.

mirroring
The process of writing the same data to multiple locations at the same time. Mirroring data protects against data loss within the recovery log.

mode A copy group attribute that specifies whether to back up a file that has not been modified since the last time the file was backed up. See *modified mode* and *absolute mode*.

modified mode
In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup only if it has changed since the last backup. A file is considered a changed file if the date, size, owner, or permissions of the file have changed. See also *absolute mode*.

mount limit
The maximum number of volumes that can be simultaneously accessed from the

same device class. The mount limit determines the maximum number of mount points. See also *mount point*.

mount point
On the Tivoli Storage Manager server, a logical drive through which volumes in a sequential access device class are accessed. For removable-media device types, such as tape, a mount point is a logical drive that is associated with a physical drive. For the file device type, a mount point is a logical drive that is associated with an I/O stream. The number of mount points for a device class is defined by the value of the mount limit attribute for that device class. See also *mount limit*.

mount retention period
The maximum number of minutes that the server retains a mounted sequential-access media volume that is not being used before it dismounts the sequential-access media volume.

mount wait period
The maximum number of minutes that the server waits for a sequential-access volume mount request to be satisfied before canceling the request.

MTU See *maximum transmission unit*.

N

Nagle algorithm
An algorithm that reduces congestion of TCP/IP networks by combining smaller packets and sending them together.

named pipe
A type of interprocess communication that permits message data streams to pass between peer processes, such as between a client and a server.

NAS See *network-attached storage*.

NAS node
A client node that is a network-attached storage (NAS) file server. Data for the NAS node is transferred by a NAS file server that is controlled by the network data management protocol (NDMP). A NAS node is also called a NAS file server node.

native file system
A file system that is locally added to the file server and is not added for space

management. The hierarchical storage manager (HSM) client does not provide space management services to the file system.

native format
A format of data that is written to a storage pool directly by the Tivoli Storage Manager server. Contrast with *non-native data format*.

NDMP
See *Network Data Management Protocol*.

NetBIOS
See *Network Basic Input/Output System*.

network-attached storage (NAS) file server
A dedicated storage device with an operating system that is optimized for file-serving functions. A NAS file server can have the characteristics of both a node and a data mover.

Network Basic Input/Output System (NetBIOS)
A standard interface to networks and personal computers that is used on local area networks to provide message, print-server, and file-server functions. Application programs that use NetBIOS do not have to handle the details of LAN data link control (DLC) protocols.

Network Data Management Protocol (NDMP)
A protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server, without installing vendor-acquired software on that file server.

network data-transfer rate
A rate that is calculated by dividing the total number of bytes that are transferred by the data transfer time. For example, this rate can be the time that is spent transferring data over a network.

node A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

node name
A unique name that is used to identify a workstation, file server, or PC to the server.

node privilege class
A privilege class that gives an administrator the authority to remotely access backup-archive clients for a specific client node or for all clients in a policy domain. See also *privilege class*.

non-native data format
A format of data that is written to a storage pool that differs from the format that the server uses for operations.

normal recall mode
A mode that causes a migrated file to be copied back to its originating file system when it is accessed.

O

offline volume backup
A backup in which the volume is locked so that no other system applications can access it during the backup operation.

online volume backup
A backup in which the volume is available to other system applications during the backup operation.

open registration
A registration process in which users can register their workstations as client nodes with the server. Contrast with *closed registration*.

operator privilege class
A privilege class that gives an administrator the authority to disable or halt the server, enable the server, cancel server processes, and manage removable media. See also *privilege class*.

options file
A file that contains processing options. On Windows and NetWare systems, the file is called dsm.opt. On AIX, UNIX, Linux, and Mac OS X systems, the file is called dsm.sys.

originating file system
The file system from which a file was migrated. When a file is recalled using normal or migrate-on-close recall mode, it is always returned to its originating file system.

orphaned stub file
A file for which no migrated file can be found on the Tivoli Storage Manager server that the client node is contacting for space management services. For example, a stub file can be orphaned when the client system-options file is

modified to contact a server that is different than the one to which the file was migrated.

out-of-space protection mode
A mode that controls whether the program intercepts out-of-space conditions. See also *execution mode*.

P

pacing
In SNA, a technique by which the receiving system controls the rate of transmission of the sending system to prevent overrun.

packet In data communication, a sequence of binary digits, including data and control signals, that is transmitted and switched as a composite whole.

page A defined unit of space on a storage medium or within a database volume.

partial-file recall mode
A recall mode that causes the hierarchical storage management (HSM) function to read just a portion of a migrated file from storage, as requested by the application accessing the file.

password generation
A process that creates and stores a new password in an encrypted password file when the old password expires. Automatic generation of a password prevents password prompting. Password generation can be set in the options file (passwordaccess option). See also *options file*.

path An object that defines a one-to-one relationship between a source and a destination. Using the path, the source accesses the destination. Data can flow from the source to the destination, and back. An example of a source is a data mover (such as a network-attached storage [NAS] file server), and an example of a destination is a tape drive.

pattern-matching character
See *wildcard character*.

physical file
A file that is stored in one or more storage pools, consisting of either a single logical file, or a group of logical files that are packaged together as an aggregate. See also *aggregate* and *logical file*.

physical occupancy
The amount of space that is used by physical files in a storage pool. This space includes the unused space that is created when logical files are deleted from aggregates. See also *physical file*, *logical file*, and *logical occupancy*.

plug-in
A self-contained software component that modifies (adds, or changes) the function in a particular system. When a plug-in is added to a system, the foundation of the original system remains intact.

policy domain
A grouping of policy users with one or more policy sets, which manage data or storage resources for the users. The users are client nodes that are associated with the policy domain.

policy privilege class
A privilege class that gives an administrator the authority to manage policy objects, register client nodes, and schedule client operations for client nodes. Authority can be restricted to certain policy domains. See also *privilege class*.

policy set
A group of rules in a policy domain. The rules specify how data or storage resources are automatically managed for client nodes in the policy domain. Rules can be contained in management classes. See also *active policy set* and *management class*.

premigrated file
A file that has been copied to Tivoli Storage Manager storage, but has not been replaced with a stub file on the local file system. An identical copy of the file resides both on the local file system and in Tivoli Storage Manager storage. Premigrated files occur on UNIX and Linux file systems to which space management has been added. Contrast with *migrated file* and *resident file*.

premigrated files database
A database that contains information about each file that has been premigrated to Tivoli Storage Manager storage. The database is stored in a hidden directory

named `.SpaceMan` in each file system to which space management has been added.

premigration
The process of copying files that are eligible for migration to Tivoli Storage Manager storage, but leaving the original file intact on the local file system.

premigration percentage
A space management setting that controls whether the next eligible candidates in a file system are premigrated following threshold or demand migration.

primary storage pool
A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files migrated from client nodes. See also *destination* and *copy storage pool*.

privilege class
A level of authority that is granted to an administrator. The privilege class determines which administrative tasks the administrator can perform. See also *node privilege class*, *operator privilege class*, *policy privilege class*, *storage privilege class*, and *system privilege class*.

profile
A named group of configuration information that can be distributed from a configuration manager when a managed server subscribes. Configuration information can include registered administrator IDs, policies, client schedules, client option sets, administrative schedules, storage manager command scripts, server definitions, and server group definitions. See also *configuration manager* and *managed server*.

Q

quota
(1) For HSM on AIX, UNIX, or Linux systems, the limit (in megabytes) on the amount of data that can be migrated and premigrated from a file system to server storage.

(2) For HSM on Windows systems, a user-defined limit to the space that is occupied by recalled files.

R

randomization
The process of distributing schedule start times for different clients within a specified percentage of the schedule's startup window.

raw logical volume
A portion of a physical volume that is comprised of unallocated blocks and has no journaled file system (JFS) definition. A logical volume is read/write accessible only through low-level I/O functions.

read-without-recall recall mode
A mode that causes hierarchical storage management (HSM) to read a migrated file from storage without storing it back on the local file system. The last piece of information read from the file is stored in a buffer in memory on the local file system. Contrast with *normal recall mode* and *migrate-on-close recall mode*.

rebind
To associate a backed-up file with a new management class name. For example, rebinding occurs when the management class associated with a file is deleted. See also *bind*.

recall
In Tivoli Storage Manager, to copy a migrated file from server storage back to its originating file system using the space management client. See also *transparent recall*, *selective recall*, and *recall mode*.

recall mode
A mode that is assigned to a migrated file with the **dsmattr** command that determines how the file is processed when it is recalled. It determines whether the file is stored on the local file system, is migrated back to Tivoli Storage Manager storage when it is closed, or is read from Tivoli Storage Manager storage without storing it on the local file system.

receiver
A server repository that contains a log of server and client messages as events. For example, a receiver can be a file exit, a user exit, or the Tivoli Storage Manager server console and activity log. See also *event*.

reclamation
The process of consolidating the remaining data from many sequential-access volumes onto fewer, new sequential-access volumes.

reclamation threshold

The percentage of space that a sequential-access media volume must have before the server can reclaim the volume. Space becomes reclaimable when files are expired or are deleted.

reconciliation

The process of synchronizing a file system with the Tivoli Storage Manager server, and then removing old and obsolete objects from the Tivoli Storage Manager server.

recovery log

A log of updates that are about to be written to the database. The log can be used to recover from system and media failures. The recovery log consists of the active log (including the log mirror) and archive logs.

register

To define a client node or administrator ID that can access the server.

registry

A repository that contains access and configuration information for users, systems, and software.

resident file

On a Windows system, a complete file on a local file system that might also be a migrated file because a migrated copy can exist in Tivoli Storage Manager storage. On a UNIX or Linux system, a complete file on a local file system that has not been migrated or premigrated, or that has been recalled from Tivoli Storage Manager storage and modified. Contrast with *stub file* and *premigrated file*. See *migrated file*.

restore

To copy information from its backup location to the active storage location for use. For example, to copy information from server storage to a client workstation.

retention

The amount of time, in days, that inactive backed-up or archived files are kept in the storage pool before they are deleted. Copy group attributes and default retention grace periods for the domain define retention.

retrieve

To copy archived information from the

storage pool to the workstation for use. The retrieve operation does not affect the archive version in the storage pool.

roll back

To remove changes that were made to database files since the last commit point.

root user

A system user who operates without restrictions. A root user has the special rights and privileges needed to perform administrative tasks.

S

SAN See *storage area network*.

schedule

A database record that describes client operations or administrative commands to be processed. See *administrative command schedule* and *client schedule*.

scheduling mode

The type of scheduling operation for the server and client node that supports two scheduling modes: client-polling and server-prompted.

scratch volume

A labeled volume that is either blank or contains no valid data, that is not defined, and that is available for use.

script A series of commands, combined in a file, that carry out a particular function when the file is run. Scripts are interpreted as they are run. Contrast with *Tivoli Storage Manager command script*.

Secure Sockets Layer (SSL)

A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

selective backup

The process of backing up certain files or directories from a client domain. The files that are backed up are those that are not excluded in the include-exclude list. The files must meet the requirement for serialization in the backup copy group of the management class that is assigned to each file. Contrast with *incremental backup*.

selective migration

The process of copying user-selected files

from a local file system to Tivoli Storage Manager storage and replacing the files with stub files on the local file system. Contrast with *threshold migration* and *demand migration*.

selective recall
The process of copying user-selected files from Tivoli Storage Manager storage to a local file system. Contrast with *transparent recall*.

serialization
The process of handling files that are modified during backup or archive processing. See *dynamic serialization*, *static serialization*, *shared static serialization*, and *shared dynamic serialization*.

server A software program or a computer that provides services to other software programs or other computers.

server options file
A file that contains settings that control various server operations. These settings affect such things as communications, devices, and performance.

server-prompted scheduling mode
A client/server communication technique where the server contacts the client node when tasks must be done. Contrast with *client-polling scheduling mode*.

server storage
The primary, copy, and active-data storage pools that are used by the server to store user files such as backup versions, archive copies, and files migrated from space manager client nodes (space-managed files). See also *active-data pool*, *primary storage pool*, *copy storage pool*, *storage pool volume*, and *volume*.

session
A logical or virtual connection between two stations, software programs, or devices on a network that allows the two elements to communicate and exchange data.

session resource usage
The amount of wait time, processor time, and space that is used or retrieved during a client session.

shared dynamic serialization
A value for serialization that specifies that a file must not be backed up or archived if it is being modified during the operation. Tivoli Storage Manager retries the backup or archive operation a number of times; if the file is being modified during each attempt, Tivoli Storage Manager will back up or archive the file on its last try. See also *serialization*. Contrast with *dynamic serialization*, *shared static serialization*, and *static serialization*.

shared library
A library device that is used by multiple storage manager servers.

shared static serialization
A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. Tivoli Storage Manager attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *static serialization*.

snapshot
An image backup type that consists of a point-in-time view of a volume.

space-managed file
A file that is migrated from a client node by the space manager client. The space manager client recalls the file to the client node on demand.

space management
The process of keeping sufficient free storage space available on a local file system for new data by migrating files to server storage. Synonymous with *hierarchical storage management*.

space manager client
A program that runs on a UNIX or Linux system to manage free space on the local file system by migrating files to server storage. The program can recall the files either automatically or selectively. Also called *hierarchical storage management (HSM) client*.

space monitor daemon
A daemon that checks space usage on all file systems for which space management is active, and automatically starts threshold migration when space usage on a file system equals or exceeds its high threshold.

sparse file
A file that is created with a length greater than the data it contains, leaving empty spaces for the future addition of data.

special file
On AIX, UNIX, or Linux systems, a file that defines devices for the system, or temporary files that are created by processes. There are three basic types of special files: first-in, first-out (FIFO); block; and character.

SSL See *Secure Sockets Layer*.

stabilized file space
A file space that exists on the server but not on the client.

stanza A group of lines in a file that together have a common function or define a part of the system. Each stanza is identified by a name that occurs in the first line of the stanza. Depending on the type of file, a stanza is ended by the next occurrence of a stanza name in the file, or by an explicit end-of-stanza marker. A stanza can also be ended by the end of the file.

startup window
A time period during which a schedule must be initiated.

static serialization
A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. If the file is in use during the first attempt, the storage manager cannot back up or archive the file. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *shared static serialization*.

storage agent
A program that enables the backup and restoration of client data directly to and from storage attached to a storage area network (SAN).

storage area network (SAN)
A dedicated storage network that is tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

storage hierarchy
(1) A logical order of primary storage pools, as defined by an administrator. The order is typically based on the speed and capacity of the devices that the storage pools use. The storage hierarchy is defined by identifying the next storage pool in a storage pool definition. See also *storage pool*.

(2) An arrangement of storage devices with different speeds and capacities. The levels of the storage hierarchy include: main storage, such as memory and direct-access storage device (DASD) cache; primary storage (DASD containing user-accessible data); migration level 1 (DASD containing data in a space-saving format); and migration level 2 (tape cartridges containing data in a space-saving format).

storage pool
A named set of storage volumes that are the destination that is used to store client data. A storage pool contains backup versions, archive copies, and files that are migrated from space manager client nodes. A primary storage pool is backed up to a copy storage pool. See also *primary storage pool*, *copy storage pool*, and *active-data pool*.

storage pool volume
A volume that has been assigned to a storage pool. See also *volume*, *active-data pool*, *copy storage pool*, and *primary storage pool*.

storage privilege class
A privilege class that gives an administrator the authority to control how storage resources for the server are allocated and used, such as monitoring the database, the recovery log, and server storage. See also *privilege class*.

stub A shortcut on the Windows file system that is generated by the hierarchical storage management (HSM) client for a migrated file that allows transparent user access. A stub is the sparse file representation of a migrated file, with a reparse point attached.

stub file
A file that replaces the original file on a local file system when the file is migrated to storage. A stub file contains the information that is necessary to recall a migrated file from Tivoli Storage Manager storage. It also contains additional

information that can be used to eliminate the need to recall a migrated file.

stub file size
The size of a file that replaces the original file on a local file system when the file is migrated to Tivoli Storage Manager storage. The size that is specified for stub files determines how much leader data can be stored in the stub file. The default for stub file size is the block size defined for a file system minus 1 byte.

subscription
In a Tivoli environment, the process of identifying the subscribers that the profiles are distributed to. For Tivoli Storage Manager, a subscription is the process by which a managed server receives configuration information associated with a particular profile on a configuration manager. See also *managed server*, *configuration manager*, and *profile*.

system privilege class
A privilege class that gives an administrator the authority to issue all server commands. See also *privilege class*.

Systems Network Architecture (SNA)
The description of the logical structure, formats, protocols, and operational sequences for transmitting information through and controlling the configuration and operation of networks.

T

tape library
A set of equipment and facilities that support an installation's tape environment. The tape library can include tape storage racks, mechanisms for automatic tape mounting, a set of tape drives, and a set of related tape volumes mounted on those drives.

tape volume prefix
The high-level-qualifier of the file name or the data set name in the standard tape label.

target node
A client node for which other client nodes (called agent nodes) have been granted proxy authority. The proxy authority allows the agent nodes to perform operations such as backup and restore on behalf of the target node, which owns the data.

TCA See *trusted communications agent*.

TCP/IP
See *Transmission Control Protocol/Internet Protocol*.

threshold migration
The process of moving files from a local file system to Tivoli Storage Manager storage based on the high and low thresholds that are defined for the file system. Contrast with *demand migration*, *selective migration*, and *migration job*.

throughput
In storage management, the total bytes in the workload, excluding overhead, that are backed up or restored, divided by elapsed time.

timeout
A time interval that is allotted for an event to occur or complete before operation is interrupted.

timestamp control mode
A mode that determines whether commands preserve the access time for a file or set it to the current time.

Tivoli Storage Manager command script
A sequence of Tivoli Storage Manager administrative commands that are stored in the database of the Tivoli Storage Manager server. The script can run from any interface to the server. The script can include substitution for command parameters and conditional logic.

tombstone object
A small subset of attributes of a deleted object. The tombstone object is retained for a specified period, and at the end of the specified period, the tombstone object is permanently deleted.

Transmission Control Protocol/Internet Protocol (TCP/IP)
An industry-standard, nonproprietary set of communication protocols that provides reliable end-to-end connections between applications over interconnected networks of different types.

transparent recall
The process that is used to automatically recall a file to a workstation or file server when the file is accessed. See also *recall mode*. Contrast with *selective recall*.

trusted communications agent (TCA)
A program that handles the sign-on password protocol when clients use password generation.

U

UCS-2 A 2-byte (16-bit) encoding scheme based on ISO/IEC specification 10646-1. UCS-2 defines three levels of implementation: Level 1-No combining of encoded elements allowed; Level 2-Combining of encoded elements is allowed only for Thai, Indic, Hebrew, and Arabic; Level 3-Any combination of encoded elements are allowed.

UNC See *Universal Naming Convention name*.

Unicode
A character encoding standard that supports the interchange, processing, and display of text that is written in the common languages around the world, plus some classical and historical texts. The Unicode standard has a 16-bit character set defined by ISO 10646.

Unicode-enabled file space
Unicode file space names provide support for multilingual workstations without regard for the current locale.

Unicode transformation format 8
Unicode Transformation Format (UTF), 8-bit encoding form, which is designed for ease of use with existing ASCII-based systems. The CCSID value for data in UTF-8 format is 1208.

Universal Naming Convention (UNC) name
A name that is used to access a drive or directory containing files shared across a network. The UNC name includes the system name and a SharePoint name that represents the shared drive or directory.

Universally Unique Identifier (UUID)
The 128-bit numerical identifier that is used to ensure that two components do not have the same identifier.

UTF-8 See *Unicode transformation format 8*.

UUID See *Universally Unique Identifier*.

V

validate
To check a policy set for conditions that can cause problems if that policy set becomes the active policy set. For

example, the validation process checks whether the policy set contains a default management class.

version
A backup copy of a file stored in server storage. The most recent backup copy of a file is the active version. Earlier copies of the same file are inactive versions. The number of versions retained by the server is determined by the copy group attributes in the management class.

virtual file space
A representation of a directory on a network-attached storage (NAS) file system as a path to that directory.

virtual volume
An archive file on a target server that represents a sequential media volume to a source server.

volume
A discrete unit of storage on disk, tape or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control. See also *scratch volume*, and *storage pool volume*.

volume history file
A file that contains information about volumes that have been used by the server for database backups and for export of administrator, node, policy, or server data. The file also has information about sequential-access storage pool volumes that have been added, reused, or deleted. The information is a copy of volume information that is recorded in the server database.

Volume Shadow Copy Service
A set of Microsoft application-programming interfaces (APIs) that you can use to create shadow copy backups of volumes, exact copies of files, including all open files, and so on.

VSS See *Volume Shadow Copy Service*.

VSS Backup
A backup operation that uses Microsoft Volume Shadow Copy Service (VSS) technology. The backup operation produces an online snapshot (point-in-time consistent copy) of Exchange data. This copy can be stored

on local shadow volumes or on Tivoli Storage Manager server storage.

VSS Fast Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Exchange database files and log files) that reside on local shadow volumes.

VSS Instant Restore

A volume-level hardware-assisted Microsoft Volume Shadow Copy Service (VSS) function where target volumes that contain the snapshot are copied back to the original source volumes.

VSS offloaded backup

A backup operation that uses a Microsoft Volume Shadow Copy Service (VSS) hardware provider (installed on an alternate system) to move IBM Data Protection for Exchange data to the Tivoli Storage Manager server. This type of backup operation shifts the backup load from the production system to another system.

VSS Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Exchange database files and log files) that reside on Tivoli Storage Manager server storage to their original location.

W

wildcard character

A special character such as an asterisk (*) or a question mark (?) that can be used to represent one or more characters. Any character or set of characters can replace the wildcard character.

workstation

A configuration of input/output equipment at which an operator works. A workstation is a terminal or microcomputer at which a user can run applications and that is usually connected to a mainframe or a network.

worldwide name

A 64-bit, unsigned name identifier that is unique.

workload partition (WPAR)

A partition within a single operating system instance.

Index

Special characters

Numerics

A

B

IBM ®

Program Number: 5608-E01, 5608-E02, 5608-E07

Printed in USA