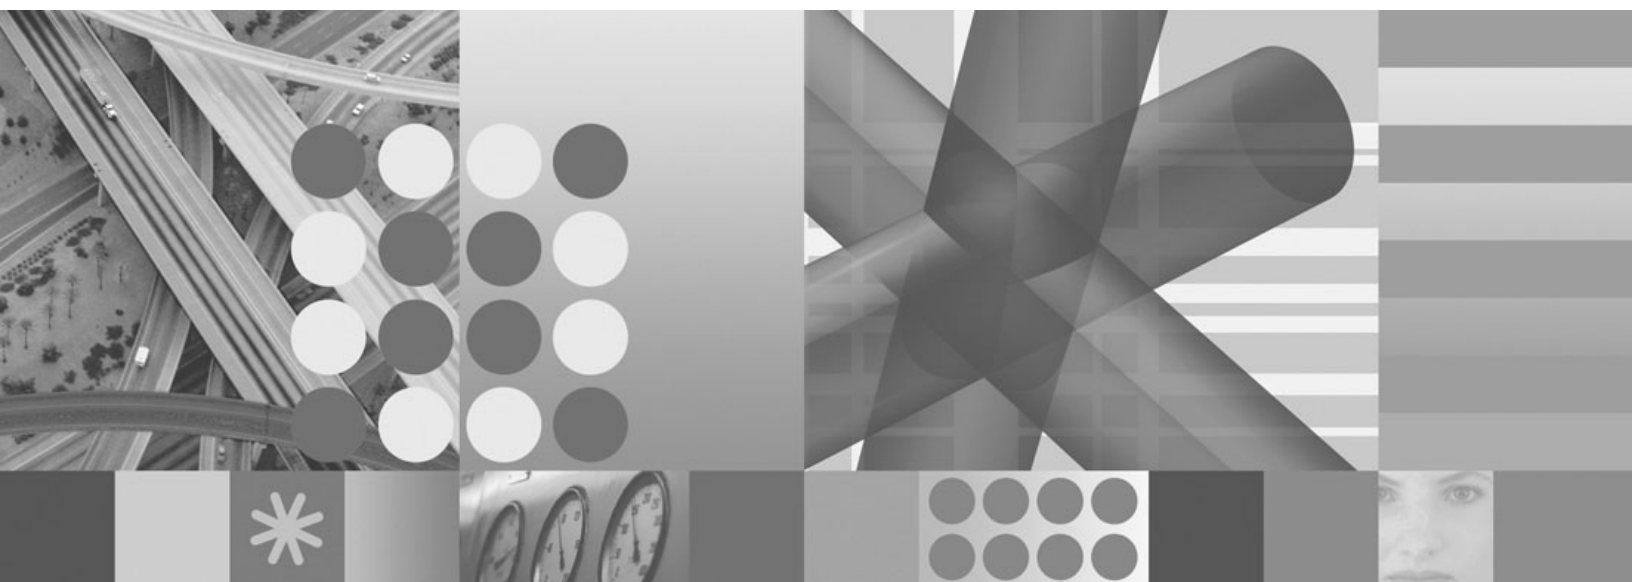




Backup-Archive Clients Installation and User's Guide



Backup-Archive Clients Installation and User's Guide

Note

Before using this information and the product it supports, read the general information under “Notices” on page 345.

Edition notice

This edition applies to Version 5.5 of IBM Tivoli Storage Manager (5608-ISM), IBM Tivoli Storage Manager Extended Edition (5608-ISX), IBM Tivoli Storage Manager for Storage Area Networks (5608-SAN), and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters.

Order publications through your IBM representative or the IBM branch office serving your locality.

Your feedback is important in helping to provide the most accurate and high-quality information. If you have comments about this manual or any other IBM Tivoli Storage Manager documentation, see “Contacting customer support” on page xiii.

© Copyright International Business Machines Corporation 1993, 2007. All rights reserved.

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

Contents

Tables vii

About this book ix

Who should read this manual	ix
IBM Tivoli Storage Manager Web site	ix
Conventions used in this book	x
Reading syntax diagrams	x
Related publications	xii
Downloading or ordering publications	xiii
Accessibility	xiii
Contacting customer support	xiii
Reporting a problem	xiv
Internet	xiv

Summary of changes for Tivoli Storage Manager xvii

Technical changes for Version 5.5 - November 2007	xvii
Technical changes for Version 5.4 - January 2007	xvii
Technical changes for Version 5.3 - March 2006	xviii
Technical changes for Version 5.3 - October 2005	xix

Chapter 1. Installing Tivoli Storage Manager 1

Migrating from earlier versions	1
Upgrade path for clients and servers	1
Migrating to the Tivoli Storage Manager Version 5.5 Macintosh client from releases prior to 5.3.2.	2
Migrating to the Unicode-enabled client	4
Macintosh client environment.	6
Client components	7
Hardware requirements.	7
Software requirements	7
Communications method	7
Installing the Macintosh client program (required system administrator task).	7
Enabling non-administrators to manage their own data (required system administrator task).	11

Chapter 2. Configuring Tivoli Storage Manager 13

Root user, authorized user, and non-authorized user tasks.	13
Using TSM Tools for Administrators	15
Creating and modifying the client system options file (required root user or authorized user task)	15
Creating a default client user options file (optional root user or authorized user task)	16
Creating a customized client user options file (optional user task).	17
Setting options in an options file	18
Setting processing environment variables (optional)	19
Setting T and C shell variables	20
Setting Bash and Bourne shell variables	20
Configuring the Web client	20

Configuring the client scheduler	21
Starting the client scheduler	22
Using the command-line client to handle scheduled events	22
Defining a schedule	23
Configuring Tivoli Storage Manager client/server communication across a firewall	23
Registering your workstation with a server (required root user or authorized user task)	25
Using closed registration	26
Using open registration	26
Creating an include-exclude list (optional root user or authorized user task)	26
Using include-exclude options	28
Controlling symbolic link and alias processing.	33
Controlling compression and encryption processing.	33
Testing an include-exclude list with the Preview command	34
Processing include and exclude options	34

Chapter 3. Getting started 39

Tivoli Storage Manager client authentication	39
Starting a GUI session.	40
Using your Tivoli Storage Manager password	40
Setup wizard	40
Starting a command-line session	41
Using batch mode	41
Using interactive mode	41
Starting: Additional considerations	42
Starting a Web client session.	42
Setting user privileges	43
Starting the client scheduler automatically	43
Changing your password.	43
Sorting file lists using the Tivoli Storage Manager Java GUI	44
Displaying online help	45
Ending a session.	45
Online forums	45
Other sources of online help.	46

Chapter 4. Backing up your data 47

Planning your backups	47
When to back up and when to archive files.	47
Pre-backup considerations	48
How Tivoli Storage Manager maintains volume names	48
Volume naming precautions	49
Volume naming precautions on dual boot systems.	49
Unicode considerations	50
Configure memory-constrained systems to run incremental backups	50
Configure systems with large numbers of files to run incremental backups	50

Using an include-exclude options list to control processing	51
Encrypting data during backup or archive operations	52
Supported file systems	53
Maximum file size for backup and restore	54
How Tivoli Storage Manager handles long user and group names	54
How Tivoli Storage Manager handles Filevault home directories	55
How Tivoli Storage Manager handles the Time Machine backup disk	55
Performing an incremental, selective, or incremental-by-date backup	56
Full and partial incremental backup	56
Incremental-by-date backup	57
Comparing full incremental, partial incremental, and incremental-by-date backups	58
Selective backup	58
Backing up data using the Java GUI	59
Backing up data using the command line	60
Deleting backup data	61
Displaying backup processing status	62
Backup with client node proxy support	63
Scheduling with client node proxy support	65
Backup: Additional considerations	66
Understanding how files are stored	66
Understanding how files are managed	67
Understanding how hard links are handled	67
Understanding how NFS hard and soft mounts are handled	67
Understanding how deleted file systems are handled	68
Backing up open files	68
Using wildcard characters	69

Chapter 5. Restoring your data 71

Restoring an active or inactive backup	71
Performing large restore operations	72
Standard query restore, no query restore, and restartable restore	72
Standard restore process	73
No query restore process	73
Restartable restore process	74
Restoring files and directories	74
Restoring data using the GUI	74
Restoring data using the command line	75
Performing point-in-time restores	76
Restoring data from a backupset	77
Restoring an entire or partial backupset	78
Restoring backupsets using the Web client GUI	78
Restoring backupsets using the command-line client	79
Related tasks and considerations	80
Authorizing another user to restore or retrieve your files	80
Restoring or retrieving files from another client node	81
Restoring or retrieving your files to another workstation	81
Restoring a disk in case of disk loss	82

Deleting file spaces	82
Overlapping icons	82

Chapter 6. Archiving and retrieving your data. 85

Archiving files	85
Maximum file size for archive and retrieve operations	85
Archiving data using the GUI	85
Archiving data using the command line	86
Archiving data with client node proxy	87
Deleting archive data	89
Archive: Advanced considerations	89
Retrieving archives	90
Retrieving data using the GUI	90
Retrieving data using the command line	91
Understanding how your archives are managed	91

Chapter 7. Automating tasks 93

Specifying scheduling options	93
Displaying information about scheduled work	93
Displaying information about completed work	95
Scheduling options for commands	96
Enabling or disabling scheduled commands	96
Client return codes	96

Chapter 8. Understanding storage management policies 99

Using policy domains and policy sets	99
Using management classes and copy groups	99
Displaying information about management classes and copy groups	101
Copy group name	102
Copy type	102
Copy frequency	102
Versions data exists	102
Versions data deleted	102
Retain extra versions	102
Retain only version	103
Copy serialization	103
Copy mode	103
Copy destination	104
Retain versions	104
Selecting a management class for files	104
Assigning a management class to files	105
Overriding the management class for archived files	106
Selecting a management class for folders	106
Binding management classes to files	106
Rebinding backup versions of files	107
Using a retention grace period	107
Using event-based policy retention protection	108
Archiving files on a data retention server	109

Chapter 9. Using processing options 111

Overview of processing options	111
Communication options	111
TCP/IP options	112
Server and node Node options	112
Server options	112
Node options	113

Backup and archive processing options	114	Language.	191
Restore and retrieve processing options.	117	Latest	192
Scheduling options	118	Localbackupset	193
Format and language options	120	Location	194
Command processing options	120	Mailprog	195
Authorization options	120	Managedservices	196
Error processing options.	121	Maxcmdretries	198
Transaction processing options	121	Memoryefficientbackup	199
Diagnostics options	122	Nfstimeout	201
Using options with commands	123	Nodename	202
Entering options with a command	123	Noprompt	203
Initial command-line-only options	126	Numberformat	204
Client options reference	127	Optfile	205
Archmc	128	Password.	206
Asnodename	129	Passwordaccess.	207
Auditlogging	131	Passworddir.	209
Auditlogname	133	Pick	210
Autofsrename	135	Pitdate.	211
Backupsetname.	137	Pittime	212
Changingretries	138	Postschedulecmd/Postnschedulecmd	213
Collocatebyfilespec	139	Preschedulecmd/Prenschedulecmd	215
Commmethod	140	Preservepath	217
Commrestartduration.	141	Queryschedperiod.	220
Commrestartinterval	142	Quiet	221
Compressalways	143	Removeoperandlimit	222
Compression	144	Replace	223
Console	145	Resourceutilization	225
Dateformat	146	Retryperiod	228
Defaultserver	148	Revokeremoteaccess	229
Deletefiles	149	Schedcmddisabled.	230
Description	150	Schedlogmax	231
Detail	151	Schedlogname	232
Dirmc	152	Schedlogretention	233
Dironly	153	Schedmode	235
Disablenqr	154	Schedrestretrdisabled.	237
Diskbuffsize	155	Scrolllines	238
Diskcachelocation	156	Scrollprompt	239
Domain	157	Servename	240
Editor	159	Sessioninitiation	242
Enablearchiveretentionprotection	160	Skipacl	244
Encryptiontype.	161	Skipaclupdatecheck	245
Encryptkey	162	Srvprepostscheddisabled	246
Errorlogmax.	164	Subdir.	247
Errorlogname	165	Tapeprompt	249
Errorlogretention	166	Tcpadminport	251
Exclude options	168	Tcpbuffsize	252
Filelist.	171	Tcpcadaddress	253
Filename	173	Tcpclientaddress	254
Filesonly	174	Tcpclientport	255
Fromdate.	175	Tcpnodelay	256
Fromnode	176	Tcpport	257
Fromowner	177	Tcpserveraddress	258
Fromtime.	178	Tcpwindowsize.	259
Groups	179	Timeformat	260
Guitreeviewafterbackup	180	Todate.	261
Httpport	181	Totime.	262
Ifnewer	182	Txnbytelimit.	263
Inactive	183	Type	264
Inlexcl	184	Users	265
Include options.	186	V2archive	266
Incrbydate	190	Verbose	267

Virtualnodename	268
Webports	269

Chapter 10. Using commands 271

Starting and ending a client command session	272
Process commands in batch mode	273
Process commands in interactive mode	273
Entering client commands	274
Command name	274
Options	274
Parameters	274
File specification syntax	275
Remembering previous commands	276
Using wildcard characters	277
Entering commands	277
Client commands reference	278
Archive	279
Cancel Restore	281
Delete Access	282
Delete Archive	283
Delete Backup	285
Delete Filespace	288
Expire	289
Help	290
Incremental	291
Loop	295
Macro	296
Preview	297
Query Access	298
Query Archive	299
Query Backup	301
Query Backupset	303
Query Filespace	305
Query Inclexcl	306
Query Mgmtclass	307
Query Node	308
Query Options	309
Query Restore	310
Query Schedule	311
Query Session	312
Query Systeminfo	313
Restart Restore	315
Restore	316

Restore Backupset	320
Retrieve	323
Schedule	326
Selective	328
Set Access	330
Set Event	332
Set Password	334

Appendix. Using the Tivoli Storage Manager central scheduler 335

Overview of the Tivoli Storage Manager scheduler	335
Handling spaces in file names in schedule definitions	336
Preferential start times for certain nodes	336
Understanding scheduler processing options	337
Handling return codes from preschedulecmd and postschedulecmd scripts	338
Using the client acceptor to manage scheduler services versus the traditional scheduler services	339
Setting the client scheduler process to run as a background task and start automatically at boot time	339
Changing processing options used by the scheduler service	340
Managing multiple schedule requirements on one system	340
Restarting the scheduler process on a remote system	341
Troubleshooting the Tivoli Storage Manager scheduler	342
Investigating abnormal scheduled events	342
Checking the server activity log	343
Inspecting the Client Schedule Log	343
Understanding the scheduler log	343

Notices 345

Trademarks	346
----------------------	-----

Glossary 349

Index 357

Tables

1. Typographical conventions	x	30. Backup and archive processing options	114
2. IBM Tivoli Storage Manager client and server publications	xii	31. Restore and retrieve processing options	117
3. Optimized option default values	xxi	32. Scheduling options.	118
4. Macintosh installation packages	8	33. Format and language options	120
5. Mac OS X authorization tools and associated Tivoli Storage Manager applications	13	34. Command processing options	120
6. Root user, authorized user, and non-authorized user tasks	14	35. Authorization options.	120
7. Option for excluding directories.	28	36. Error processing options	121
8. Options for controlling backup and archive processing	29	37. Transaction processing options.	121
9. Wildcard and other special characters.	31	38. Diagnostics options	122
10. Using wildcard characters with include and exclude patterns	32	39. Client command options.	124
11. Options for controlling symbolic link and alias processing	33	40. Options that are valid on the initial command line only	126
12. Options for controlling compression and encryption processing	33	41. Entering commands	271
13. Working with your files using the Tivoli Storage Manager GUI	44	42. Commands	271
14. Backup: Primary tasks	47	43. Command recall and edit functions	276
15. Planning your backups.	47	44. Wildcard characters	277
16. Supported file systems.	53	45. Archive command: Related options	279
17. Maximum file size for backup and restore	54	46. Delete Archive command: Related options	283
18. Command line backup examples	60	47. Delete Backup command: Related options	285
19. Client command line informational messages	62	48. Delete Filespace command: Related options	288
20. Restore: Primary tasks	71	49. Expire command: Related options.	289
21. Command line restore examples	75	50. Incremental command: Related options	293
22. Archive and retrieve: Primary tasks	85	51. Query Archive command: Related options	299
23. Maximum file size for archive and retrieve	85	52. Query Backup command: Related options	301
24. Command line archive examples	86	53. Query Backupset command: Related options	303
25. command-line examples of retrieving archives	91	54. Query Filespace command: Related options	305
26. Client return codes and their meanings	97	55. Query Mgmtclass command: Related options	307
27. Default values in the standard management class	101	56. Query Node command: Related options	308
28. TCP/IP options	112	57. Query Options command: Related options	309
29. Server and node options	113	58. Query Systeminfo command: Related options	314
		59. Restore command: Related options	317
		60. Restore Backupset command: Related options	321
		61. Retrieve command: Related options	323
		62. Schedule command: Related options	326
		63. Selective command: Related options	329
		64. Scheduled client event status codes	342
		65. Schedule log information	344

About this book

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. See “Related publications” on page xii for available server publications.
- 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. For more information about the Administrative client, see “Related publications” on page xii for available Tivoli Storage Manager Administrator’s Reference publications.

Who should read this manual

This manual provides instructions for an end-user to install, configure, and use the Tivoli Storage Manager client. For installation information and supported operating system levels, see Chapter 1, “Installing Tivoli Storage Manager,” on page 1. For configuration information, see Chapter 2, “Configuring Tivoli Storage Manager,” on page 13.

This manual assumes that you are familiar with your workstation, your operating system, and your basic system administration.

Tivoli Storage Manager tasks that can only be performed by authorized users and root users are identified by the phrases, **authorized user** and **root user**. 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 and whose owner execution permission bit is set to **s**. In the following example, the user **tivoli** is an authorized user while running **dsmd** since the **dsmd** owner execution permission bit is set to **s**:

See “Root user, authorized user, and non-authorized user tasks” on page 13 for more information about **authorized user** and **root user** tasks.

IBM Tivoli Storage Manager Web site

Technical support information and publications are available at the following address:

<http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>

By accessing the Tivoli Storage Manager home page, you can access subjects that interest you. You can also access current Tivoli Storage Manager product information.

Conventions used in this book

This book uses the following typographical conventions:

Table 1. Typographical conventions

Example	Description
<code>dsmc.nlm</code>	A series of lowercase letters with an extension indicates Tivoli Storage Manager program file names.
archive	Boldface type indicates a command that you type at a workstation, such as a command you type on a command line.
<i>dateformat</i>	Boldface italic type indicates a Tivoli Storage Manager option. The bold type is used to introduce the option, or used in an example. Occasionally, file names are entered in boldface italic for emphasis.
<i>filespec</i>	Italicized type indicates either the name of a parameter, a new term, or a placeholder for information that you provide. Italics are also used for emphasis in the text.
<code>maxcmdretries</code>	Monospaced type represents fragments of a program or information as it would display on a screen.
plus sign (+)	A plus sign between two keys indicates you should press both keys at the same time.

Reading syntax diagrams

This section describes how to read the syntax diagrams used in this manual. To read a syntax diagram, follow the path of the line. Read from left to right, and top to bottom.

- The ►— symbol indicates the beginning of a syntax diagram.
- The —→ symbol at the end of a line indicates the syntax diagram continues on the next line.
- The ►— symbol at the beginning of a line indicates a syntax diagram continues from the previous line.
- The —→◄ symbol indicates the end of a syntax diagram.

Syntax items, such as a keyword or variable, can be:

- On the line (required element)
- Above the line (default element)
- Below the line (optional element).

Syntax diagram description	Example
----------------------------	---------

Abbreviations:

Uppercase letters denote the shortest acceptable truncation. If an item appears entirely in uppercase letters, it cannot be truncated.

You can type the item in any combination of uppercase or lowercase letters.

In this example, you can enter KEYWO, KEYWORD, or KEYWOrd.



Symbols:

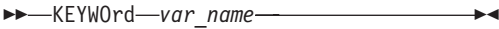
Enter these symbols exactly as they appear in the syntax diagram.

*	Asterisk
{ }	Braces
:	Colon
,	Comma
=	Equal Sign
-	Hyphen
()	Parentheses
.	Period
	Space

Variables:

Italicized lowercase items (*var_name*) denote variables.

In this example, you can specify a *var_name* when you enter the KEYWORD command.



Repetition:

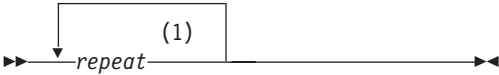
An arrow returning to the left means you can repeat the item.



A character or space within the arrow means you must separate repeated items with that character or space.



A footnote by the arrow references the number of times you can repeat the item.



Notes:





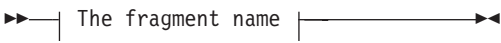

- Specify *repeat* as many as 5 times.

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.

Syntax diagram description	Example
Optional choice: When an item is below the line, that item is optional. In the first example, you can choose A or nothing at all. 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.	 
Defaults: Defaults are above the line. The default is selected unless you override it. You can override the default by including an option from the stack below the line. In this example, A is the default. You can override A by choosing B or C. You can also specify the default explicitly.	
Repeatable choices: A stack of items followed by an arrow returning to the left means you can select more than one item or, in some cases, repeat a single item. In this example, you can choose any combination of A, B, or C.	
Syntax fragments: Some diagrams, because of their length, must fragment the syntax. The fragment name appears between vertical bars in the diagram. The expanded fragment appears between vertical bars in the diagram after a heading with the same fragment name.	 The fragment name: 

Related publications

Table 2 lists the IBM Tivoli Storage Manager client and server publications that are referred to in this manual.

Table 2. IBM Tivoli Storage Manager client and server publications

Publication title	Order number
Client publications	
<i>IBM Tivoli Storage Manager Messages</i>	SC32-0140
<i>IBM Tivoli Storage Manager for Windows Backup-Archive Clients Installation and User's Guide</i>	SC32-0146
<i>IBM Tivoli Storage Manager for NetWare Backup-Archive Clients Installation and User's Guide</i>	SC32-0144
<i>IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide</i>	SC32-0145
Server publications	

Table 2. IBM Tivoli Storage Manager client and server publications (continued)

Publication title	Order number
<i>IBM Tivoli Storage Manager for AIX Installation Guide</i>	SC32-0134
<i>IBM Tivoli Storage Manager for HP-UX Installation Guide</i>	SC32-0135
<i>IBM Tivoli Storage Manager for Linux Installation Guide</i>	SC32-0136
<i>IBM Tivoli Storage Manager for z/OS Installation Guide</i>	SC32-0139
<i>IBM Tivoli Storage Manager for Sun Solaris Installation Guide</i>	SC32-0137
<i>IBM Tivoli Storage Manager for Windows Installation Guide</i>	SC32-0138

Downloading or ordering publications

All Tivoli publications are available for electronic download or order from the IBM Publications Center: <http://www.ibm.com/shop/publications/order/>.

If you have questions or comments regarding Tivoli publications and product documentation, visit <http://www.ibm.com/software/tivoli/contact.html> to send an e-mail.

The International Technical Support Center (ITSC) publishes Redbooks®, which are books on specialized topics such as using Tivoli Storage Manager to back up databases. You can order publications through your IBM representative or the IBM branch office serving your locality. You can also search for and order books of interest to you at the IBM Redbooks Web site at this address:

<http://www.ibm.com/redbooks/>

Tivoli Field Guides are designed to address specific technical scenarios or concepts that are often complex to implement or difficult to understand. All completed field guides are available free to registered customers and internal IBM employees at the following Web site:

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

Accessibility

Accessibility features help users with a physical disability, such as restricted mobility or limited vision, to use software products successfully. With this product, you can use assistive technologies to hear and navigate the interface. You can also use the keyboard instead of the mouse to operate all features of the graphical user interface.

Contacting customer support

For support for this or any Tivoli product, you can contact Tivoli Customer Support in one of the following ways:

- Visit the Tivoli Storage Manager technical support Web site at:
<http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>
- Visit the IBM Software Support Toolbar Web site at:
<http://www.ibm.com/software/support/toolbar/>
- Submit a problem management record (PMR) electronically at **IBMSERV/IBMLINK**. You can access from the IBM Web site at:
<http://www.ibm.com/ibmlink/>
- Submit a problem management record (PMR) electronically from the IBM Web site at:

<http://www.ibm.com/software/support/probsub.html>.

Customers in the United States can also call 1-800-IBM-SERV (1-800-426-7378).

International customers should consult the Web site for customer support telephone numbers.

You can also review the *IBM Software Support Guide*, which is available on our Web site at <http://techsupport.services.ibm.com/guides/handbook.html>.

When you contact IBM Software Support, be prepared to provide identification information for your company so that support personnel can readily assist you. Company identification information is needed to register for online support available on the Web site.

The support Web site offers extensive information, including a guide to support services (IBM Software Support Guide); frequently asked questions (FAQs); and documentation for all IBM Software products, including Release Notes, Redbooks, and white papers, defects (APARs), and solutions. The documentation for some product releases is available in both PDF and HTML formats. Translated documents are also available for some product releases.

We are very interested in hearing about your experience with Tivoli products and documentation. We also welcome your suggestions for improvements. If you have comments or suggestions about our documentation, complete our customer feedback survey at the following website, by selecting the Feedback link in the left navigation bar:

<http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>

Reporting a problem

Have the following information ready when you report a problem:

- The Tivoli Storage Manager server version, release, modification, and service level number. You can get this information by entering the **query status** command at the Tivoli Storage Manager administrative command line.
- It is recommended that you use the Tivoli Storage Manager client **query systeminfo** command with the *filename* option to gather Tivoli Storage Manager system information and output this information to a file. This information is intended primarily as an aid for IBM support to assist in diagnosing problems.
- The Tivoli Storage Manager client version, release, modification, and service level number. You can get this information by entering **dsmc** at the command line.
- The communication protocol (for example, TCP/IP), version, and release number you are using.
- The activity you were doing when the problem occurred, listing the steps you followed before the problem occurred.
- The exact text of any error messages.

Internet

To search multiple Internet resources for your product, go to the support web site for the product:

- <http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>

and look for the section to search the support knowledge base. From this section, you can search a variety of resources including:

- IBM technotes
- IBM downloads
- IBM Redbooks
- Forums and newsgroups

You can also check the following IBM Software Support Toolbar Web site, and choose **Tivoli** → **Tivoli developerWorks®** → **Forums**. <http://www.ibm.com/software/support/toolbar/>

A newsgroup, *listserv@marist.edu*, is implemented by a third party. IBM supports this newsgroup on a best-effort basis only. See “Online forums” on page 45 for more information.

Summary of changes for Tivoli Storage Manager

This section summarizes changes made to the Tivoli Storage Manager product and this publication. Technical changes to the text are indicated by vertical lines to the left of the change.

Technical changes for Version 5.5 - November 2007

The following changes have been made to the product for this edition:

Transparent encryption support

Support has been added for transparent encryption, using the *encryptkeygenerate* option. Refer to “Encryptkey” on page 162 for more information about this option.

TCP/IP Version 6 support

The IBM Tivoli Storage Manager client now supports the TCP/IP Version 6 (IPv6) protocol and can be used in a dedicated IPv6, IPv4, or intermixed IP environment.

Auditlogging and auditlogname options

The *auditlogging* option generates an audit log with an entry for each file that is processed during an incremental, selective, archive, restore, or retrieve operation. Refer to “Auditlogging” on page 131 for more information about this option.

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. Refer to “Auditlogname” on page 133 for more information about this option.

Technical changes for Version 5.4 - January 2007

The following changes have been made to the product for this edition:

Password file format changes

The format of the Tivoli Storage Manager password file changed in Version 5.3.4. Your existing password file will be migrated to the new format during the first session with a Tivoli Storage Manager server after you install the client. See “Upgrade path for clients and servers” on page 1 for more information about the changes.

Universal Binary support

You can now install and use Tivoli Storage Manager on an Intel-based Macintosh system. You can also do the following:

- Restore your current backup data from a PPC system to your Intel system with no loss of data or attributes.
- Restore your Intel backup data to a PPC system with no loss of data or attributes.

Changed defaults for the *managedservices* option

The *managedservices* option default has been changed to include *schedule*. By default the Client Acceptor Daemon will manage the Web client and schedules. A new option, *none*, has been added to support running *dsmc schedule*.

Option file and error log names

The Macintosh client has been changed to use the same option file and error log names as the other Tivoli Storage Manager Clients. The TSM User Preferences file is now the *dsm.opt* file and the TSM System Preferences file is now the *dsm.sys* file. During installation, old option files are renamed.

Support for the Web Interface

The Macintosh client now supports the Web Interface for remote back up and restore capability.

Support for *srprepoptscheddisabled* and *schedrestretrdisabled* options

The *srprepoptscheddisabled* and *schedrestretrdisabled* options are now supported, in addition to the existing *schedcmddisabled* option. These options help to prevent Tivoli Storage Manager administrators from executing inadvertent or malicious operations on Tivoli Storage Manager client nodes.

See “Srvprepoptscheddisabled” on page 246 and “Schedrestretrdisabled” on page 237 for information about these new options.

Reducing the memory used during the incremental backup function

The *memoryefficient diskcachemethod* option uses local disk space, which reduces the amount of memory required to complete an incremental backup. See “Memoryefficientbackup” on page 199 for more information about using *memoryefficient diskcachemethod*.

The *diskcachelocation* option specifies the location where the disk cache database will be created if the option *memoryefficientbackup=diskcachem* is set during an incremental backup. See “Diskcachelocation” on page 156 for information about how to use this option.

Support for *include.fs*

The *include.fs* option controls how Tivoli Storage Manager processes your file space for incremental backup. See “Include options” on page 186 for more information about using the *include.fs* option.

Exclude.fs option support

The *exclude.fs* option is now supported.

Dropped support for Mac OS 10.3

Beginning with this release of Tivoli Storage Manager, Mac OS 10.3 is no longer supported.

Technical changes for Version 5.3 - March 2006

The following changes have been made to the product for this edition:

Access Control Lists support

Beginning with Mac OS X 10.4.3, Access Control Lists are fully supported by the IBM Tivoli Storage Manager Version 5.3.3 Macintosh Client.

New options, *Skipacl* and *Skipaclupdatecheck*

Skipacl and *Skipaclupdatecheck* are new options for controlling the processing of access control lists (ACL) by Tivoli Storage Manager.

Skipacl=yes skips ACL processing completely on backup.

Skipaclupdatecheck=yes prevents the Tivoli Storage Manager client from performing checksum and size comparisons before and after backup, and during incremental processing. See “Skipacl” on page 244 and “Skipaclupdatecheck” on page 245 for more information.

Technical changes for Version 5.3 - October 2005

The following changes have been made to the product for this edition:

Changing HFS paths to UNIX paths

The Tivoli Storage Manager Macintosh Client now uses UNIX file paths for options, schedules, and file system names on the Tivoli Storage Manager server. By default, the UNIX mount point will be used to create file spaces on the server. See “Configuring the client scheduler” on page 21, “Migrating from earlier versions” on page 1, Chapter 9, “Using processing options,” on page 111, and “Using the Tivoli Storage Manager central scheduler,” on page 335 for more information about changing file paths from HFS to UNIX.

Note: There are required migration steps that must be completed before using Tivoli Storage Manager Version 5.3. This applies to all users migrating from any previous version of Tivoli Storage Manager. If these steps are not completed, Tivoli Storage Manager Schedules and Include/Exclude processing might not work as expected. See “Migrating from earlier versions” on page 1 for more information.

Extended Attribute support

The Tivoli Storage Manager Macintosh client can back up and restore Extended Attribute data. The Extended Attribute data is automatically backed up with each object that has Extended Attributes. See “Supported file systems” on page 53 for more information.

Event-based policy retention protection

In addition to traditional backup archive processing to a Tivoli Storage Manager server, the Tivoli Storage Manager Backup-Archive Clients can now archive data to an IBM System Storage Archive Manager server by using the *enablearchiveretentionprotection* option. When connected to any Tivoli Storage Manager server Version 5.2.2 and above, the Tivoli Storage Manager Backup-Archive Clients can take advantage of event-based archive management policies. With the new **set event** command the client can also signal events and request deletion hold and deletion release for specific files. This feature is intended to help customers meet regulatory requirements. See “Using event-based policy retention protection” on page 108, “Enablearchiveretentionprotection” on page 160 and “Set Event” on page 332 for more information.

New links from the backup-archive client GUI to the Tivoli Storage Manager and Tivoli Home Pages

From the backup-archive client GUI, you can now directly access the IBM Tivoli Storage Manager web site, the IBM Tivoli web site, the IBM Tivoli Storage Manager Support web site, and the Tivoli Storage Manager Publications web site.

From the backup-archive command-line interface help menu screen, the web site address is displayed. You can enter the web site address in a browser to access the IBM Tivoli Storage Manager Support web site for technical support on Tivoli Storage Manager.

New options, Errorlogmax and Schedlogmax, and DSM_LOG environment variable changes

Errorlogmax and *Schedlogmax* are new options for controlling log size. *Errorlogmax* specifies the maximum size for the error log, in megabytes.

Schedlogmax specifies the maximum size for the scheduler log, in megabytes. See Chapter 9, “Using processing options,” on page 111 for more information.

DSM_LOG environment variable changes have been made to prevent a security or data integrity problem. Logs will no longer be created in the installation directory. In addition, if the client is unable to open a required log for writing, the client process will terminate. The Tivoli Storage Manager client must have a writeable dsmerror.log in order to run. See Chapter 2, “Configuring Tivoli Storage Manager,” on page 13, Chapter 9, “Using processing options,” on page 111, and Chapter 10, “Using commands,” on page 271 for more information.

Enhanced encryption

Tivoli Storage Manager now supports AES (Advanced Encryption Standard) 128-bit data encryption to encrypt data during backup and archive operations using the *include.encryption* option.

Include-exclude enhancements

A preview function allows you to view the objects that would be backed up or archived according to the include-exclude list. This enhancement is useful for validating the include-exclude list.

The Tivoli Storage Manager Client GUI directory tree shows detailed information of included and excluded objects.

The directory tree in the Tivoli Storage Manager Client GUI allows you to select files and directories to include or exclude.

See Chapter 2, “Configuring Tivoli Storage Manager,” on page 13, Chapter 4, “Backing up your data,” on page 47, Chapter 6, “Archiving and retrieving your data,” on page 85, and Chapter 10, “Using commands,” on page 271 for more information.

Enhancements to query schedule command

To more accurately determine the status of scheduled events, the **query schedule** command on a Tivoli Storage Manager Version 5.3 and above client reports new parameters. See Chapter 7, “Automating tasks,” on page 93 and Chapter 10, “Using commands,” on page 271 for more information about these enhancements. Also see the Administrator’s Guide and Administrator’s Reference for your operating system.

Dynamic client tracing

A new command-line utility, **dsmtrace**, is available to enable tracing, disable tracing, or change trace flags while the client is running. See the Problem Determination Guide for information about using this utility.

Note: Tracing is an advanced diagnostic feature intended for use only at the recommendation of IBM support and development, or as outlined in the Problem Determination Guide.

Client-node proxy support

Backup time can be reduced and clustered configurations can store data with client-node proxy support. Tivoli Storage Manager nodes, when authorized as agent nodes, can be directed 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. The concurrent operations can reduce backup and restore times in environments such as Apple’s Xsan. Client nodes can also be configured with proxy node authority to support many of the systems

that can be configured to support clustering failover. The *asnodename* option also allows data to be restored from a different system than the one which performed the backup. See “Asnodename” on page 129 for more information.

Support for deleting individual backups from a server file space

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. This can be useful in the following situations:

- You need to delete sensitive files which were mistakenly backed up.
- You need to delete a subset of backup files which were inadvertently backed up.
- You need to delete files that are found to contain viruses.

See “Deleting backup data” on page 61 for more information.

Optimized option default values

Table 3. Optimized option default values

Option	Change	Page
<i>diskbuffsize</i>	Allows you to specify I/O buffer size (in kilobytes) that the client can use to optimize backup, archive, or HSM client performance. Recommendation: Use the default value instead of specifying the <i>diskbuffsize</i> option.	155
<i>largecommbuffers</i>	This option has been replaced by the <i>diskbuffsize</i> option. At this time, <i>largecommbuffers</i> will continue to be accepted by the Tivoli Storage Manager client in order to ease the transition to the new option. However, the value specified by <i>largecommbuffers</i> will be ignored in favor of the <i>diskbuffsize</i> setting. Recommendation: Discontinue the use of <i>largecommbuffers</i> because future releases of Tivoli Storage Manager might not accept this option.	155
<i>tcpadminport</i>	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. If this option is not specified, the default value is the value of the <i>tcpport</i> option.	251

Table 3. Optimized option default values (continued)

Option	Change	Page
<i>tcpbuffsize</i>	The default value was changed from 16 kilobytes to 32 kilobytes.	252
<i>tcpnodelay</i>	<p>The default value was changed from <i>no</i> to <i>yes</i>. <i>tcpnodelay yes</i> disables the TCP/IP Nagle algorithm. This algorithm is used to reduce the number of small segments sent across the network, but in some environments this might negatively impact Tivoli Storage Manager performance.</p> <p>Recommendation: Use the default of <i>yes</i>, unless you fully understand the effects of the TCP Nagle algorithm on network transmissions and how its use affects the performance of Tivoli Storage Manager in your environment.</p>	256
<i>tcpwindowsize</i>	The default value was changed from 32 kilobytes to 63 kilobytes.	259
<i>txnbytelimit</i>	The default value was changed from 2048 kilobytes to 25600 kilobytes.	263

Chapter 1. Installing Tivoli Storage Manager

The Tivoli Storage Manager backup-archive client helps you protect information on your workstation. Using Tivoli Storage Manager, you can maintain backup versions of your workstation files that you can restore if the original files are damaged or lost. You can also archive workstation files that you do not currently need, or preserve them in their current state, and retrieve them when necessary.

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

You can access Tivoli Storage Manager backup and archive features:

- Locally through the native Graphical User Interface. The main window displays buttons labeled: **Backup**, **Restore**, **Archive**, and **Retrieve**. To initiate a task, click the appropriate button.
- Locally or remotely through the client command-line interface.

This chapter provides migration, environment, and installation information for the Tivoli Storage Manager client. After installation, see Chapter 2, "Configuring Tivoli Storage Manager," on page 13 for required configuration tasks to perform before using Tivoli Storage Manager.

Migrating from earlier versions

The following sections explain what you need to do if you are migrating to Tivoli Storage Manager Version 5.5 from a previous version.

Upgrade path for clients and servers

As part of a migration plan from Tivoli Storage Manager Version 5.4 to Tivoli Storage Manager Version 5.5, 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 5.4 client can perform backup, restore, archive, and retrieve functions to a Tivoli Storage Manager Version 5.5 server.
- A Tivoli Storage Manager Version 5.5 client can perform backup, restore, archive, retrieve, and query functions to a Tivoli Storage Manager Version 5.4 server.
- If you back up or archive data from a Tivoli Storage Manager Version 5.3 or Version 5.4 client to any Tivoli Storage Manager server, you can restore or retrieve that data using a Tivoli Storage Manager Version 5.5 client.
- If you back up or archive data from a Tivoli Storage Manager Version 5.5 Macintosh client, you cannot restore or retrieve that data using a Tivoli Storage Manager Version 5.4 Macintosh client.
- All administrative clients can administer Tivoli Storage Manager Version 5.4 and Version 5.5 servers.
- If you use a Tivoli Storage Manager client prior to Version 5.5 to restore the Version 5.5 data, the file attributes will be set to zero. This only affects the client dependant attributes, such as date, time, and owners.
- If you are migrating from Tivoli Storage Manager Version 5.3, the format of the Tivoli Storage Manager password file has changed. Your existing password file will be migrated to the new format during the first session with a Tivoli Storage

Manager server after you install the Tivoli Storage Manager Version client. After the password file has been migrated, previous versions of the Tivoli Storage Manager client will not process the updated password file correctly. If you use Tivoli Storage Manager client data encryption and you ever need to go back to an older client version, follow these steps to avoid problems restoring your data with the updated password file and the older client:

1. Before installing the new version, make a copy of the password file. See “Passworddir” on page 209 for information about finding the location of the password file.
2. After installing the new version, run a backup and enter the encryption key password when prompted.
If data encryption is used, you will be prompted for the encryption key password on the first backup after the upgrade, to ensure that the correct encryption key is now used.
3. If you need to go back to an older version, follow these steps:
 - a. Uninstall the client.
 - b. Restore the saved passwords file.
 - c. Install the older version.

Migrating to the Tivoli Storage Manager Version 5.5 Macintosh client from releases prior to 5.3.2

The 5.3.2 release of the Tivoli Storage Manager Macintosh client had several changes that must be managed as part of the migration from previous Tivoli Storage Manager versions:

- Option file changes
- Scheduler changes
- Extended attribute changes

Option file changes

The Tivoli Storage Manager client now uses the UNIX[®] path names internally. This reduces the amount of path translation that must be done. However, this affects how paths are stored in the Tivoli Storage Manager option files and Tivoli Storage Manager Schedule definitions.

One set of options that takes a path are the include and exclude options. For example, consider the option to exclude the files in the 'Documents' folder for 'User1' on the startup volume 'Macintosh HD'. With the Tivoli Storage Manager Version 5.2.4 and earlier clients, this entry would be in the TSM System Preferences file:

```
exclude "Macintosh HD:Users:User1:Documents:*
```

Beginning with the Version 5.3.2 client, to exclude the files, the entry in the TSM System Preferences would be:

```
exclude /Users/User1/Documents/*
```

See “Using include-exclude options” on page 28 for more examples of how the include/exclude options must now be set.

All of the options listed in 2 have a path for a parameter. If any of these options are set in your current option files, they need to be converted to UNIX paths before the Tivoli Storage Manager Version 5.4 client can be used.

- DOMain

- ERRORLOGName
- INCLExcl
- MAILprog
- PASSWORDDIR
- All exclude and include options
- POSTSchedulecmd
- POSTNSchedulecmd
- PRESchedulecmd
- PRENSchedulecmd
- SCHEDLOGName

Follow these steps to migrate your option paths:

1. Rename your existing system option files (dsm.sys) and client user options files (dsm.opt).
2. Open the renamed files in a text editor so you can view your current option settings. The default location for the option files is:
/Library/Preferences/Tivoli Storage Manager
3. Start the Tivoli Storage Manager application. The setup assistant will start automatically.
4. Using your old option files as a reference, fill in the fields presented by the Setup Assistant. Mark each option in the old option files to keep track of the options that have been set.
5. After the Setup Assistant has finished, open the preference editor and set any options that were not available in the Setup Assistant.
6. Mark each option in the old option files to keep track of the options that have been set.
7. Quit the client so the options can take affect.
8. Review the marks in the old option files to see if any option has not been set with either the Setup wizard or Preference Editor. These options can be added to the option files with a text editor.

Scheduler changes

The scheduler now takes the UNIX path for the 'object' option. Any schedule that specifies a path must be changed to use the UNIX path. For example, this schedule has an object that must be updated:

```
Policy Domain Name: STANDARD
Schedule Name: INCR_MAC_HD
Description: INCR_MAC_HD
Action: Incremental
Options:
Objects: "Macintosh HD"
Priority: 5
Start Date/Time: 05/14/2005 15:21:32
Duration: 1 Hour(s)
Period: 1 Day(s)
Day of Week: Any
Expiration:
Last Update by (administrator): ADMIN
Last Update Date/Time: 05/14/2005 15:22:37
Managing profile:
```

This is the same schedule, with a new name, and an updated object:

```
Policy Domain Name: STANDARD
Schedule Name: INCR_UNIX_MAC_HD
Description: INCR_UNIX_MAC_HD
Action: Incremental
Options:
```

Objects: /
Priority: 5
Start Date/Time: 05/14/2005 15:21:32
Duration: 1 Hour(s)
Period: 1 Day(s)
Day of Week: Any
Expiration:
Last Update by (administrator): ADMIN
Last Update Date/Time: 05/14/2005 15:24:04
Managing profile:

Contact your Tivoli Storage Manager System Administrator and have them update any schedules that specify a path to use the UNIX versions of those paths.

Extended attributes changes

If you back up or archive data with a Tivoli Storage Manager Version 5.5 client to any Tivoli Storage Manager server and that data has Extended Attributes, the data cannot be restored by the Tivoli Storage Manager Version 5.2, or earlier, client.

Migrating to the Unicode-enabled client

Unicode is a universal character encoding standard that supports the interchange, processing, and display of text that is written in any of the languages of the modern world. Unicode-enabled file spaces provide support for multilingual workstations without regard for the current locale. Beginning with Tivoli Storage Manager Version 5.2, the Macintosh client is Unicode-enabled. The Unicode-enabled client supports Unicode for file, directory, and file space names.

Benefits of using the Unicode-enabled client

The following list includes some benefits of using the Unicode-enabled client.

- If you back up or archive objects with a Unicode-enabled client, you can restore or retrieve these objects with a Unicode-enabled client in any other supported language environment. For example, a Japanese Unicode-enabled client can restore files backed up by a German Unicode-enabled client.
- The Unicode-enabled client is cross language, meaning that it works with files using different code pages. The Unicode-enabled client can handle a mixture of files from different code pages or locales in the same file space.
- Without the support for storing Unicode-enabled file spaces, some earlier Macintosh clients have experienced backup failures. This occurs when file spaces contain names of directories or files in multiple languages or have names that cannot be converted to the client's current code page.

How to migrate to the Unicode-enabled client

New clients storing data on the server for the first time require no special set up. In this case, the server automatically stores data in Unicode-enabled file spaces. However, if your client already stores data on the server, you need to plan for the migration of your file spaces to Unicode-enabled file spaces. The migration involves renaming your file spaces on the server and creating new Unicode-enabled file spaces on the server using the *autofsrename* option.

The Tivoli Storage Manager administrator can specify three possible values using the server *autofsrename* option: yes, no, or client. Note that the server setting for this option overrides the setting on the client.

- If the server specifies no, the existing file spaces for your node are not renamed or Unicode-enabled, even if you have a Unicode-enabled client installed.

- If the server specifies *yes*, and the client performs an incremental backup with *domain=all-local*, all file spaces specified in the current operation that are not Unicode-enabled are renamed. The Unicode-enabled client then creates new Unicode-enabled file spaces.
If the client performs a selective backup, only the file spaces you specify in the current operation are renamed and Unicode-enabled.
- If the server specifies *client*, the client can set the *autofsrename* option in the client user options file (*dsm.opt*) to *prompt*, *yes*, or *no*. If the client sets the option to *no*, the file spaces that are not Unicode-enabled are not renamed and are not Unicode-enabled. If the client specifies *prompt*, the client prompts you whether to rename file spaces. The client prompts you only once to rename a particular file space.

Note: The Web GUI will not prompt you if the *autofsrename* option is set to *prompt*. The option will be ignored and the file space will remain non-Unicode.

Optionally, your Tivoli Storage Manager administrator can rename the file spaces on the server that are not Unicode-enabled.

For options and results when using the *autofsrename* option, see the table below.

Server option	Client option	Result
Yes	Prompt, Yes, No	Renamed
No	Prompt, Yes, No	Not Renamed
Client	No	Not Renamed
Client	Yes	Renamed
Client	Prompt	CLI: Based on User Response Web GUI: Not renamed

For detailed information about the *autofsrename* option, see “Autofsrename” on page 135. For additional information about working with the Macintosh clients, see:

- “Considerations for Unicode-enabled clients” on page 184
- “Incl excl” on page 184

Considerations for migrating to the Unicode-enabled client

When migrating to the Tivoli Storage Manager Macintosh Unicode-enabled client, consider the following guidelines:

- Ensure that you install the fonts for the languages you want to display. If you do not have the necessary fonts installed, the file names and archive descriptions might not display properly. However, the Unicode-enabled clients can still back up, restore, archive, and retrieve the files and directories.
- The name of the Unicode-enabled file space might be unreadable 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. On the GUI main window, select the **File details** option from the **View** menu to display the fsID of a file space.
- After installing a Unicode-enabled client, perform a full incremental backup to 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.

- Renamed file spaces remain on the server and are managed as stabilized file spaces. *These file spaces contain all the original data, which can still be restored until the file spaces are deleted.*
- 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, if necessary.
- 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 other than 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*, Tivoli Storage Manager will *expire* them in the file space that is not Unicode-enabled. However, the Unicode-enabled client can back up or archive these files to a Unicode-enabled file space.
- When migrating from a client that is not Unicode-enabled to a Unicode client, file names with double-byte characters mixed with single-byte characters might be restored in mixed single-byte character set (SBCS) format. This is only a display issue.
- When a client performs a selective backup of an object and the original file space is renamed, the new Unicode-enabled file space contains only that object. Tivoli Storage Manager will back up all other directories and file spaces during the next full incremental backup.
- The Unicode-enabled client has a file and directory name size limitation of 504 bytes. Depending on the UTF-8 encoding (which can be 2 bytes, 3 bytes, or 4 bytes) of the Unicode character, this translates to a maximum of 126 to 252 characters.

Attention: If you do not follow the migration instructions properly, you might have two file spaces, one Unicode and one non-Unicode, with different file space identifiers (fsID) for the same client volume. In this case, the Tivoli Storage Manager client uses the non-Unicode file space as the default file space for backup and restore operations.

Macintosh client environment

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

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/sysmgmt/products/support/IBMTivoliStorageManager.html>

Note: Due to differences between this and prior versions of Tivoli Storage Manager, the installation will remove prior versions of Tivoli Storage Manager. The prior installations that will be removed include the following:

- /Applications/Tivoli Storage Manager v5.1
- /Applications/Tivoli Storage Manager v5.2
- /Applications/Tivoli Storage Manager v5.2.2
- /Applications/Tivoli Storage Manager v5.2.3
- /Applications/Tivoli Storage Manager v5.2.4
- /Applications/Tivoli Storage Manager v5.2.5

If versions prior to Version 5.3.2 are installed in some other location, the following must be done before installation:

1. Remove or rename the existing *TSM System Preferences* and *TSM User Preference* files in */Library/Preferences/Tivoli Storage Manager*.
2. Delete the prior versions.

Client components

The following components are included with the Tivoli Storage Manager Client for Macintosh:

- Backup-Archive client (command line)
- Administrative client (command line)
- Backup-Archive Java™ Graphical User Interface (GUI)
- Web Backup-Archive client

For more information about the Administrative client, see “Related publications” on page xii for available Tivoli Storage Manager Administrator’s Reference publications.

Hardware requirements

- A PowerPC® or Intel® based Macintosh
- Memory: 512 MB
- Disk space: 100 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

- Mac OS X, version 10.5.x, where $x \geq 0$

Communications method

To use this communication method:	Install this software:	To connect to these Tivoli Storage Manager servers:
TCP/IP	Sockets for Mac X, which comes standard with the OS.	AIX®, HP-UX, Linux®, Solaris, Windows®, z/OS®

Installing the Macintosh client program (required system administrator task)

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

Installing from the Tivoli Storage Manager CD

To install the Macintosh client from the Tivoli Storage Manager CD, follow these steps:

1. Go to the *tsmcli/mac* directory, for the Macintosh client installer.
2. Double-click the *5.5.0.0-TIV-TSMBAC-Mac.dmg* file, to mount the disk image.
3. Double-click the “Tivoli Storage Manager” installation package icon.
4. Follow the on-screen prompts to complete the installation.

Installing from the command line

This 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 display

attached. It is done from the command line, can be scripted and requires no user interaction. To install Tivoli Storage Manager from the command line, follow these steps:

1. Change directories to where the Tivoli Storage Manager installer is located.
2. The 'custom' install package can be installed with 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 /
```

The packages listed in Table 4 are part of the Tivoli Storage Manager install distribution used by the Mac OS X Install application. They are listed here to assist those users that need to do a fully customized command-line installation.

Table 4. Macintosh installation packages

Package	Package Name	Package Description
Tivoli Storage Manager Backup-Archive Java Bundle	TsmBaJavaBundle.pkg	<p>Contains the following component:</p> <ul style="list-style-type: none"> • Tivoli Storage Manager Java GUI Client <p>NOTE: Requires that the following packages be installed first:</p> <ul style="list-style-type: none"> • 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 ICC Library • Tivoli Storage Manager Backup-Archive Help and Messages

Table 4. Macintosh installation packages (continued)

Package	Package Name	Package Description
Tivoli Storage Manager Backup-Archive Web Applet	TsmBaWebApplet.pkg	<p>Contains the following component:</p> <ul style="list-style-type: none"> • Tivoli Storage Manager Web Client <p>NOTE: Requires that the following packages be installed first:</p> <ul style="list-style-type: none"> • 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 ICC Library • Tivoli Storage Manager Backup-Archive Help and Messages
Tivoli Storage Manager Administrative Client	TsmBaDsmadmcc.pkg	<p>Contains the Tivoli Storage Manager Administrative Client</p> <p>NOTE: Requires that the following packages be installed first:</p> <ul style="list-style-type: none"> • Tivoli Storage Manager Backup-Archive Client Common Files • Tivoli Storage Manager Backup-Archive ICC 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 used by all of the Tivoli Storage Manager B/A client components.
Tivoli Storage Manager Backup-Archive Client Common Java GUI Files	TsmBaJavaCommon.pkg	Contains the files used by all of the Tivoli Storage Manager B/A client Java components.
Tivoli Storage Manager Backup-Archive Java GUI Files	TsmBaJavaGui.pkg	Contains the files used by the Tivoli Storage Manager B/A Java GUI component.

Table 4. Macintosh 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 used by all of the Tivoli Storage Manager B/A client components.
Tivoli Storage Manager Backup-Archive ICC Library	TsmBaIcc.pkg	Contains the Tivoli Storage Manager ICC Library used by all Tivoli Storage Manager B/A client components.

Uninstalling the Macintosh client

If the Tivoli Storage Manager Scheduler is configured as a Startup Item use TSM Tools for Administrators or StopCad.sh to stop and uninstall it before proceeding.

To assist with uninstallation, a shell script is provided. The shell script name is `uninstall.sh` and is located in the default installation directory: `/Library/Application Support/tivoli/tsm/client/ba/bin`. The `sudo` command must be used to run the script. The following steps can be done instead of using the script.

Move the following folders to the trash:

1. `/Applications/Tivoli Storage Manager`
2. `/Library/Application Support/tivoli`

Remove the following symbolic links:

1. `/usr/bin/dsmc`
2. `/usr/bin/dsmtca`
3. `/usr/bin/dsmcad`
4. `/usr/bin/dsmadm`
5. `/usr/bin/dsmtrace`
6. `/usr/bin/dsmagent`
7. `/usr/lib/libxmlutil-5.5.0.0.dylib`
8. `/usr/lib/libtsm550xerces-c1_6_0.dylib`

Move the following receipts to the trash:

1. `/Library/Receipts/TsmBaClientCommon.pkg`
2. `/Library/Receipts/TsmBaDsmadm.pkg`
3. `/Library/Receipts/TsmBaDsmc.pkg`
4. `/Library/Receipts/TsmBaHelpCat.pkg`
5. `/Library/Receipts/TsmBaIcc.pkg`
6. `/Library/Receipts/TsmBaJavaBundle.pkg`
7. `/Library/Receipts/TsmBaJavaCommon.pkg`
8. `/Library/Receipts/TsmBaJavaGui.pkg`
9. `/Library/Receipts/TsmBaSmp.pkg`
10. `/Library/Receipts/TsmBaWebApplet.pkg`

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, option files, and receipts is optional.

Enabling non-administrators to manage their own data (required system administrator task)

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:

- Set the *passwordaccess* option to *generate* in the system options file (**dsm.sys**).
- The system administrator must generate the initial password by using Tivoli Storage Manager one time.

Chapter 2. Configuring Tivoli Storage Manager

After installation, *required* configuration tasks include the following:

Task	Page
Creating and modifying the client system options file (required root user or authorized user task)	15
Creating a default client user options file (optional root user or authorized user task)	16
Registering your workstation with a server (required root user or authorized user task)	25

Optional configuration tasks include the following:

Task	Page
Setting processing environment variables (optional)	19
Creating an include-exclude list (optional root user or authorized user task)	26

Root user, authorized user, and non-authorized user tasks

Throughout this publication, the phrases **root user** and **authorized user** identify Tivoli Storage Manager tasks that either a system administrator, non-administrator, or both can perform. Tivoli Storage Manager commands and options that you can use to perform these tasks are also identified in this way.

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. Table 5 lists the authorization tools to use.

Table 5. 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

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 will have an account type of **Admin**.

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 and whose owner execution permission bit is set to s. In the following example, the user `tivoli` is an authorized user while running `dsmc` since the `dsmc` owner execution permission bit is set to s:

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

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.
- Have a user account that has not been configured as a system administrator.

See “Enabling non-administrators to manage their own data (required system administrator task)” on page 11 for more information.

Table 6 shows the tasks that can be performed by the root user, authorized user, and the non-authorized user.

Table 6. Root user, authorized user, and non-authorized user tasks

Task	Root User	Authorized User	Non-Authorized User
Registering new nodes with Tivoli Storage Manager server.	Yes.	Yes.	No, even when the registration is set to <i>open</i> 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 a client user options file.	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 will prevent writing to the same location if the file has read only permission. Ownership of all restored objects will be changed to authorized user.	Yes, if you own the file or you are granted access. The ownership of the restored file will change 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 will prevent writing to the same location if the file has read only permission. Ownership of all retrieved objects will be changed to authorized user.	Yes, if you archived the file. Ownership of all retrieved objects will be changed to the non-authorized user.
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.

Using TSM Tools for Administrators

TSM Tools for Administrators is an Automator application that enables System Administrators to run Tivoli Storage Manager as a root process, start the Client Acceptor Daemon, and stop the Client Acceptor Daemon. All three of these tasks require a System Administrator account and password.

Follow these steps to use TSM Tools for Administrators:

1. Double-click the icon.
2. Select the application to run.
3. Press the **OK** button.
You will be prompted to authenticate with a System Administrator name and password.
4. Press the **OK** button and the selected application will start.

Note: While the selected application is running, you will see 'Run AppleScript' on the menu bar.

Creating and modifying the client system options file (required root user or authorized user task)

During installation, a sample system option file (dsm.sys.smp) is placed in the installation directory.

If you are a **root user** or authorized user, copy the file to /Library/Preferences/Tivoli Storage Manager/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.

Attention: 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:

Servename	server_a
COMMMethod	TCPip
TCPPort	1500
TCPServeraddress	node.domain.company.com

Note: If you want to use the Web client, you must also specify the *passwordaccess=generate* option. See "Passwordaccess" on page 207 for more information.

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

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

Note: See Chapter 9, “Using processing options,” on page 111 for more information about these 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.

Note: 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.

See “Setting options in an options file” on page 18 for information on how to set options in the `dsm.sys` file.

Creating a default client user options file (optional root user or authorized user task)

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

This file contains 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

See Chapter 9, “Using processing options,” on page 111 for more information about these 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 command line:

- Change to the directory containing the sample file.
- Copy the `dsm.opt.smp` file to `dsm.opt` or to a new file name of your choice.
 - **For the default client user options file:** You can store your default client user options file as `/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 can set the DSM_CONFIG environment variable or use the optfile option to point to your new client user options file.

- **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 to which you have write access.

Recommendation: Use full path names instead of relative path names when you set environment variables.

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 **Client 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.

Note: 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.

The Preferences editor uses the environment variable DSM_DIR or the default location to locate the dsm.sys file, and DSM_CONFIG or the default location 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.

Creating a customized client user options file (optional user task)

If you are a user and 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 the options that can be set in the default user options file.

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.
2. Copy dsm.opt.smp to ~/Library/preferences/Tivoli Storage Manager 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. If your file name is *not* dsm.opt, set the DSM_CONFIG environment variable to point to your new client user options file. For instructions to set this variable, see section, “Setting processing environment variables (optional)” on page 19.
4. Edit your dsm.opt file as appropriate for your system or use the Tivoli Storage Manager Preferences editor by selecting **Edit** → **Client Preferences** from the Tivoli Storage Manager GUI.

See “Setting options in an options file” on page 18 for more information about setting options in a file.

Note: 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 client user option file as a plain text file. Do not add the .txt extension.

Setting options in an options file

This section describes how to set options in your dsm.sys file or client user options file (dsm.opt).

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

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

To set an option in these files, 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:

```
Domain / /Volumes/fs2 /Volumes/fs2 /Volumes/fs3 /Volumes/fs4
```
- Enter one or more blank spaces between parameters.
- The maximum number of characters in a file name is 255. The maximum combined length of the file name and path name is 1024 characters. The maximum length of a file name is limited to 504 bytes (not characters). Because the Unicode representation of a character can occupy several bytes, 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. See “Query Options” on page 309 for more information.

Note: The client option files must be a plain text file. By default, TextEdit does not save files as plain text. Select **Format** → **Make PlainText** to save the client user options file as a plain text file. When saving, do not add the .txt extension.

Setting processing environment variables (optional)

Generally, setting the environment variables is an optional task. Setting them will make it more convenient for you to use the command line. However, you must set the environment variables in either of the following circumstances:

1. You want to invoke Tivoli Storage Manager from a directory other than the directory where Tivoli Storage Manager is installed.
2. You want to specify a different user preferences file for the backup-archive client, the administrative client, or both.

Note: You can also specify an alternate user preferences file for the command-line client (not the administrative client) using the *optfile* option. See “Optfile” on page 205 for more information.

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*, and the resource files reside. You *cannot* specify the root (/) directory for *DSM_DIR*. If *DSM_DIR* is not set, the files are expected in the default installation directory.

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*.

Note: The default installation directory is: `/Library/Application Support/tivoli/tsm/client/ba/bin`.

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.

DSM_CONFIG *cannot* be used to specify an options file in the root (/) directory for. The file must reside in 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 both of these requirements:

1. The options file must be named *dsm.opt*.
2. The file must reside in either `~/Library Preferences/Tivoli Storage Manager` or `/Library Preferences/Tivoli Storage Manager`.

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.

If you do not define *DSM_LOG*, Tivoli Storage Manager writes messages to the *dsmerror.log* and *dsmsched.log* files.

Notes:

1. The *errorlogname* and *schedlogname* options override *DSM_LOG*.
2. The log files cannot be symbolic links. Tivoli Storage Manager will detect any such links, delete the links, then exit the operation. This prevents Tivoli Storage Manager from overwriting protected data. The affected logs will be created as files in a subsequent operation.

Setting T and C shell variables

For the tcsh shell, enter the environment variables in the `.tcshrc` file for T shell and `.cshrc` for C shell in your user's \$HOME directory. For example:

```
setenv DSM_DIR "/Library/ApplicationSupport/tivoli/tsm/client/ba/bin"
setenv DSM_CONFIG "/Users/user1/dsm.opt"
setenv DSM_LOG /Users/user1/logs
```

Where:

- The `/Library/Application Support/tivoli/tsm/client/ba/bin` directory is where the executable file and resource files reside.
- `/Users/user1/dsm.opt` is the path and file name for your `dsm.opt` file.
- The `/Users/user1/logs` directory is where the `dsmerror.log` and `dsmsched.log` files reside.

Setting Bash and Bourne shell variables

For the Bourne shell, add the `DSM_CONFIG`, `DSM_LOG` and `DSM_DIR` variables to the `.profile` file in your user's \$HOME directory. For example:

```
export DSM_DIR="/Library/Application Support/tivoli/tsm/client/ba/bin"
export DSM_CONFIG="/Users/user1/dsm.opt"
export DSM_LOG="/Users/user1/logs"
```

Where:

- The `/Library/Application Support/tivoli/tsm/client/ba/bin` directory is where the executable file and resource files reside.
- `/Users/user1/dsm.opt` is the path and file name for your `dsm.opt` file.
- The `/Users/user1/logs` directory is where the `dsmerror.log` and `dsmsched.log` files reside.

Configuring the Web client

To configure the Web client, 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. For more information on *passwordaccess*, see "Passwordaccess" on page 207.
3. To generate the Tivoli Storage Manager password, start the backup-archive client by entering:

```
dsmc query session
```

when prompted, enter your user ID and password.

Note: A password can also be generated by using "TSM Tools for Administrators" and starting Tivoli Storage Manager.

4. Start the client acceptor daemon by entering:

```
dsmcad
```

Note: The client acceptor can also be started with "TSM Tools for Administrators" and selecting "Start the TSM Connect Agent"

Port 1581 is the default port number. You can set a different port number using the *httpport* option. See "Httpport" on page 181 for more information.

After installing and configuring the Web client on your workstation you can use the Web client to perform backup, archive, restore, and retrieve operations from any browser with JRE (Java Runtime Environment) 1.4.1 or higher. See “Starting a Web client session” on page 42 for more information.

Configuring 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:

- Display information about available schedules.
- Display information about work that the schedule has completed.
- Modify scheduling options in the dsm.sys file. See “Scheduling options” on page 118 for more information.

See Chapter 7, “Automating tasks,” on page 93 for more information. See “Using the Tivoli Storage Manager central scheduler,” on page 335 for supplemental information about the Tivoli Storage Manager central scheduler.

The Tivoli Storage Manager client acceptor daemon can manage the scheduler. In this case, 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. This reduces the number of background processes on your workstation and resolves memory retention problems that can occur when running the scheduler service without client acceptor daemon management. It is recommended that you use the Client Acceptor daemon to manage the client scheduler.

Use the *managedservices* option in your client system options file (dsm.sys) to specify whether the client acceptor daemon manages the scheduler. See “Managedservices” on page 196 for more information.

Perform the following steps to configure the client acceptor daemon to manage the client scheduler:

1. Install the Scheduler. See “Starting the client scheduler” on page 22 for more information.
2. From the Tivoli Storage Manager GUI, select **Edit** → **Client Preferences**. Then select the **Web Client** category. Check the **Schedule** option in the **ManagedServices** options section. If you wish to run the Web client also, check the **Both** option.
3. Start the Client Acceptor. See “Configuring the Web client” on page 20 for more information.

Notes:

1. If you include files for encryption, ensure that *ENCRYPTKEY SAVE* is set in the system options (dsm.sys) 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 the encryption prompt will be given and the key can be saved.

2. For more information about scheduling options, changing the scheduling mode, specifying the TCP/IP address or port number, or running commands before or after a schedule, see “Scheduling options” on page 118.
3. See Chapter 7, “Automating tasks,” on page 93 for information about the following tasks:
 - Modifying scheduling options in the client options file.
 - Displaying information about available schedules.
 - Displaying information about work that the schedule has completed.

To start the client scheduler on your client node and connect to the server schedule, change to the Tivoli Storage Manager installation directory and enter the following command:

```
dsmc schedule
```

Starting the client scheduler

You can run scheduling services by using the command-line client.

Using the command-line client to handle scheduled events

System Administrator

You can configure Tivoli Storage Manager to use the command-line client interface to handle scheduled events.

Before starting the client scheduler, you must complete the following steps before the client acceptor daemon will work.

1. Ensure that the *managedservices* option includes *schedule* in the client systems options (dsm.sys) file. See “Managedservices” on page 196 for more information.
2. Set the *passwordaccess* option to *generate* in the client systems options (dsm.sys) file. See “Passwordaccess” on page 207 for more information.

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 will be given and the key can be saved.

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 during system boot up.

Method 1 (recommended)

- Use “TSM Tools for Administrators” and “Start the TSM Connect Agent”. This will install the client acceptor daemon as a system startup item so

the client acceptor daemon will start after system restarts. The client acceptor daemon is also started immediately, so a reboot is not required to handle scheduled events.

Method 2

- In `/Library/Application Support/tivoli/tsm/client/ba/bin` is a shell script that you can use to install the client acceptor daemon as a Startup Item. The script name is `install_cad.sh`. See this script for more details.

Notes:

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

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

2. To check if 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 will have the path `/usr/bin/dsmcad`.

Defining a schedule

To define a schedule for a Macintosh node that is using any release prior to Tivoli Storage Manager Version 5.3.2, the schedule object must be specified in HFS format. For example, to back up the Macintosh HD volume, the schedule object would be `Macintosh HD:`. To back up the Users folder on the Macintosh HD volume, the schedule object would be `Macintosh HD:Users:`.

Beginning with the Tivoli Storage Manager Version 5.3.2 Macintosh Client, the schedule object must be specified with the UNIX path. For example, assume that the Macintosh HD is the startup disk for the system. The correct schedule object for the Macintosh HD volume is `/`. To back up the Users folder on the Macintosh HD volume, the schedule object would be `/Users/`.

If Macintosh HD is *not* the startup disk, the correct path is likely to be `/Volumes/Macintosh HD`. The correct schedule object for the Macintosh HD volume is `"/Volumes/Macintosh HD/"`. To back up the Users folder on the Macintosh HD volume, the schedule object would be `"/Volumes/Macintosh HD/Users"`.

Notes:

1. If the system only has one volume, then `/` is the correct UNIX path to use.
2. Most secondary file system paths start with `/Volumes/<name of volume>`.
3. The `df` command will show the correct UNIX paths for all volumes.

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. Because every firewall is different, 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. See “Tcpport” on page 257 for more information. This will allow Tivoli Storage Manager scheduler communications in both *polling* and *prompted* mode, client acceptor daemon-managed schedulers, and regular backup-archive client operations.

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

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. See “Webports” on page 269 for more information about the *webports* option.

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. See “Tcadminport” on page 251 for more information.

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 will 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.

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.

Notes:

1. Using the *sessioninitiation* option requires a Tivoli Storage Manager Version 5.2 or higher server and client.
2. 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 will still attempt to initiate sessions, but are blocked by the Tivoli Storage Manager server for nodes which have the *sessioninitiation* option set to *serveronly*.
3. When installing Tivoli Storage Manager scheduler using the setup wizard, and the Tivoli Storage Manager server is behind a firewall, the node password will not get stored on the client system. As a result, the scheduler service be unable to authenticate to the server when the server contacts the client to execute a schedule. In this case, you can

run the scheduler from the command line (`dsmc schedule`), wait until a scheduled operation starts, 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 run the scheduler from the command line (`dsmc schedule`), wait until a scheduled backup starts, and enter the encryption key when prompted. After the password and encryption key are updated, you must restart the scheduler.

4. When configuring the Tivoli Storage Manager scheduler on a client system for the first time, the scheduler service might be unable to authenticate to the server when the server contacts client scheduler to execute 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 scheduler from the command line (`dsmc schedule`), wait until a scheduled operation starts, and enter the password for your node when prompted.

Note: 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 system to the server system. See 23 for more information about setting up this communication. As soon as 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 will initiate 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 be used to prompt the client to connect to the server.

See “Sessioninitiation” on page 242 for more information about the *sessioninitiation* option.

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 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.
- The server cannot log events to a Tivoli Enterprise Console® server across a firewall.

Registering your workstation with a server (required root user or authorized user task)

Root user or authorized user

Before you can use Tivoli Storage Manager, 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.

Using 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 the following information to your Tivoli Storage Manager administrator:

- Your node name (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.

In addition to possibly defining certain options in your options file, 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.

Using 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.

Note: You can only do Open Registration through the command-line interface.

Creating an include-exclude list (optional root user or authorized user task)

System Administrator

This is an optional task but an important one. 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. For information on management classes and policy domains, see Chapter 8, "Understanding storage management policies," on page 99.

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 folders 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.

Attention: There are some system files that you should exclude. See “Excluding system files” on page 30 for more information.

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.

See “Inclexcl” on page 184 for important detailed information about specifying an include-exclude file using the *inclexcl* option.

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.

Note: You can specify a Unicode file as an include-exclude file using the Preferences editor. However, you cannot create the Unicode file using the Preferences editor. See “Considerations for Unicode-enabled clients” on page 184 for instructions on creating a Unicode include-exclude file.

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.

Note: Each server stanza must have its own include-exclude list.

3. Enter your *include* and *exclude* statements using the appropriate include-exclude options as described in “Using include-exclude options” on page 28. Tivoli Storage Manager evaluates all *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.cp
```

4. Save the file and close it. 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.

Using include-exclude options

This section provides the following information:

- Brief descriptions of the *include* and *exclude* options that you can specify in your dsm.sys file. See table references for more information about each option.
- A minimum include-exclude list that excludes system files.
- A list of supported wildcard characters.
- Examples of how you might use wildcard characters with *include* and *exclude* patterns.

Excluding directories

Use *exclude.fs* and *exclude.dir* statements to exclude all files and sub-directories in the specified directory from processing. 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 directories and files from the list of objects available for processing. The *exclude.fs* and *exclude.dir* statements override all include statements that match the pattern.

Table 7. Option for excluding directories

Option	Description	Page
<i>exclude.fs</i>	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. See “Exclude options” on page 168 for more information.	168

Table 7. Option for excluding directories (continued)

Option	Description	Page
<i>exclude.dir</i>	Excludes a directory, its files, and all its subdirectories and their files from backup processing. For example, the statement <i>exclude.dir /test/dan/data1</i> excludes the <i>/test/dan/data1</i> directory, its files, and all its subdirectories and their files. Using the <i>exclude.dir</i> option is preferable over the standard <i>exclude</i> option to exclude large directories containing many files that you do not want to back up. You cannot use <i>include</i> options to override an <i>exclude.dir</i> statement. Only use <i>exclude.dir</i> when excluding an entire directory branch.	168
Notes:		
1. If you define an exclude statement without using a drive letter, such as <i>exclude.dir dirname</i> , this will exclude from processing any directory named <i>dirname</i> on any drive.		
2. Use the following statements to exclude volumes <i>/Volumes/disk2</i> altogether from backup processing. Note that the volume (<i>/Volumes/disk2</i>) will still be backed up, but all other directories on <i>/Volumes/disk2</i> will be excluded.		
<pre>exclude /Volumes/disk2/* exlcude.dir /Volumes/disk2/*</pre>		
3. An alternative method for excluding an entire volume from domain incremental backup is to use a domain statement to exclude the volume. For example:		
<pre>domain "-/Volumes/disk2"</pre>		
This alternative still permits selective backup processing of files on <i>/Volumes/disk2</i> .		

Controlling backup and archive processing

After Tivoli Storage Manager evaluates all statements, the following options are evaluated against the remaining list of objects available for processing.

Table 8. Options for controlling backup and archive processing

Option	Description	Page
Backup processing		
<i>exclude</i> <i>exclude.backup</i> <i>exclude.file</i> <i>exclude.file.backup</i>	<i>These options are equivalent.</i> Use these options to exclude a file or group of files from backup services.	168
<i>include</i> <i>include.backup</i> <i>include.file</i>	Use these options to include files or assign management classes for backup processing.	186
<i>include.fs</i>	Controls how Tivoli Storage Manager processes your file space for incremental backups.	186
Archive processing		
<i>exclude.archive</i>	Excludes a file or group of files from archive services.	168
<i>include</i> <i>include.archive</i>	<i>These options are equivalent.</i> Use these options to include files or assign management classes for archive processing.	186

Excluding system files

The Tivoli Storage Manager client will add 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 will occur.

```
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
```

Notes:

1. Volumes with ... in their name are not supported. Tivoli Storage Manager uses the ... sequence as part of include-exclude processing. Tivoli Storage Manager reports an invalid include-exclude statement if a volume has ... in the name. The volume *must* be renamed.
2. Objects that have a type of *rhap* and a creator of *lcmt* will be excluded from processing. Generally, these are special file system objects, but can also be created with the **mknod** command or are UNIX mount points. The objects or mount points will need to be manually recreated as part of a full system restore.

Recommendation: You should have the following minimum include-exclude list in your include-exclude options file:

```

EXCLUDE ".../Norton FileSaver Data"
EXCLUDE ".../Norton VolumeSaver Data"
EXCLUDE ".../Norton VolumeSaver Index"
EXCLUDE ".../TSM Sched*"
EXCLUDE ".../TSM Error*"
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 /.../.Trash
EXCLUDE.DIR /.../.Spotlight-*
EXCLUDE.DIR /.../.TemporaryItems
EXCLUDE.DIR /.../Library/Caches
EXCLUDE.DIR /.../.fsevents
EXCLUDE.DIR /.../Backups.backupdb
EXCLUDE.DIR /Trash
EXCLUDE.DIR /TheVolumeSettingsFolder
EXCLUDE.DIR "/Temporary Items"
EXCLUDE.DIR /Users/.../Library/Caches
EXCLUDE.DIR "/System Folder/Preferences/netscape A/Cache"
EXCLUDE.DIR "/System Folder/Preferences/Explorer/Temporary Files"
EXCLUDE.DIR "/System Folder/Preferences/cache-cache"
EXCLUDE.DIR "/System Folder/Preferences/netscape users/.../Cache"

```

These are system files that cannot be recovered without possibly corrupting the operating system, or temporary files with data that you can easily recreate.

Including and excluding groups of files

To specify groups of files that you want to include or exclude, use the wildcard characters listed in Table 9. This table applies to *include* and *exclude* statements *only*. For information about using wildcard characters in Tivoli Storage Manager commands, see “Using wildcard characters” on page 277.

Note: 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 9. Wildcard and other special characters

Character	Function
?	<p>The match one character matches any single character <i>except</i> the directory separator; it does not match the end of the string. For example:</p> <ul style="list-style-type: none"> 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 abkkrs.
*	<p>The match-all character. For example:</p> <ul style="list-style-type: none"> 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.

Table 9. Wildcard and other special characters (continued)

Character	Function
/...	The match- <i>n</i> 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.

Examples using wildcards with include and exclude patterns

Table 10 contains examples of ways you might use wildcard characters with *include* and *exclude* patterns.

Table 10. Using wildcard characters with include and exclude patterns

Task	Pattern
Exclude all files that end with .doc, except those found on volume La Pomme in the Documents folder.	EXCLUDE /.../*.doc INCLUDE "/Volumes/La Pomme/Documents/*.doc
Exclude all files during backup with an extension of bak, except those found on the in the dev directory.	exclude /.../*.bak include /usr/dev/*.bak
Exclude all files and folders under any Documents folder that might exist, <i>except</i> for the Current file in the Documents folder on the La Pomme volume. Include this file. Assume that La Pomme is not the startup disk.	EXCLUDE /.../Documents/.../* INCLUDE "/Volumes/La Pomme/Documents/Current"
Exclude all files and directories in any tmp directory that might exist, <i>except</i> for the file . Include this file.	exclude /.../tmp/.../* include /home/tmp/save.fil
Exclude any .cpp file in any folder on the Vol1, Vol2, Vol3, and Vol4 volumes.	EXCLUDE /Volumes/Vol[1-4]/.../*.cpp
Exclude any .cpp file in any folder on the Vol1, Vol2, Vol3, and Vol4 volumes.	EXCLUDE /Volumes/Vol[1-4]/.../*.cpp
Exclude any .o file in any directory on the /usr1, /usr2, and /usr3 file systems.	exclude /usr[1-3]/.../*.o
Exclude the .obj files found in the root directory in the <i>only</i> .	
Exclude any file that resides under the tmp directory found .	exclude /.../tmp/.../*
Exclude the entire directory structure /var/spool from all processing.	exclude.dir /var/spool
Exclude a single file system from backup processing.	exclude.fs home:
Exclude all file systems mounted anywhere in the /Volumes/svt1 and /Volumes/svt2 directory tree from backup processing.	exclude.fs /Volumes/svt1/* exclude.fs /Volumes/svt2/.../*
Exclude the /home/mydir/test1 directory and any files and subdirectories under it.	exclude.dir /home/mydir/test1

Table 10. Using wildcard characters with include and exclude patterns (continued)

Task	Pattern
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 all symbolic links or aliases from backup processing, except for the Docs folder for user1. Assume that Jaguar is the startup disk.	EXCLUDE.ATTRIBUTE.SYMLINK /.../* INCLUDE.ATTRIBUTE.SYMLINK /Users/user1/Docs/*

Controlling symbolic link and alias processing

After Tivoli Storage Manager evaluates all *exclude.fs* and *exclude.dir* statements and removes the excluded directories, 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. Table 11 defines options for controlling symbolic link and alias processing.

Table 11. Options for controlling symbolic link and alias processing

Option	Description	Page
<i>exclude.attribute.symlink</i>	Excludes a file or a group of files that are symbolic links or aliases from backup processing only.	168
<i>include.attribute.symlink</i>	Includes a file or a group of files that are symbolic links or aliases within broad group of excluded files for backup processing only.	186

Controlling compression and encryption processing

After Tivoli Storage Manager evaluates *exclude.dir* and any other include-exclude options controlling symbolic links, backup, and archive processing, it uses the following options to determine which files undergo compression and encryption processing.

Table 12. Options for controlling compression and encryption processing

Option	Description	Page
Compression processing		
<i>exclude.compression</i>	Excludes files from compression processing if <i>compression=yes</i> is specified. This option applies to backups and archives.	168
<i>include.compression</i>	Includes files for compression processing if <i>compression=yes</i> is specified. This option applies to backups and archives.	186
Encryption processing		
<i>exclude.encrypt</i>	Excludes files from encryption processing.	168

Table 12. Options for controlling compression and encryption processing (continued)

Option	Description	Page
<i>include.encrypt</i>	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. Attention: The <i>include.encrypt</i> option is the only way to enable encryption on the Backup-Archive client. If no <i>include.encrypt</i> statements are used encryption will not occur.	186

Testing an include-exclude list with the Preview command

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.

Assume that you want to back up the files on your /Users/home file space. You 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. You scroll down the tree and notice that all of the *.o files in your /Volumes/home/mary/myobjdir will be backed up. You don't want to back up these files, so you right click a .o file, and choose "View File Details" from the popup menu. The dialog shows that these files are included. click the "Advanced" button and create a rule to exclude all .o files from the /home file space. 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. 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.

See "Preview" on page 297 for more information about the **Preview** command.

Processing include and exclude options

The Tivoli Storage Manager server can define include-exclude options using the *incl excl* parameter in a client option set. The include-exclude statements specified by the server are evaluated along with those in the dsm.sys 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 *incl excl* options that specify include-exclude files, the include-exclude statements in these files are placed in the list position occupied by the *incl excl* option and processed accordingly.

When performing an incremental backup, Tivoli Storage Manager evaluates all *exclude.dir* statements *first*, and removes the excluded directories and files from the list of objects available for processing. See "Excluding directories" on page 28 and "Exclude options" on page 168 for more information about the *exclude.dir* option.

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. See Chapter 9, "Using processing options," on page 111 for more information about the order in which all options are processed.

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. See "Query Inclexcl" on page 306 for more information.

The client program processes the list of include-exclude statements according to the following rules:

1. Folders *are not* checked; files are checked. Therefore, if you have a volume (such as "/Volumes/La Pomme") defined using the *domain* option, all folders of that volume are backed up *regardless* of what you defined for your *include* and *exclude* options. This is the case *even if* the files in those folders are not backed up.
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: 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.

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 current file being processed is /Volumes/La Pomme/Foo/Dev/test.cpp. The 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 current file name being processed.
2. Processing moves to Rule 2 and checks. This time, pattern /Volumes/La Pomme/Foo/.../*.cpp matches the current file name being processed. Processing stops, the option is checked, and it is *include*.
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 current file being processed is /Volumes/La Pomme/Widget/Sample File.
The 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 Volumes/La Pomme/Widget/Sample File is implicitly included and is backed up.

Example 3

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 current file being processed is /La Pomme/Lib/Src/Module1.cpp. The 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 *exclude*.
5. File /La Pomme/Lib/Src/Module1.cpp is not backed up.

Example 4

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 current file being processed is: /Users/user1/documents/printf.cpp.
Processing follows these steps:

1. The *exclude.attribute.symlink* statement is checked first. If the printf.cpp file is a symbolic link or alias it will be 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 *exclude*.
6. File /Users/user1/documents/printf.cpp is not backed up.

Example 5

Assume that La Pomme is the startup disk, is defined as the domain, and that you defined the following statements for the include and exclude options.

```
EXCLUDE ../../*.cpp
INCLUDE /Foo/../../*.cpp
EXCLUDE /Foo/Junk/*.cpp
```

The current file being processed is: /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 `/Foo/Dev/test.cpp` does not match the current file name being processed.
2. Processing moves to Rule 2 and checks. This time, pattern `/Foo/.../*.cpp` matches the current file name being processed. Processing stops, the option is checked, and it is included.
3. File `/Foo/Dev/test.cpp` is backed up.

Chapter 3. Getting started

This chapter includes instructions for the following tasks:

Task	Page
Starting a GUI session	40
Starting a command-line session	41
Starting a Web client session	42
Starting the client scheduler automatically	43
Changing your password	43
Sorting file lists using the Tivoli Storage Manager Java GUI	44
Displaying online help	45
Ending a session	45

Tivoli Storage Manager client authentication

When using the backup-archive GUI, command-line client, or the Web client, you can log on using a node name and password *or* administrative user ID and password. Tivoli Storage Manager prompts for your user ID and compares it to the configured node name. If they match, Tivoli Storage Manager 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 user options file (dsm.opt). See “Revokeremoteaccess” on page 229 for more information.

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.

Starting a GUI session

Start a GUI session using one of the following methods:

- Double-click the Tivoli Storage Manager application. This will start Tivoli Storage Manager without System Administrator privileges and allow you to 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 will start with System Administrator privileges. This allows access to files owned by all users on the system.
- Use the following command to run the Java GUI in the foreground:

```
dsmj
```

- Use the following command to run the Java GUI in the background:

```
dsmj
```

Tivoli Storage Manager locates the *client options* file (dsm.opt) and starts with the options specified in that file. See Chapter 2, “Configuring Tivoli Storage Manager,” on page 13 for more information about the client options file.

Note: dsmj is installed in the following location: /Library/Application Support/tivoli/tsm/client/ba/bin.

Using your Tivoli Storage Manager 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. For information about changing your password, see “Changing your password” on page 43.

Setup wizard

When the GUI client starts, it checks to see whether a client system options file exists. If the client system options file does not exist (which usually happens after you have installed the client for the first time on your system), the setup wizard will automatically start and guide you through the configuration process.

Starting a command-line session

You can start a command-line session using one of the following methods.

Note: If the 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.

- 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, to execute a single command (*batch mode*). If the symbolic link `/usr/bin/dsmc` exists, you can enter the **dsmc** command from any directory. Otherwise you can enter the fully qualified name.
- On the command line enter **dsmc**. This places you in *interactive mode*, permitting you to run several commands without preceding each with **dsmc**.

Your Tivoli Storage Manager administrator can require you to use a password to connect to the server.

You can start a client command session in either batch or interactive mode.

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 archive the file `/Users/user1/Documents/file1.txt`, enter the command:

```
dsmc archive /Users/user1/Documents/file1.txt
```

To issue the **incremental** command, enter the following at the prompt:

```
dsmc incremental
```

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. See “Passwordaccess” on page 207 for more information.

When you type your password and press Return, the password is not displayed on your screen.

Using interactive mode

Use *interactive* mode when you want to issue a series of commands. Because Tivoli Storage Manager establishes the connection to the server only once for interactive mode, 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**

When you press **Return**, this 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 named `/Users/user1/file1.txt`, enter the command:
`archive /Users/user1/file1.txt`

Depending upon the current setting of your *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 type your password and press Return, the password is not displayed on your screen.

See “Options handling in interactive mode” on page 274 for a discussion of how options are handled in interactive mode.

See Chapter 10, “Using commands,” on page 271 for more information on how to start and use the command-line client.

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
```

Starting a Web client session

After installing the Web client on your workstation (see “Configuring the Web client” on page 20) you can use the Web client to perform backup, archive, restore, and retrieve operations from any browser that is running Java Runtime Environment (JRE) 1.5 or 1.4.x, where *x* is greater than or equal to 1. 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.

The Web client runs in the following browsers:

- Safari
- Firefox

You can back up and restore your own data, or a Tivoli Storage Manager administrator can centralize the backup or restore operations of many clients.

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 HTTPport 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.

Setting 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. See “Tivoli Storage Manager client authentication” on page 39 for more information about these authorities.

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. See “Revokeremoteaccess” on page 229 for more information about the *revokeremoteaccess* option. See “Tivoli Storage Manager client authentication” on page 39 for more information about access authorities.

Starting 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). See Chapter 7, “Automating tasks,” on page 93 for more information about the client scheduler.

You can also use the Tivoli Storage Manager Client Acceptor daemon to manage the scheduler. See “Starting the client scheduler” on page 22 for more information.

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.

To change your password from the GUI:

1. Start Tivoli Storage Manager with TSM Tools for Administrators.
2. Enter your current and new passwords, and enter your new password again in the **Verify password** field.
3. Click **Change**.

To change your password from the command-line client, enter:

```
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. See “Password” on page 206 for additional password information.

The following are additional password information sources:

- “Starting the client scheduler automatically” on page 43
- “Password” on page 206
- “Set Password” on page 334

Sorting file lists using the Tivoli Storage Manager Java GUI

Table 13. 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	<ul style="list-style-type: none"> • Select one of the Sort by items from the View menu bar. • Click the appropriate column heading in the File List box.
Display active and inactive backup versions	<ul style="list-style-type: none"> • 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.	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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.

Notes:

1. Using the Tivoli Storage Manager GUI, 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 (see “Performing point-in-time restores” on page 76).
2. 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 the *inactive* and *pick* options with query and restore commands to display both active and inactive objects. See “Inactive” on page 183 and “Pick” on page 210 for more information.

Displaying online help

You can display online help in any of the following ways:

- **From the Tivoli Storage Manager GUI:** You can click the GUI question mark (?) icon, which displays online information about the current operation.
- **From the Web client:**
 - Select the **Help** menu.
 - Click the **Help** button in current window.
- **From the dsmc command line:** Enter the **help** command. A menu of topics is displayed for which help is available. See “Help” on page 290 for more information about the **Help** command.

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 will be sorted by relevance. Use the left and right arrow keys in the lower right-hand corner to navigate between help pages.

Ending a session

You can end a Tivoli Storage Manager client session in any one of the following ways:

- **From the Tivoli Storage Manager backup-archive Java GUI or Web client:**
 - Open the **File** menu and select **Quit**.
 - Press Command+Q.
 - **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.

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.software.ibm.com/storage](ftp://ftp.software.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. These servers are:

[ftp.rz.uni-karlsruhe.de](ftp://ftp.rz.uni-karlsruhe.de) (mirror - Germany)
[ftp.wu-wien.ac.at](ftp://ftp.wu-wien.ac.at) (mirror - Austria)
[ftp.cac.psu.edu](ftp://ftp.cac.psu.edu) (mirror - Pennsylvania)

You can get maintenance information from the Tivoli Storage Manager support page at:

<http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>

Also see “Contacting customer support” on page xiii for product support information.

Chapter 4. Backing up your data

Use Tivoli Storage Manager to store copies of your workstation files on a Tivoli Storage Manager server. These copies are called backup versions or backups. You can restore these backup versions from the server to your workstation if the original workstation files are lost or damaged. See Chapter 5, “Restoring your data,” on page 71 for more information.

As a quick reference, this chapter includes instructions for the following tasks.

Table 14. Backup: Primary tasks

Task	Page
Planning your backups	47
Pre-backup considerations	48
Performing an incremental, selective, or incremental-by-date backup	56
Backing up data using the Java GUI	59
Backing up data using the command line	60
Deleting backup data	61

Planning your backups

If you are a first-time user, or if you only back up files occasionally, you wish to use Table 15 as a checklist of preliminary steps to consider before performing a backup.

Table 15. Planning your backups

—	Decide whether you want to back up files or archive them. See “When to back up and when to archive files” for more information.
—	See “Pre-backup considerations” on page 48 for important considerations before you back up your files and directories.
—	Do you need to exclude files from backup services? See “Using an include-exclude options list to control processing” on page 51 for more information.
—	Decide what type of backup you want according to your needs. See the following sections for more information: <ul style="list-style-type: none">• “Performing an incremental, selective, or incremental-by-date backup” on page 56• “Backup with client node proxy support” on page 63
—	For further backup considerations, see “Backup: Additional considerations” on page 66.

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.

Recommendation: If archive is run multiple times against a file using the same archive description, a new copy of the file will be 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.

Note: If you frequently create archives for the same data, consider using instant archives (backup sets) instead. Frequent archive operations can create a large amount of metadata in the server database increasing database growth and decreasing performance for operations such as expiration. See “Restoring data from a backupset” on page 77 for more information on how backup sets can be generated and restored.

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. See Chapter 6, “Archiving and retrieving your data,” on page 85 for more information about archiving and retrieving files.

Pre-backup considerations

This section discusses important processing and performance considerations before you back up your data.

How Tivoli Storage Manager maintains 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.

Volume naming precautions

Beginning with Version 5.3.2, 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. For example:

```
/Volumes/La Pomme  
/Volumes/la pomme
```

Notes:

1. 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.
2. Volumes with ... in their name are not supported. Tivoli Storage Manager uses the ... sequence as part of include-exclude processing. Tivoli Storage Manager reports an invalid include-exclude statement if a volume has ... in the name. The volume *must* be renamed.

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.

The file system to which Tivoli Storage Manager will back up the data can be checked. In the TSM Backup window, select a file system. Then choose **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.

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 will 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
```

This means that when Tivoli Storage Manager does a backup or archive, the file space names will also depend upon which version of Mac OS X is running.

The real issue is that both versions of Mac OS X will backup to the / file system on the Tivoli Storage Manager server. If this happens, the system files will be intermixed.

Attention: To avoid any potential problems on dual-boot systems, you must do ONE of the following:

1. Select one version of Mac OS X on which to install and run Tivoli Storage Manager. This will ensure that the UNIX mount points are the same each time Tivoli Storage Manager does a backup. This is the recommended approach.
2. Configure each version of Mac OS X with a unique TSM node name. Then exclude the other version of Mac OS X from backup processing with a '-domain' statement in the system options (dsm.sys) file. For example, if Tiger is the startup disk, add this option to the system options (dsm.sys) file:

```
DOMAIN -/Volumes/Panther
```

If Panther is the startup disk, add this option to the system options (dsm.sys) file:

```
DOMAIN -/Volumes/Tiger
```

Unicode considerations

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. For migration information, see "Migrating to the Unicode-enabled client" on page 4.

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. For detailed information about this option, see "Autofsrename" on page 135.

Configure memory-constrained systems to run incremental backups

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 (dsm.opt) or in your dsm.sys 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* will degrade your backup performance.

Configure systems with large numbers of files to run incremental backups

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 will require, on average, approximately 300 MB of memory. The exact amount of memory used per object will vary, 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 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*.
- 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 will limit the process to incremental backup of a single file system at a time. This will reduce 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, 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.
- 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.
- 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.

Using an include-exclude options list to control processing

There might be 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. See Chapter 7, “Automating tasks,” on page 93 for more information on management classes and how to assign them to files.

For more information about creating an include-exclude list in your system options (dsm.sys) file, see “Creating an include-exclude list (optional root user or authorized user task)” on page 26.

Encrypting data during backup or archive operations

Tivoli Storage Manager Version 5.3.2 and higher supports 128-bit Advanced Encryption Standard (AES) data encryption, using the *encryptiontype* option, for the strongest possible encryption. The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received. See “Encryptiontype” on page 161 for more information.

Attention: 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.

For more information about the *include.encrypt* option, see “Include options” on page 186. For more information about the *exclude.encrypt* option, see “Exclude options” on page 168.

Use the *include.encrypt* and *exclude.encrypt* options to select files for encryption processing. By default, files are not encrypted unless they are explicitly included using the *include.encrypt* option. For more information about the *exclude.encrypt* option, see “Exclude options” on page 168. For more information about the *include.encrypt* option, see “Include options” on page 186.

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 will prompt 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.

For more information about this option, see “Encryptkey” on page 162.

Supported file systems

Tivoli Storage Manager supports HFS, HFS+, HFSX, UFS, UDF, and ISO 9600 file systems, as defined in Table 16.

Table 16. Supported file systems

File system	Tivoli Storage Manager support
Mac OS Standard (HFS)	<ul style="list-style-type: none"> • Case-insensitive, but case preserving • Supports creation and modification dates as meta data • Supports Aliases • 32 character limit for folder and file names
MAC OS Extended (HFS+)	<ul style="list-style-type: none"> • Case-insensitive, but case preserving • Supports creation and modification dates as meta data • Supports Aliases • 255 Character limit for folder and file names
MAC OS Extended, Case-sensitive (HFSX)	<ul style="list-style-type: none"> • Case-sensitive • Supports creation and modification dates as meta data • Supports Aliases • 255 Character limit for folder and file names
Mac OS Xsan (XSAN)	<ul style="list-style-type: none"> • Case-sensitive • Supports modification dates as meta data • Supports Aliases and symbolic links • 255 Character limit for folder and file names
UNIX File System (UFS)	<ul style="list-style-type: none"> • Case-sensitive • Supports modification dates as metadata • Supports aliases and symbolic links • 255 character limit for folders and file names
Universal Disk Format (UDF)	The Universal Disk Format for DVD volumes. Recognized as removable media.
ISO 9660	The standard format for CD-ROM volumes. Recognized as removable media.

The UFS and HFSX file systems are sensitive to case 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) 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.

Attention: 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 Manager 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 Macintosh HD and this system is repartitioned with a Case-Sensitive HFS+ file system. Either the Macintosh HD 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 will mix 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 the symbolic links point to.

When files backed up from an HFS volume are restored to a UFS volume, the resource forks will not have the correct owner. This can be corrected by using the **chown** command on the resource fork file.

Maximum file size for backup and restore

Table 17 specifies the maximum file sizes for backup and restore operations.

Note: 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. As long as the file system allows creation of the file, the Tivoli Storage Manager client will back up or archive the file.

Table 17. Maximum file size for backup and restore

Platform	Max file size (in bytes)
Mac OS X	HFS - 2,147,485,648 (2GB) HFS+, HFSX, XSAN, and UFS - 9,223,372,036,854,775,808 (8EB)

How Tivoli Storage Manager handles 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.

Recommendation: Do not exceed the 64 character limit for user and group names. If you do, Tivoli Storage Manager will shorten the name to fall within this limit by using the following transformation: Take the first 53 characters, append a /, and then the numeric id as a character string.

An error message will be logged containing 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 will either have to find the error message containing the transformation, or construct the name using the above outlined rule.

How Tivoli Storage Manager handles FileVault home directories

The FileVault feature compresses a user's home directory into a single file. Tivoli Storage Manager cannot access any user files inside a FileVault home directory.

Tivoli Storage Manager can back up the user's home directory as a single file. Due to how FileVault home directories are managed by Mac OS X, it is recommended that any FileVault home directories be backed up when the user is not logged in. When a user is logged in, the home directory is mounted in a temporary folder in /Users/.<username>, and as a result of this, the sparse file will be backed up again under this temporary folder.

Tivoli Storage Manager can restore a user's home directory if the user is not logged in. This means that a user cannot restore their own home directory. User home directories can be restored by the system administrator.

A second system administrator account must be created so that the primary system administrator account can be restored.

How Tivoli Storage Manager handles the Time Machine backup disk

Due to the unique nature of how Time Machine backs up data, consider the following before using Tivoli Storage Manager to back up the Time Machine data:

1. The 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 Time Machine is 5 GB, the first backup will copy all 5 GB of data to the Time Machine backup disk.

Subsequent backups will only copy the files that have changed since the previous backup. All files and directories that have not changed will be hard linked with the version copied over during the previous backup.

The Finder will show each backup as 5 GB, for a total size of 10 GB. However, because of the use of hard links, the total disk usage will only be slightly larger than 5 GB.

Due to how Tivoli Storage Manager handles hard linked data, it will back up all hard linked objects that are not already on the Tivoli Storage Manager Server.

For example, 10 GB of data would be sent to the Tivoli Storage Manager server.

2. During restore of files that are hard linked, Tivoli Storage Manager will attempt to recreate the original hard link. This can only be done if ALL files that are hard linked are restored at the same time. This would not be practical for a large Time Machine backup disk.
3. When Time Machine 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 would still have these restrictive ACLs in place.

| **Recommendation:** Due to how Tivoli Storage Manager handles the Time Machine
| back up data, exclude the Time Machine back up data. All Time Machine data is in
| a directory named Backups.backupdb. An example exclude statement is included in
| “Excluding system files” on page 30.

Performing an incremental, selective, or incremental-by-date backup

Your administrator might have set up schedules to back up files on your workstation automatically. See Chapter 7, “Automating tasks,” on page 93 for information on checking and running the schedules available to you. The following sections discuss how to back up files without using a schedule.

Full and partial incremental 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 folders on the volumes 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 will 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 will cause 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 will 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. See Chapter 8, “Understanding storage management policies,” on page 99 for more information about storage management policies.

To perform a full incremental backup see “Backing up data using the Java GUI” on page 59 or “Backing up data using the command line” on page 60.

During an incremental backup, the client queries the server to determine the exact state of your files since the last incremental backup. The client uses this information to:

- Back up new files
- Back up files whose contents changed since the last backup. Changes include *any* of the following:
 - 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:

- Locked attribute
- File permissions
- File owner
- Group ID
- Icon location
- Type or creator
- 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 will be 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.

Attention: 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

Folders are counted in the number of objects backed up. To exclude directories and their contents from backup, use the *exclude.dir* option. For more about *exclude.dir*, see “Exclude options” on page 168.

Understanding which files are backed up

When you request a backup, Tivoli Storage Manager backs up a file if all of the following requirements are met:

- The selected management class contains a backup copy group. See Chapter 8, “Understanding storage management policies,” on page 99 for more information on management classes and backup copy groups.
- 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. See “Using management classes and copy groups” on page 99 for more information.
- 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. See “Using management classes and copy groups” on page 99 for more information.
- 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. See “Using management classes and copy groups” on page 99 for more information.

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.

To perform an incremental-by-date backup, see “Backing up data using the Java GUI” on page 59 or “Backing up data using the command line” on page 60.

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 volume, the server does not update the date of the last full incremental backup. In this case, the next incremental-by-date backup will back up these files again.

Tip: 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.

Comparing full incremental, partial incremental, and incremental-by-date backups

Full incremental, partial incremental, and incremental-by-date all back up new and changed files. 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.

Selective backup

Use a selective backup when you want to back up specific files or folders 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.

To perform a selective backup using the Tivoli Storage Manager GUI, see “Backing up data using the Java GUI” on page 59. To perform a selective backup using the Tivoli Storage Manager command line, see “Backing up data using the command line” on page 60.

Unlike incremental backups, a selective backup:

- 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.

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 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 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.

To search:

- a. Enter your search criteria in the Find Files (Backup) window.
- b. Click the **Search** button. The Matching Files (Backup) window appears.
- c. Click the selection boxes next to the files you want to back up and close the Matching Files (Backup) window.

To filter:

- a. Enter your filter criteria in the Find Files (Backup) window.
 - b. Click the **Filter** button. The Backup window displays the filtered files.
 - c. Click the selection boxes next to the filtered files or directories you want to back up.
3. Select one of the following backup types from the pull-down menu:
 - To run an incremental backup, click **Incremental (complete)**.
 - To run an incremental-by-date backup, click **Incremental (date only)**.
 - To run a selective backup, click **Always backup**.
 4. Click **Backup**. The Backup **Task List** window displays the backup processing status.

Considerations:

- 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 GUI or Web client GUI to view the backup policies defined by the Tivoli Storage Manager server for your client node. See Chapter 8, “Understanding storage management policies,” on page 99 for more information on management classes, how an authorized user assigns management classes to files, and storage management policies.
- 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*. See “Domain” on page 157 for more information.
- You can use the Preferences editor to exclude file systems in your default domain from backup processing.

Backing up data using the command line

You can use the **incremental** or **selective** commands to perform backups. Table 18 shows examples of using these commands to perform different tasks. See “Incremental” on page 291 and “Selective” on page 328 for more information about these commands.

Table 18 shows examples of using commands to perform different tasks.

Table 18. Command line backup examples

Task	Command	Considerations
<i>Incremental backups</i>		
Perform an incremental backup of your client domain.	dsmc incremental	See “Incremental” on page 291 for more information about the incremental command. See “Full and partial incremental backup” on page 56 for detailed information about incremental backups.
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 157 for more information about the <i>domain</i> option.
Back up all local file systems defined in your client domain <i>except</i> for the jaguar file system.	dsmc incremental -domain="all-local -/Volumes/jaguar"	You cannot use the (-) operator in front of the domain keyword all-local. See “Domain” on page 157 for more information.
Back up <i>only</i> the /Volumes/fs1 and /Volumes/fs2 file systems.	dsmc incremental /Volumes/fs1 /Volumes/fs2	None
Back up all files in the /Users directory and all its subdirectories.	dsmc incremental /Users/ -subdir=yes	See “Subdir” on page 247 for more information about the <i>subdir</i> option.
<i>Incremental-by-date backup</i>		
Perform an incremental-by-date backup of your default client domain.	dsmc incremental -incrbydate	Use the <i>incrbydate</i> 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 190 for more information about the <i>incrbydate</i> option.
<i>Selective backups</i>		
Back up all files in the /Users/van/Documents directory.	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 328 for more information about the selective command.

Table 18. Command line backup examples (continued)

Task	Command	Considerations
Back up all files in the /Users/van/Documents directory and all its subdirectories.	dsmc selective /Users/van/Documents/ -subdir=yes	<p>If you specify <i>-subdir=yes</i> when backing up a specific path and file, Tivoli Storage Manager recursively backs up <i>all</i> subdirectories under that path, and any instances of the specified file that exist under <i>any</i> of those subdirectories.</p> <p>If a subdirectory is a mounted file system, Tivoli Storage Manager does not back up the files in that subdirectory when you use the <i>subdir=yes</i> option. See “Subdir” on page 247 for more information about the <i>subdir</i> option.</p>
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	<p>If you specify the <i>removeoperandlimit</i> 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 222 for more information about this option.</p>
Back up all files listed in the /Users/filelist.txt file.	dsmc selective -filelist=/Users/ filelist.txt	<p>Use the <i>filelist</i> option to process a list of files. See “Filelist” on page 171 for more information.</p>

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.

Attention: 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 will prompt 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: A directory will be deleted only if you select **Delete All Objects**.

To delete backup copies using the Tivoli Storage Manager command line client, use the **delete backup** command. See “Delete Backup” on page 285 for more information.

Displaying 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 file’s size, path, file name, total number of bytes transferred, and whether the backup attempt was successful. 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.

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. Informational messages might display as follows:

Table 19. Client command line informational messages

Informational message	Meaning
Directory-->	Indicates the directory that you back up.
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.
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 “Binding management classes to files” on page 106 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 “Full and partial incremental backup” on page 56 for more information.
Total number of objects failed:	Objects can fail for several reasons. Check the <code>dsmerror.log</code> for details.
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 19. Client command line informational messages (continued)

Informational message	Meaning
Aggregate data transfer rate:	<p>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.</p> <p>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 <i>resourceutilization</i> option, Tivoli Storage Manager attempts to improve performance and load balancing by using multiple sessions when it backs up a 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.</p>
Objects compressed by:	<p>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%.</p>
Elapsed processing time:	<p>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.</p>
Total number of bytes transferred:	As indicated.

Backup with client node proxy support

Backups of multiple nodes which share storage can be consolidated to a common target nodename on the Tivoli Storage Manager server. This is useful when the system responsible for performing the backup can change over time, such as with an Xsan or 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 a Tivoli Storage Manager Version 5.3 and higher server and client. To enable this option, follow these steps:

- Install the backup-archive client on all nodes in the shared data environment.
- Register each node with the Tivoli Storage Manager server, if it does not exist. Register the common target nodename to be shared by each of the agent nodes used in your shared data environment.

- Register each of the nodes in the shared data environment with the Tivoli Storage Manager server. This is the agent nodename which is used for authentication purposes. Data will not be stored using the nodename when the *asnodename* option is used.
- The Tivoli Storage Manager administrator must 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.
- Use the **query proxynode** administrative client command to display the authorized user's client nodes, granted by the **grant proxynode** command.

Recommendation: Do not use the *asnodename* option when a node is backing up or restoring local file systems. Only use this option when backing up or restoring to Xsan or cluster file systems. If this recommendation is not followed, it is possible for the local file system data of each client node to become mixed on the Tivoli Storage Manager Server. This will happen for all client nodes have the same local file system names.

Follow these steps 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** → **Client 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.
5. Quit and restart the Tivoli Storage Manager client so that the new options can take effect.

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 as the node name in the tree view, 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.

Suggestion:

- All agent nodes in the multiple node environment should be of the same platform type.

Restrictions enforced within a proxied session:

- You cannot access another node (either from the GUI drop down or use of the Fromnode option).

Follow these steps to set up encryption with the *encryptkey=save* option:

1. Specify *encryptkey=save* in the options file.
2. Add an include.encrypt statement in the system options (dsm.sys) file for the files that need to be encrypted.
3. 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 on all participating agent node systems to ensure that the encryption key is the same.
 - Use the same encryption key for all files backed up in the shared node environment.

See “Asnodename” on page 129 for more information.

Scheduling with client node proxy support

Multiple nodes can be used to perform backup operations using the scheduler. By granting proxy authority to the agent nodes, they will 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.

The following examples show the administrative client server commands using the scheduler on multiple nodes.

1. The administrator registers all of the nodes to be used:

```
register node NODE-A
register node NODE-B
register node NODE-C
register node NODE-Z
```

2. The administrator grants proxy authority to each agent node:

```
grant proxynode target=NODE-Z agent=NODE-A
grant proxynode target=NODE-Z agent=NODE-B
grant proxynode target=NODE-Z agent=NODE-C
```

3. The administrator defines the schedules:

```
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
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
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 use the client acceptor daemon (dsmcad), with managedservices set to schedule in the system options (dsm.sys) file.

Notes:

1. Each schedule can be started from a different system.
2. After running the schedules, any proxied client will be able to query and restore all of the backed up data.

You can also set the schedules using the following client command:

```
dsmc sched
```

The following steps must be taken to enable scheduling of multiple nodes:

1. All agent nodes must have proxy authority over common target node
2. All agent nodes must have a schedule defined on the server

```
def sched domain_name sched_name options='-asnode=target'
```

3. Each agent node must have its schedule associated with node

```
def association domain_name schedule_name <agentnodename>
```

Backup: Additional considerations

This section discusses some advanced considerations for using backup. You do not need to understand this information to use Tivoli Storage Manager for basic work.

Understanding how files are stored

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 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 new 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 new mount point was defined. However, it does not locate the file named doc1 because that file was backed up before the new 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 new 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.

Understanding how files are managed

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** in the client GUI to view the backup policies defined by the Tivoli Storage Manager server for your client node. See Chapter 8, “Understanding storage management policies,” on page 99 for more information on management classes, how an authorized user assigns management classes to files, and storage management policies.

Understanding how hard links are handled

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 will back 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 will be hard-linked. 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, Tivoli Storage Manager will respect the current file system and not restore the hard link.

If you do not back up and restore all files that are hard-linked at the same time, problems will occur. To ensure that hard-linked files remain synchronized, back up all hard links at the same time and restore those same files together.

Understanding how NFS hard and soft mounts are handled

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 will try repeatedly to contact the server. The NFS daemon retries will not time out, will affect system performance, and you cannot interrupt them, but control will return to Tivoli Storage Manager when the *nfstimeout* value is reached.

Soft mount

— If the NFS file system is soft mounted, NFS will try 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.

Understanding how deleted file systems are handled

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
- 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.

See Chapter 8, "Understanding storage management policies," on page 99 for further information about policy management.

Backing up open files

Some files on your system be in use, or open, when you try backing them up. Because an open file can change during a backup, the backup might not accurately reflect the contents of the file. If this is the case, you need to consider the following options:

- If the file is unimportant, or can be easily rebuilt (for example, a temporary file), you might not care that the file is not backed up.
- If the file is important:
 - If the backup is not scheduled, close the file and the application before you start the backup.

If the backup is scheduled, use the *preschedulecmd* option to issue a command that closes the file before the backup occurs. For example, if the open file is a database, use the database's quiesce command to shut down the database. You can use the *postschedulecmd* option to restart the application that uses the file after the backup completes.

- You can back up the file even if it is open and changes during the backup, by assigning the file a management class with the serialization **dynamic** or **shared dynamic**. This is only useful if the file is usable even if it changes during the backup. See “Selecting a management class for files” on page 104 for information on assigning management classes, and “Displaying information about management classes and copy groups” on page 101 for information on determining which management classes are available to you.

Using wildcard characters

You can use the operating system's 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 "/Users/user1/Documents/Ledger.*"
```

enclose the parameter containing the asterisk in quotes to ensure the system does not interpret the wildcard character and produce unexpected results. Wildcard character information is covered in the following table.

<i>* (Asterisk)</i>	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 .
<i>? (Question mark)</i>	One character that matches all files with:
?.*cpp	The extension cpp with only one character in the file name
hm?.cpp	Three-character names that begin with hm and have the cpp extension

Attention: 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.

<i>* ? (Asterisk and question mark)</i>	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 will recognize those characters only as wildcard characters.

Chapter 5. Restoring your data

You can use Tivoli Storage Manager to restore the following:

- Backup versions of specific files
- A group of files with similar names
- Entire folders
- Backupsets

For more information on restoring backupsets, see “Restoring data from a backupset” on page 77.

This chapter includes instructions for the following tasks.

Table 20. Restore: Primary tasks

Task	Page
Restoring data using the GUI	74
Restoring data using the command line	75
Performing point-in-time restores	76
Restoring data from a backupset	77
Authorizing another user to restore or retrieve your files	80
Restoring or retrieving files from another client node	81
Restoring or retrieving your files to another workstation	81
Deleting file spaces	82

Restoring an active or inactive backup

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 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. See “Inactive” on page 183 for more information.

Performing large restore operations

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 /Users file system, enter:

```
dsmc restore /Users/ -subdir=yes -replace=all -tapeprompt=no
```

However, if you enter multiple commands for the directories in the /Users file space, you can restore the files faster.

For example, you could enter these commands:

```
dsmc restore /Users/user11/ -subdir=yes -replace=all -tapeprompt=no  
dsmc restore /Users/user21/ -subdir=yes -replace=all -tapeprompt=no  
dsmc restore /Users/user31/ -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 -dironly -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 network utilization and how much memory you have.

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.

If your administrator is using collocation, the number of sequential access media mounts required for restore operations is also reduced.

Standard query restore, no query restore, and restartable restore

The following sections describe the standard (or classic) restore and the no query restore processes.

Standard restore process

The standard restore process is also known as classic restore, and is outlined below.

1. The client queries the server for a list of files backed up for the client file space you want to restore.
2. 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.
3. 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.
4. The client tells the server to restore file data and directory objects.
5. 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. See “Resourceutilization” on page 225 for more information.

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*, or *volinformation*, 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/*"
```

An example of a restricted wildcard file specification would be:

```
/Users/user1/Documents/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 will fail 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 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 will restart 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.

For more information on using the command line to begin restartable restores, see "Restore" on page 316. 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 files and directories

You can restore specific files, a group of files with similar names, or directories.

You can locate the files you want to restore by searching and filtering. Filtering displays only the files that match the filter criteria for your restore operation. Files that do not match the filter criteria do not display. The filter process searches files in the specified directory but does not include subdirectories.

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.

Restoring data using the GUI

To run a restore, perform the following steps:

1. Click **Restore** from the GUI main window. The Restore window appears.
2. Expand the folder tree. Click the object that you want to restore. To search or filter files, click the **Edit** menu → **Find** item. To perform a point-in-time restore that restores files to the state that existed at a specific date and time, click the **Point-in-Time** button.
3. To modify specific restore options click the **Options** button. The options you select are effective during the current session *only*.
4. Click **Restore**. The Restore Destination window appears. Enter the appropriate information.
5. Click **Restore**. The Task List window displays the restore processing status. The Restore Report window displays processing statistics.

Considerations:

- 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.
- When restoring files with the *replace* option set to *no*, existing files will not be overwritten, but existing directories will. To leave existing directories intact during a restore operation, select the **Options** button → **All selected files and directories** dropdown menu → **Files only** option.
- When folders are restored from a UFS or HFSX file system to a HFS file system and they differ only in case, the client will restore the contents of both folders to one folder.

Restoring data using the command line

Use the **restore** command to restore files. See “Restore” on page 316 for more information about the **restore** command. Table 21 shows examples of using the **restore** command to restore objects from Tivoli Storage Manager server storage. See “Restore” on page 316 for additional examples.

Table 21. 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.	<code>dsmc restore /Users/monnett/Documents/h1.doc -latest</code>	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 <i>latest</i> 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 192 for more information.
Display a list of active and inactive backup versions of files from which you can select versions to restore.	<code>dsmc restore "/Users/monnett/Documents/*"-pick -inactive</code>	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 210 and “Inactive” on page 183 for more information.
Restore the /Users/monnett/Documents/h1.doc file to its original directory.	<code>dsmc restore /Users/monnett/Documents/h1.doc</code>	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.	<code>dsmc restore /Users/monnett/Documents/h1.doc /Users/gordon/Documents/h2.doc</code>	None
Restore the files in the /Users directory and all of its subdirectories.	<code>dsmc restore /Users/ -subdir=yes</code>	When restoring a specific path and file, Tivoli Storage Manager recursively restores <i>all</i> subdirectories under that path, and any instances of the specified file that exist under <i>any</i> of those subdirectories. See “Subdir” on page 247 for more information about the <i>subdir</i> option.

Table 21. 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 211 and “Pittime” on page 212 for more information about the <i>pitdate</i> and <i>pittime</i> 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 will 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 171 for more information about restoring a list of files.

Performing point-in-time restores

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.

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.

Notes:

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. See Chapter 8, “Understanding storage management policies,” on page 99 for more information about the *versions data deleted* attribute.

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 will restore 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.

Restoring data from a backupset

Your Tivoli Storage Manager administrator can generate a backupset (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.

It is possible to generate a backupset as a number of special files if the device class the Tivoli Storage Manager administrator specifies when creating it is *file*. These files can be stored locally (on the client) to provide more restore flexibility.

You can restore backupsets from the following locations:

- Directly from the server
- From portable media on a device attached to your client workstation
- From a backupset file on your client workstation

Backupsets can provide you with instant archive and rapid recovery capability as described below:

Instant archive

This capability allows an administrator to create an archive collection from backup versions already stored on the server.

Rapid recovery

When you are away from your office without a network connection and you lose data, you can restore the data from the backupset.

Note: If you would like to restore a backupset from portable media on a device attached to your client workstation, check with your Tivoli Storage Manager administrator to ensure that the portable media was created on a device using a format that is compatible with your client device.

Restoring an entire or partial backupset

Tivoli Storage Manager considers a backupset as one object containing the whole file structure. You can restore the entire backupset or just select portions. The backupset 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 higher 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.

When a backup set is stacked, you can only restore data for your own node. Data for all other nodes will be 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.

Attention: Due to the portability of local backup sets, additional steps should be taken to secure local backup sets on portable media. The backup set media should be physically secured since 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 backupset data from the server, individual files, directories or entire backupset data can be restored in a single operation from the GUI or the command line. When you restore backupset data locally, the GUI can only display and restore an entire backupset. The command line can be used to restore individual files or directories stored in a backupset locally. Refer to the *backupsetname* option on the **restore** and **query** commands for more information.

Restoring backupsets using the Web client GUI

Attention: Before you begin a restore, be aware that backupsets 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 backupset, perform the following steps:

- Click **Restore** from the GUI main window. The Restore window appears.
- Locate the **Backup Sets** directory tree object and expand it by clicking the disclosure triangle (+) beside it.
 - To restore the backupset from a local device, expand the **Local** object and the Specify backupset location dialog is displayed. On the dialog, select **File**

name: from the dropdown list and then enter the file name location. You can also click the **Browse** button to open a file selection window and select a backupset.

- To restore data from backupset from the server, first expand the **Server** object and then either **Filelevel** or **Image**, depending on the type of restore requested.
- Click the selection box next to the backupset or directory or file within the backupset that you want to restore.
You can select files from within a backupset if that backupset is located at the server and it was generated with a table of contents.
- Click **Restore**. The Restore Destination window appears. Enter the appropriate information.
- Click **Restore**. The Task List window displays the restore processing status.

Notes:

1. If the object you want to restore is part of a backupset generated on a node, and the node name is changed on the server, any backupset 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 backupset was generated.
2. The client can be used to restore a backupset on an attached device with or without a server connection. If the server connection fails, a prompt will appear to continue for purposes of local backupset restore. Also, the `localbackupset` option can be used to tell the client not to attempt the connection to the server.
3. Certain local devices require device drivers to be set up prior to performing a restore. See the device manual for assistance with this task. You will also need to know the device address in order to perform the restore.
4. The following features of a backupset restore from the server are not available when restoring locally:
 - a. 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 backupset.

Restoring backupsets using the command-line client

The **restore backupset** command restores a backupset from the Tivoli Storage Manager server or a local file. To restore individual files, folders, and images from a backupset, use the **restore** command with the *backupsetname* option instead. Refer to the *backupsetname* option on the **restore backupset**, **restore**, and **query** commands for more information.

Use the *location* option with the **restore backupset**, **restore**, and **query** commands to specify where Tivoli Storage Manager searches for a backupset during a query or restore operation. You can use this option to locate backupsets on the server or local files. Refer to the *location* option on the **restore backupset**, **restore**, and **query** commands for more information.

Use the **query backupset** command to query a backupset from a local file or the Tivoli Storage Manager server. The **query backupset** command displays the backupset name, generation date, retention, and description. Refer to the **query backupset** command for more information.

Related tasks and considerations

This section discusses some advanced considerations for restoring files. You do not need to understand this information to perform basic Tivoli Storage Manager tasks.

Authorizing another user to restore or retrieve your files

You can authorize a user on another node to restore your backup versions or retrieve your archive copies. This allows you to share files with other people or with other workstations that you use with a different Tivoli Storage Manager node name.

To authorize another node to restore or retrieve your files:

1. From the Tivoli Storage Manager main window select the **Utilities** menu → **Node Access List** item. The Node Access List window displays.
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 Backed up Objects or Archived Objects.
4. Type the node name of the user in the **Grant Access to Node** field.
5. Type the name of the user on the node who can access your data in the **User Name** field. This field can contain the name of any user on the node, or an asterisk (*) to grant access to all users logged in under the specified node name.
6. In the **Filespace Name** 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. The entry can 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. 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 new access rule appears in the list box in the Node Access List window. When you finish 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.

On the command-line client, 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. For more information about these commands, see:

- “Set Access” on page 330
- “Query Access” on page 298
- “Delete Access” on page 282

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 another user's file spaces on the Tivoli Storage Manager server, restore the other user's backup versions, or retrieve the other user's archive copies to your local drives.

1. From the Tivoli Storage Manager main window select the **Utilities** menu → **Access Another Node** item. The Access Another User window displays.
2. Enter the user's node name in the **Node Name** field.
3. Enter the user name of the user who owns the files in the **User Name** field.
4. Recreate any UNIX mount points that were manually created prior to the disk loss. This step is required for those systems configured to mount file systems at locations other than /Volumes. This step is NOT needed for mount points automatically created by Mac OS.
5. 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 from the /Users/devel/proja directory belonging to *usermike* on system *puma*, and place the restored files in a directory you own named /Users/id/proja, enter:

```
dsmc restore -fromowner=usermike -fromnode=puma /Users/devel/proja/  
/Users/id/proja/
```

You must specify a destination, or Tivoli Storage Manager attempts to restore the files to the node from which the files were backed up. See "Fromnode" on page 176 for more information about the *fromnode* option. See "Restore" on page 316 for more information about using the *fromnode* and *fromowner* options with the **restore** command.

You can also use the *fromnode* and *fromowner* options with the *retrieve* command, see "Retrieve" on page 323 for more information.

Use the **query filespace** command to display 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
```

See "Query Filespace" on page 305 for more information.

See Table 39 on page 124 for a list of other commands for which you can use the *fromnode* and *fromowner* options.

Restoring or retrieving your files to another workstation

Tivoli Storage Manager lets you restore or retrieve files you back up from your own workstation, when you are using a different workstation. You just need to specify your node name, because your backups and archives are stored according to your node, not your specific system. Your Tivoli Storage Manager password protects your data.

To restore or retrieve files to another workstation, either change the *nodename* option in your system options (dsm.sys) file to match the current workstation, or type it in when prompted for the nodename and password. To do this, select **File** → **Login** from the Tivoli Storage Manager GUI and enter the node name, userid, and password. You can then restore or retrieve files as if you were working on your original workstation.

If you do not want to restore or retrieve the files to the same folder name on the alternate workstation, use a different destination.

Restoring a disk in case of disk loss

Tivoli Storage Manager 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, perform the following steps:

1. Record the IP address, subnet mask, default router, and name servers in the TCP/IP Control Panel. You can use this information to reconfigure TCP/IP when you re-install Mac OS.
2. Record your nodename, TCP Server address, and TCP Port number from the **Communications** panel of the Preferences editor. The nodename information comes from the **General** panel in the Preferences editor. You will need this information to configure Tivoli Storage Manager so you can restore your data.
3. Re-install Mac OS from your Macintosh installation CD to access the network.
4. Install the Tivoli Storage Manager for Macintosh installation program. You can get this from: <http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>

Deleting file spaces

Authorized User

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 can delete individual backup versions by using the delete backup command. See “Delete Backup” on page 285 for more information.

You can delete file spaces using the Tivoli Storage Manager GUI or command-line clients.

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 filesystem** command. See “Delete Filespace” on page 288 for more information.

Overlapping icons

Tivoli Storage Manager remembers a file’s position on your desktop when it backs up the file and tries to restore it back to its original position. If a file is in the top, left corner of your desktop when you back it up, Tivoli Storage Manager tries to put it there when you restore it.

But if something else is in that same position when you restore, it is possible that folders and icons can overlap. If this occurs, in the Finder, select the **View** menu → **View Options** item → **by Name** item, to list objects alphabetically, with no overlapping.

Chapter 6. Archiving and retrieving your data

Archiving and retrieving files is similar to backing up and restoring them; many of the windows and concepts are the same. See Chapter 4, “Backing up your data,” on page 47 for more information. See “When to back up and when to archive files” on page 47 for a discussion of the differences between backups and archives.

As a quick reference, this chapter includes instructions for the following tasks:

Table 22. Archive and retrieve: Primary tasks

Task	Page
Archiving data using the GUI	85
Archiving data using the command line	86
Deleting archive data	89
Retrieving archives	90
Retrieving data using the command line	91

Archiving 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. See Chapter 7, “Automating tasks,” on page 93 for information on checking and running the schedules available to you. The following sections cover how to archive files without using a schedule.

Maximum file size for archive and retrieve operations

Note: 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. As long as the file system allows creation of the file, the Tivoli Storage Manager client will back up or archive the file.

Table 23 specifies the maximum file sizes for archive and retrieve operations.

Table 23. Maximum file size for archive and retrieve

Platform	Max file size (in bytes)
Mac OS X	HFS - 2,147,485,648 (2GB) HFS+, HFSX, XSAN, and UFS - 9,223,372,036,854,775,808 (8EB)

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.

To search:

- a. Enter your search criteria in the Find Files (Archive) window.
- b. Click the **Search** button. The Matching Files (Archive) window appears.
- c. Click the selection boxes next to the files you want to archive and close the Matching Files (Archive) window.

To filter:

- a. Enter your filter criteria in the Find Files (Archive) window.
 - b. Click the **Filter** button. The Archive window displays the filtered files.
 - c. Click the selection boxes next to the filtered files or directories you want to archive.
3. 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.
 4. 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.
 5. Click **Archive**. The Archive **Task List** window displays the archive processing status.

Archiving data 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. 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. Use the **archive** command to archive files. See “Archive” on page 279 for more information about the **archive** command.

Table 24 shows examples of using the **archive** command to archive objects. See “Archive” on page 279 additional examples.

Table 24. Command line archive examples

Task	Command	Considerations
Archive all files in the /Users/ann/Documents directory with a file extension of .txt.	dsmc archive "/Users/ann/Documents/*.txt"	Use wildcards to archive more than one file at a time.
Archive all files in the /Users/jones/Documents directory and delete the files on your workstation.	dsmc archive /Users/jones/Documents/ -deletefiles	Retrieve the archived files to your workstation whenever you need them again. See “Deletefiles” on page 149 for more information about the <i>deletefiles</i> option.

Table 24. Command line archive examples (continued)

Task	Command	Considerations
Archive the /Users/jones/Documents/h1.doc and /Users/jones/Documents/test.doc files.	<code>dsmc archive /Users/jones/Documents/h1.doc /Users/jones/Documents/test.doc</code>	If you specify the <i>removeoperandlimit</i> 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 222 for more information about this option.
Archive a list of files in the /Users/avi/Docs/filelist.txt file.	<code>dsmc archive -filelist=/Users/avi/Docs/filelist.txt</code>	Use the <i>filelist</i> option to process a list of files. See “Filelist” on page 171 for more information.
Archive the /Users/jones/Documents/ch1.doc file and assign a description to the archive.	<code>dsmc archive /Users/jones/Documents/ch1.doc -description="Chapter 1, first version"</code>	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 150 for more information about the <i>description</i> option.
Archive all of the files in the /Users/jones/Documents directory and its subdirectories.	<code>dsmc archive /Users/jones/Documents/ -subdir=yes</code>	See “Subdir” on page 247 for more information about the <i>subdir</i> option.
Use the <i>v2archive</i> option with the archive command to archive only files in the /Users/ann/Documents directory.	<code>dsmc archive "/Users/ann/Documents/*" -v2archive</code>	Tivoli Storage Manager archives only files in the /Users/ann/dir1 directory. Directories that exist in the path are not processed. See “V2archive” on page 266 for more information about the <i>v2archive</i> option.
Use the <i>archmc</i> option with the archive command to specify the available management class for your policy domain to which you want to bind your archived files.	<code>dsmc archive -archmc=ret2yrs /Users/ann/Documents/budget.jan</code>	See “Archmc” on page 128 for more information about the <i>archmc</i> option. See Chapter 8, “Understanding storage management policies,” on page 99 for more information about management classes.

Archiving data with client node proxy

Archives of multiple nodes which share storage can be consolidated to a common target nodename on the Tivoli Storage Manager server. This is useful when the system responsible for performing the archive can change over time, such as with a Xsan or cluster.

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 a Tivoli Storage Manager Version 5.3 and higher server and client. To enable this option, follow these steps:

- Install the backup-archive client on all nodes in the shared data environment.
- Register each node with the Tivoli Storage Manager server, if it does not exist. Register the common target nodename to be shared by each of the agent nodes used in your shared data environment.
- Register each of the nodes in the shared data environment with the Tivoli Storage Manager server. This is the agent nodename which is used for authentication purposes. Data will not be stored using the nodename when the *asnodename* option is used.
- The Tivoli Storage Manager administrator must 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** administrative client command.
- Use the **query proxynode** administrative client command to display the authorized user's client nodes, granted by the **grant proxynode** administrative client 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.
3. **Ensure that users of the agent nodes in the multiple node environment are using the same encryption key.**

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 multiple node operation from the GUI:

1. Verify that the client node has proxy authority to a target node using the **query proxynode** administrative client command.
2. Select **Edit → Client 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.
5. Quit and restart the client so that the new options can take effect.

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 as the node name at the top of the directory tree.
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 Client Preferences** tab.

Attention:

- All agent nodes in the multiple node environment must be of the same platform type.

Restrictions enforced within a proxied session:

- You cannot access another node or owner (either from GUI drop down or use of the *fromnode*, *fromowner* options).

See “Asnodename” on page 129 for more information.

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 main menu. Your authority status is provided in the **Delete Archive Files** field.

To delete archive copies from the Tivoli Storage Manager GUI:

1. Select **Delete Archive Data** from the **Utilities** menu. The Archive Delete window appears.
2. Expand the Directory tree by clicking on the disclosure triangle 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.
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. See “Delete Archive” on page 283 for more information.

Archive: Advanced considerations

This section covers some advanced considerations in archiving files. You do not need to understand this information in order to use Tivoli Storage Manager for basic work.

Saving access permissions

When you archive your files, Tivoli Storage Manager also saves Access Control Lists (ACLs).

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.

Understanding how hard links are handled

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 will archive the file data twice.

When you retrieve 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 retrieve both files, they will be hard-linked. The one 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 will respect the current file system and not retrieve the hard link.

If you do not archive and retrieve all files that are hard-linked at the same time, problems will occur. To ensure that hard-linked files remain synchronized, archive all hard links at the same time and retrieve those same files together.

Retrieving 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.

Note: When retrieving a directory, its modification date and time is set to the date and time of the retrieve, 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.

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 will be retrieved. If there is an existing copy on your client system, you are prompted to replace, skip, or cancel.

Retrieving data using the GUI

To retrieve an archived file:

1. Click **Retrieve** from the main window. The Retrieve window appears.
2. Expand the folder tree. Select the object you want to retrieve. To search or filter files, click the **Edit** menu **Find** item.

To search:

- a. Enter your search criteria in the Find Files (Retrieve) window.
- b. Click the **Search** button. The Matching Files (Retrieve) window appears.
- c. Click the selection boxes next to the files you want to retrieve and close the Matching Files (Retrieve) window.

To filter:

- a. Enter your filter criteria in the Find Files (Retrieve) window.
 - b. Click the **Filter** button. The Retrieve window displays the filtered files.
 - c. Click the selection boxes next to the filtered files or directories you want to retrieve.
3. To modify specific retrieve options, click the **Options** button. Any options you change are effective during the current session *only*.
 4. Click **Retrieve**. The Retrieve Destination window appears. Enter the appropriate information in the Retrieve Destination window.
 5. Click **Retrieve**. The Retrieve **Task List** window displays the retrieve processing status.

Considerations:

- When retrieving files with the *replace* option set to *no*, existing files will not be overwritten, but existing directories will. To leave existing directories intact during a retrieve operation, select the **Options** button → **All selected files and directories** dropdown menu → **Files only** option.

- 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 will retrieve the contents of both folders to one folder.

Retrieving data using the command line

Use the **retrieve** command to retrieve files from storage to your workstation.

Table 25 shows examples of using the **retrieve** command. See “Retrieve” on page 323 for additional examples, and detailed information about the **retrieve** command.

Table 25. command-line examples of retrieving archives

Task	Command	Considerations
Retrieve the /Users/jones/Documents/h1.doc file to its original directory.	dsmc retrieve /Users/jones/Documents/h1.doc	If you do not specify a destination, the files are retrieved to their original location.
Retrieve the /Users/jones/Documents/h1.doc file under a new name and directory.	dsmc retrieve /Users/jones/Documents/h1.doc /Users/smith/Documents/h2.doc	None
Retrieve all files from the /Users/jones/Documents directory that end with the characters .bak to the /Users/smith/Documents directory.	dsmc retrieve "/Users/jones/Documents/*.bak" /Users/smith/Documents/	None.
Use the <i>pick</i> option to display a list of archives from which you can select files to retrieve.	dsmc retrieve "/Users/jones/Documents/*" -pick	See “Pick” on page 210 for more information about the <i>pick</i> option.
Retrieve files listed in retrievelist.txt.	dsmc retrieve -filelist=Usrs/av/Doc/retrievelist.txt	See “Filelist” on page 171 for more information about retrieving a list of files.

Understanding how your archives are managed

As with backing up files, 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 specifically assign a management class to a file with an *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 GUI. Then you can use the **Utilities Preview Include-Exclude** function to preview the include-exclude list prior to sending data to the server. For more information about the *preview* command, see “Preview” on page 297.

For information on the various management class attributes used to manage your archives, see “Displaying information about management classes and copy groups” on page 101. See “Assigning a management class to files” on page 105 for information about using the include-exclude options list.

Chapter 7. Automating tasks

Your 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 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 also:

- Display information about available schedules.
- Display information about work that the schedule has completed.
- Modify scheduling options in the dsm.sys. See “Scheduling options” on page 118 for more information.

Specifying scheduling options

Most options that apply to scheduling are in the system options (dsm.sys) file and can be set in the Preferences editor. However, server options can override option settings established in client option files.

For more information about scheduling options, changing the scheduling mode, or specifying the TCP/IP address or port number, see “Scheduling options” on page 118.

Displaying 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. Figure 1 on page 94 displays 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
```

Figure 1. Sample classic query schedule output

The schedule name, **DAILY_INC**, starts a daily incremental backup. The next incremental backup will start 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. Figure 2 on page 95 displays 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
```

Figure 2. Sample enhanced query schedule output

For more information about the enhanced **query schedule**, see the following publications:

- *IBM Tivoli Storage Manager for AIX Administrator's Reference*
- *IBM Tivoli Storage Manager for AIX Administrator's Guide*
- *IBM Tivoli Storage Manager for HP-UX Administrator's Reference*
- *IBM Tivoli Storage Manager for Linux Administrator's Reference*
- *IBM Tivoli Storage Manager for Sun Solaris Administrator's Reference*
- *IBM Tivoli Storage Manager for Sun Solaris Administrator's Guide*
- *IBM Tivoli Storage Manager for Windows Administrator's Reference*
- *IBM Tivoli Storage Manager for Windows Administrator's Guide*

Displaying information about completed work

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 will 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 command's return code 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.

Scheduling 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 will be 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).

Enabling or disabling 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. See “Schedcmddisabled” on page 230 for more information.

You can use the *schedrestretrdisabled* option to prevent the Tivoli Storage Manager Server administrator from executing restore or retrieve schedule operations. See “Schedrestretrdisabled” on page 237 for more information.

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. See “Srvprepostscheddisabled” on page 246 for more information.

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 Tivoli Storage Manager’s scheduler, the return codes are shown in the output of the QUERY EVENT administrative command. See the *IBM Tivoli Storage Manager Administrator’s Reference* for your operating system for more information about QUERY EVENT.

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 will be 0.
- If the highest severity message is a warning (ANSnnnnW), then the return code will be 8.
- If the highest severity message is an error (ANSnnnnE or ANSnnnnS), then the return code will be 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 will be 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 Table 26.

Table 26. 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: <ul style="list-style-type: none"> • 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” on page 103.
8	The operation completed with at least one warning message. For scheduled events, the status will be <i>Completed</i> . 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 will be <i>Failed</i> . 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 from being processed yields return code 12. When a file is not found the operation yields return code 12.
other	For scheduled operations where the scheduled action is COMMAND, the return code will be the return code from the command that was executed. If the return code is 0, the status of the scheduled operation will be <i>Completed</i> . If the return code is nonzero, then the status will be <i>Failed</i> . <p>Some commands might issue a nonzero return code to indicate success. For these commands, you can avoid a <i>Failed</i> status by <i>wrapping</i> 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.</p>

The return code for a client macro will be the highest return code issued among the individual commands that comprise the macro. For example, suppose a macro consists of these commands:

```
selective t9\sys:u/devel/* -subdir=yes  
incremental t9\sys:u/devel/TestDriver/*" -subdir=yes  
archive t9\sys: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 will be 8.

Also see "Handling return codes from preschedulecmd and postschedulecmd scripts" on page 338.

Chapter 8. Understanding 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.

This chapter 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.

Using 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.

Using 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.

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 (or bound) to a file with a management class, Tivoli Storage Manager uses the default management class in the active policy set.
- For 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 retention period.
- 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. See “Selecting a management class for files” on page 104 for more information. In your Preferences file, you can associate folders with a management class, using the *dirmc* option. See “Selecting a management class for folders” on page 106 for more information.

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 folders.
- How many backup versions the server keeps of your files and folders.
- How long the server keeps backup versions of your files and folders.
- 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

When the server is unable to rebind a file to an appropriate management class, the server uses one of two values to determine the number of days to retain the file. If it is a backup version, the server uses *backup retention grace period*. 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. For more information about grace periods, see “Using a retention grace period” on page 107.

Displaying information about management classes and copy groups

Before you select the management classes you want to use, click **View policy information** from the Utilities menu. The **Display policy information** window is displayed. You can then determine which management classes are available.

The **Display policy information** window provides the following information:

- The name of the default management class.
- The name of the policy domain to which the management class belongs.
- The policy set that is currently active.
- The date and time that this policy set became active.
- The number of backup versions which are maintained for files which still exist on your workstation.
- The number of backup versions which are maintained for files which have been deleted from your workstation.
- The number of days to keep inactive backup versions.
- The number of days to keep the last backup version.
- The management class name and a description.

You can also use the *detail* option on the **query mgmtclass** command to view the available management classes.

Table 27 shows the default values for the backup and archive copy groups in the standard management class. Each attribute is discussed in more detail immediately following the table.

Table 27. Default 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

Copy group name

The name of the copy group. The default value for both backup and archive is *Standard*.

Copy type

The type of copy group. The value for backup is always *Backup*, and the value for archive is always *Archive*.

Copy frequency

Copy frequency 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 folder 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.

Versions data exists

The *Versions Data Exists* attribute specifies the maximum number of different backup versions retained for files and folders currently on your volume. 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

The *Versions Data Deleted* attribute specifies the maximum number of different backup versions retained for files and folders that you erased from your volume. This parameter is ignored as long as the file or folder remains on your volume.

If you erase the file or folder, the next time you run an incremental backup, the active backup version is changed to inactive and the oldest versions are erased that exceed 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

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

The *Retain Only Version* attribute specifies the number of days the last remaining inactive version of a file or folder 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 will be deleted from the server in 10 days.

Copy serialization

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 folder 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 folder 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 Preferences file. If the object is changed during every backup or archive attempt, it is not backed up or archived.
- **Dynamic.** A file or folder is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.
- **Shared dynamic.** A file or folder is backed up or archived regardless of whether it changes during a 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 Preferences file without the file changing during the attempt. The file is backed up or archived on the last try even if it has changed.

Attention: Be careful when you select a management class containing a copy group that specifies shared dynamic or dynamic for serialization backup. 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 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.

If you restore or retrieve a file that contains a fuzzy copy, the file might not be usable. You should not use dynamic or shared dynamic serialization to back up files, unless you are absolutely certain that a restore of a fuzzy copy will be usable.

Copy mode

The *Copy Mode* attribute determines whether a file or folder 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 meta-data changes (such as access permissions), but the file data does not change, Tivoli Storage Manager might back up only the meta-data.
- **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

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

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.

Selecting 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.
- Do the copy groups specify either static or shared static for serialization?
If serialization is shared dynamic or dynamic, you might get fuzzy backups or archive copies. Verify that this is acceptable. For example, you might want to use shared dynamic or dynamic serialization for a file to which log records are continuously added. If you used static or shared static serialization, the file might never be backed up because it is constantly in use. With shared dynamic or dynamic serialization, the file is backed up, but the backup version of the file might contain inconsistent data. Do not use shared dynamic or dynamic serialization for a file if it is very important that the backup version or archive copy contain all changes.

- 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?

Assigning 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 folders, use the *dirmc* option in your dsm.sys. See “Dirmc” on page 152 for more information.

You can assign a management class for a file or file group by using an **include** statement in your Preferences file. Management class names are not case-sensitive.

Using the command-line client, to associate the file /Users/jones/budget.jan with the management class budget, you would enter:

```
include /Users/jones/budget.jan budget
```

To specify a management class named **managall** to use for all files to which you do not explicitly assign a management class, enter:

```
include /* managall
```

The example below demonstrates how to use a management class:

```
exclude /*sno
include "/Volumes/La Pomme/winter/*Ice"      mcweekly
include "/Volumes/La Pomme/winter/December/*Ice" mcdaily
include "/Volumes/La Pomme/winter/January/*Ice" mcmonthly
include "/Volumes/La Pomme/winter/winter/white sno"
```

Processing follows these steps:

1. The file name white sno is backed up following bottom-up processing rules. Because you did not specify a management class, the file is assigned to the default management class.
2. Any file ending with ice in the /Volumes/La Pome/winter/January folder is assigned to the management class, **mcmonthly**.
3. Any file ending with ice in the /Volumes/La Pome/winter/December folder is assigned to the management class, **mcdaily**.
4. Any files ending with ice in any directory under /Volumes/La Pomme/winter/December are assigned to the management class, **mcweekly**.
5. Any file ending with sno (except (for /Volumes/La Pomme/winter/winter/white sno) in any directory is excluded from backup.

To specify your own default management class for files that are not explicitly included, specify:

```
include /* mgmt_class_name
```

as the first include or exclude option defined. See “Include options” on page 186 for more information about the *include* option.

Overriding the management class for archived files

When you archive a file, you can override the assigned management class using the 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 `/Users/van/Documents/budget.jan` with the management class **ret2yrs**, you would enter:

```
dsmc archive -archmc=ret2yrs /Users/van/Documents/budget.jan
```

Selecting a management class for folders

If the management class in your active policy set containing the longest "Retain only version" (RETONLY) setting meets your backup requirements for folders, it might not be necessary to take any action to associate folders with that management class. Tivoli Storage Manager does it automatically when it backs up your folders.

If there is more than one management class with the longest RETONLY setting, then the Tivoli Storage Manager client selects the management class whose name is first or 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 will need to choose a different management class. You should keep directories at least as long as you keep the files associated with those folders.

For backup directories, use the *dirmc* option to specify the management class to which directories will be 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 : 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.

Binding 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 the 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 file's inactive objects 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.

Rebinding 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.

Using a 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 used when:

- 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.

Using 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 will be 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` will remain on the server forever, and `event.file` will be expired 70 days after it was created (60 days after its event occurred). However, let's say that 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 will be expired in 30 days (the hold does not extend the retention period).

Archiving 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 will be deleted. If the files were archived on a data retention server, the same command will fail both files. The data retention server forces the user to keep archives until the stated retention criteria are met.

Now let's see 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.

Refer to “Set Event” on page 332 for more information about the **set event** command.

See the *IBM Tivoli Storage Manager Administrator's Guide* for your operating system, for more information about RETINIT.

Chapter 9. Using processing options

You can use defaults for processing options or you can tailor the processing options to meet your specific needs. This chapter:

- Provides an overview of processing options.
- Includes an options reference section that provides detailed information about each option. See “Using options with commands” on page 123 for more information.

Overview of processing options

Tivoli Storage Manager uses *processing options* that you specify in your dsm.sys or client user options file (dsm.opt) to control communications, backup-archive processing, and other types of processing.

This section provides an overview of the following types of options that you can use:

- Communication options
- Server and node options
- Backup and archive processing options
- Restore and retrieve processing options
- Scheduling options
- Format and language options
- Authorization options
- Error processing options
- Transaction processing option
- Diagnostics options

See Chapter 2, “Configuring Tivoli Storage Manager,” on page 13 for information on how to create and modify your dsm.sys or client user options file (dsm.opt) file.

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. For a complete list of command-line options, a description, and where to go in this book for more information, see Table 39 on page 124.

Communication options

You use communication options to specify how your client node communicates with a Tivoli Storage Manager server.

Macintosh only supports the TCP/IP communication protocol. Use the *commmethod* option to specify the TCP/IP communication protocol. For more information, see “Commmethod” on page 140.

Ask your Tivoli Storage Manager administrator for assistance in setting your communication options.

TCP/IP options

To use the TCP/IP communication protocol, you must include the *tcpserveraddress* option in your (dsm.sys). The other TCP/IP options have default values that you can modify if you want to change the default value.

Table 28. TCP/IP options

Option	Description	Page
<i>tcpbuffsize</i>	Specifies the size, in kilobytes, of the Tivoli Storage Manager internal TCP/IP communication buffer.	252
<i>tcpnodelay</i>	Specifies whether the server or client disables the delay of sending successive small packets on the network.	256
<i>tcpadminport</i>	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.	251
<i>tcpport</i>	Specifies the TCP/IP port address for a Tivoli Storage Manager server.	257
<i>tcpserveraddress</i>	Specifies the TCP/IP address for a Tivoli Storage Manager server.	258
<i>tcpwindowsize</i>	Specifies the size, in kilobytes, of the TCP/IP sliding window for your client node.	259
<i>webports</i>	Enables the use of the scheduler program outside a firewall by specifying the TCP/IP port number used by the Client Acceptor daemon.	269

Server and node Node options

Use the following options to specify the server to contact for backup-archive services, and the client node for which to request backup-archive services.

Server options

Use the *servername* option in your (dsm.sys) 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 148 for more information.

Use the *servername* option in the client user options file (dsm.opt) to specify a server to contact for backup-archive services. This overrides the default server specified in your (dsm.sys).

Figure 3 shows a sample dsm.sys.

```

DEFAULTServer          server2

SErvername      server1
  NODename      node1
  COMMMethod    TCPip
  TCPPort       1500
  TCPServeraddress node.domain.company.com
  PASSWORDAccess generate
  MAILprog      /usr/bin/xsend root
  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

```

Figure 3. Sample client system options file

Node options

Table 29. Server and node options

Option	Description	Page
<i>defaultserver</i>	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.	148
<i>asnodename</i>	Use the <i>asnodename</i> 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.	129
<i>nodename</i>	Use the <i>nodename</i> option in your dsm.sys file to identify your workstation to the server to establish communications.	202
<i>servername</i>	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.	240
<i>virtualnodename</i>	The <i>virtualnodename</i> option specifies the node name of your workstation when you want to restore or retrieve files to a different workstation.	268

Backup and archive processing options

You can use the following options to control some aspects of backup and archive processing.

Table 30. Backup and archive processing options

Option	Description	Page
<i>archmc</i>	Use the <i>archmc</i> option with the archive command to specify the available management class for your policy domain to which you want to bind your archived files.	128
<i>asnodename</i>	Use the <i>asnodename</i> 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.	129
<i>autofsrename</i>	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.	135
<i>changingretries</i>	Specifies the number of retries when attempting to back up or archive a file that is in use.	138
<i>compressalways</i>	The <i>compressalways</i> option specifies whether to continue compressing an object if it grows during compression. Use this option with the <i>compression</i> option.	143
<i>compression</i>	The <i>compression</i> option compresses files <i>before</i> you send them to the server. Compressing your files reduces data storage for backup versions and archive copies of your files.	144
<i>deletefiles</i>	Use the <i>deletefiles</i> option with the archive command to delete files from your workstation after you archive them.	149
<i>description</i>	The <i>description</i> option assigns or specifies a description for files when performing archive, delete, retrieve, query archive or query backupset operations.	150
<i>detail</i>	Use the <i>detail</i> option to display management class, file space, backup, and archive information depending on the command with which it is used.	151
<i>dirmc</i>	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.	152
<i>dirsonly</i>	Backs up, restores, archives, retrieves, or queries directories <i>only</i> .	153
<i>diskcachelocation</i>	Specifies the location where the disk cache database will be created if the option <i>memoryefficient=diskcachemethod</i> is set during an incremental backup.	156

Table 30. Backup and archive processing options (continued)

Option	Description	Page
<i>domain</i>	Specifies the volumes to include in your default client domain for an incremental backup.	157
<i>enablearchiveretentionprotection</i>	Allows the client to connect to a data retention server.	160
<i>exclude</i> <i>exclude.backup</i> <i>exclude.file</i> <i>exclude.file.backup</i>	Use these options to exclude a file or group of files from backup services.	168
<i>encryptiontype</i>	Allows you to select AES 128-bit data encryption, providing a stronger form of data encryption than DES 56-bit data encryption.	161
<i>encryptkey</i>	Specifies whether to save the encryption key password locally when performing a backup-archive operation or whether to prompt for the encryption key password.	162
<i>exclude.archive</i>	Excludes a file or a group of files that match the pattern from archive services <i>only</i> .	168
<i>exclude.attribute.symlink</i>	Excludes a file or a group of files that are symbolic links or aliases from backup processing <i>only</i> .	168
<i>exclude.compression</i>	Excludes files from compression processing if you set the <i>compression</i> option to <i>yes</i> . This option applies to backups and archives.	168
<i>exclude.dir</i>	Excludes a directory, its files, and all its subdirectories and their files from backup processing.	168
<i>exclude.encrypt</i>	Excludes specified files from encryption processing.	168
<i>encryptiontype</i>	Specifies whether to use AES 128-bit or DES 56-bit data encryption during backup and archive operations. AES 128-bit encryption is the default.	161
<i>filelist</i>	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.	171
<i>filesonly</i>	Backs up, restores, retrieves, or queries files <i>only</i> .	174
<i>guitreeviewafterbackup</i>	Specifies whether the GUI client is returned to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.	180
<i>incl excl</i>	Specifies the path and file name of an include-exclude options file.	184
<i>include</i> <i>include.backup</i> <i>include.file</i>	Use these options to include files or assign management classes for backup processing.	186
<i>include.archive</i>	Includes files or assigns management classes for archive processing.	186
<i>include.compression</i>	Includes files for compression processing if you set the <i>compression</i> option to <i>yes</i> . This option applies to backups and archives.	186

Table 30. Backup and archive processing options (continued)

Option	Description	Page
<i>include.encrypt</i>	Includes the specified files for encryption processing. By default, Tivoli Storage Manager does not perform encryption processing.	186
<i>include.fs</i>	Use the <i>include.fs</i> option to control how Tivoli Storage Manager processes your file space for incremental backup.	186
<i>incrbydate</i>	Use with the incremental command to request an incremental backup by date.	190
<i>memoryefficientbackup</i>	Specifies a memory-saving backup algorithm for incremental backups when used with the incremental command.	199
<i>noprompt</i>	suppresses the confirmation prompt that is presented by the delete group , delete archive , expire , restore image , and set event commands.	203
<i>optfile</i>	Specifies the alternative options file you want to use when you start a Tivoli Storage Manager session.	205
<i>removeoperandlimit</i>	Specifies that Tivoli Storage Manager removes the 20-operand limit for the Mac OS X platform. If you specify the <i>removeoperandlimit</i> 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.	222
<i>skipacl</i>	Specifies whether to skip ACL processing completely.	244
<i>skipaclupdatecheck</i>	Specifies whether to perform checksum and size comparisons before and after backup and during incremental processing.	245
<i>subdir</i>	Specifies whether to include subdirectories of a named directory.	247
<i>tapeprompt</i>	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.	249
<i>type</i>	Use the <i>type</i> option with the query node command to specify the type of node to query.	264
<i>v2archive</i>	Use the <i>v2archive</i> 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.	266

Restore and retrieve processing options

The following options relate to restore and retrieve processing.

Table 31. Restore and retrieve processing options

Option	Description	Page
<i>dirsonly</i>	Qualifies the operation (backup, archive, restore, retrieve) to process directories alone.	153
<i>disablenqr</i>	Specifies whether the Tivoli Storage Manager Backup-Archive Client can use the no-query restore method for restoring files and directories from the server.	154
<i>filelist</i>	Specifies a file containing a list of files to be processed by the specified command.	171
<i>filesonly</i>	Qualifies the operation (backup, archive, restore, retrieve) to process files alone.	174
<i>fromdate</i>	Use the <i>fromdate</i> option with the <i>fromtime</i> option to specify a date and time from which you want to search for backups or archives during a restore, retrieve, or query operation.	175
<i>fromnode</i>	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.	176
<i>fromowner</i>	Displays file spaces for an alternate owner. Also specifies an alternate owner from which to restore or retrieve files.	177
<i>fromtime</i>	Use the <i>fromtime</i> option with the <i>fromdate</i> option to specify a beginning time from which you want to search for backups or archives during a restore, retrieve or query operation.	178
<i>guitreeviewafterbackup</i>	Specifies whether the GUI client is returned to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.	180
<i>ifnewer</i>	Replaces an existing file with the latest backup version only if the backup version is newer than the existing file.	182
<i>inactive</i>	Displays a list of active and inactive files when used with the <i>pick</i> option.	183
<i>latest</i>	Restores the most recent backup version of a file whether it is active or inactive.	192
<i>localbackupset</i>	Specifies whether the Tivoli Storage Manager GUI bypasses initial logon with the server to restore a local backup set on a standalone workstation.	193
<i>location</i>	Specifies where Tivoli Storage Manager searches for the backup set during a query or restore operation.	194
<i>noprompt</i>	suppresses the confirmation prompt that is presented by the delete group , delete archive , expire , restore image , and set event commands.	203
<i>optfile</i>	Specifies the you want to use when you start a Tivoli Storage Manager session.	205

Table 31. Restore and retrieve processing options (continued)

Option	Description	Page
<i>pick</i>	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 <i>inactive</i> option to view both active and inactive objects.	210
<i>pitdate</i>	Use the <i>pitdate</i> option with the <i>pittime</i> option to establish a point in time for which you want to display or restore the latest version of your backups.	211
<i>pittime</i>	Use the <i>pittime</i> option with the <i>pitdate</i> option to establish a point in time for which you want to display or restore the latest version of your backups.	212
<i>preservepath</i>	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.	217
<i>replace</i>	Specifies whether to overwrite an existing file, or to prompt you for your selection when you restore or retrieve files.	223
<i>subdir</i>	Specifies whether you want to include subdirectories of a named directory.	247
<i>tapeprompt</i>	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.	249
<i>todate</i>	Use the <i>todate</i> option with the <i>totime</i> 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.	261
<i>totime</i>	Use the <i>totime</i> option with the <i>todate</i> 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.	261
<i>type</i>	Use the <i>type</i> option with the query node command to specify the type of node to query.	264

Scheduling options

You can use the following options to regulate central scheduling. Tivoli Storage Manager uses scheduling options only when the Scheduler is running.

Table 32. Scheduling options

Option	Description	Page
<i>managedservices</i>	Specifies whether the Storage Manager Client Acceptor daemon manages the scheduler program.	196
<i>maxcmdretries</i>	Specifies the maximum number of times the client scheduler attempts to process a scheduled command that fails.	198
<i>postschedulecmd</i> , <i>postnschedulecmd</i>	Specifies a command to process after running a schedule.	213

Table 32. Scheduling options (continued)

Option	Description	Page
<i>preschedulecmd, prenschedulecmd</i>	Specifies a command to process before running a schedule.	215
<i>queryschedperiod</i>	Specifies the number of hours the client scheduler waits between attempts to contact the server for scheduled work.	220
<i>retryperiod</i>	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.	228
<i>schedcmddisabled</i>	Specifies whether to disable the scheduling of generic commands specified by your Tivoli Storage Manager administrator.	230
<i>schedcmduser (server defined only)</i>	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.	96
<i>schedlogmax</i>	Specifies the maximum size of the scheduler log, in megabytes.	231
<i>schedlogname</i>	Specifies the path and file name where you want to store schedule log information.	232
<i>schedlogretention</i>	Specifies the number of days to keep log file entries in the schedule log, and whether to save pruned entries.	233
<i>schedmode</i>	Specifies which schedule mode to use, <i>polling</i> or <i>prompted</i> .	
<i>schedrestretrdisabled</i>	Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing restore or retrieve schedule operations.	237
<i>sessioninitiation</i>	Use the <i>sessioninitiation</i> option to control whether the server or client initiates sessions through a firewall. The default is that the client can initiate sessions.	242
<i>svrprepostscheddisabled</i>	Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing pre-schedule and post-schedule commands when performing scheduled operations.	246
<i>tcpclientaddress</i>	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 will use this address when it begins the server prompted scheduled operation. See <i>schedmode prompted</i> ("Schedmode" on page 235) for details.	254
<i>tcpclientport</i>	Specifies a TCP/IP port number for the server to contact the client when the server begins the server prompted scheduled operation. See <i>schedmode prompted</i> ("Schedmode" on page 235) for details.	255

Format and language options

You can use the following options to select different formats for date, time, numbers, and for different languages if you have the appropriate client language pack installed for that language.

Table 33. Format and language options

Option	Description	Page
<i>dateformat</i>	Specifies the format for displaying dates.	146
<i>language</i>	Specifies the language used for messages.	191
<i>numberformat</i>	Specifies the format for displaying numbers.	204
<i>timeformat</i>	Specifies the format for displaying time.	260

Command processing options

The following options apply when you use Tivoli Storage Manager commands.

Table 34. Command processing options

Option	Description	Page
<i>editor</i>	Specifies if the command-line interface editor and command retrieve capability is turned on or off.	159
<i>quiet</i>	Limits the number of messages that are displayed on your screen during processing. This option can be overridden by the server.	221
<i>scrolllines</i>	Specifies the number of lines of information that are displayed on your screen at one time. Use this option only when <i>scrollprompt</i> is set to <i>yes</i> .	238
<i>scrollprompt</i>	Specifies whether you want Tivoli Storage Manager to stop and wait after displaying the number of lines of information you specified with the <i>scrolllines</i> option, or scroll through and stop at the end of the information list.	239
<i>verbose</i>	Specifies that processing information should be displayed on your screen. The alternative is <i>quiet</i> . This option can be overridden by the server.	267

Authorization options

These options control access to a Tivoli Storage Manager server.

Table 35. Authorization options

Option	Description	Page
<i>groups</i>	Specifies the groups on your workstation that you want to authorize to request Tivoli Storage Manager services from the server.	179
<i>mailprog</i>	Specifies the program and user ID where you want to send a newly-generated password when the old one expires.	195
<i>password</i>	Specifies a Tivoli Storage Manager password.	206

Table 35. Authorization options (continued)

Option	Description	Page
<i>passwordaccess</i>	Specifies whether you want to use a generated password or be prompted for a password each time you start the client.	207
<i>passworddir</i>	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.	209
<i>users</i>	Authorizes specific users on your workstation to request services from a server.	265

Error processing options

These options specify the name of the error log file and how Tivoli Storage Manager treats the entries in the log file.

Table 36. Error processing options

Option	Description	Page
<i>errorlogmax</i>	Specifies the maximum size of the error log, in megabytes.	164
<i>errorlogname</i>	Specifies the fully qualified path and file name of the file where you want to store information about errors that occur during processing.	165
<i>errorlogretention</i>	Specifies how many days to maintain error log entries before pruning, and whether to save the pruned entries.	166

Transaction processing options

These options control how Tivoli Storage Manager processes transactions between the client and server.

Table 37. Transaction processing options

Option	Description	Page
<i>collocatebyfilespec</i>	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 <i>collocatebyfilespec</i> option to <i>yes</i> 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).	139
<i>commrestartduration</i>	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.	141

Table 37. Transaction processing options (continued)

Option	Description	Page
<i>commrestartinterval</i>	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.	142
<i>diskbuffsize</i>	Specifies the maximum disk I/O buffer size (in kilobytes) that the client can use when reading files.	155
<i>largecommbuffers</i>	<p>This option has been replaced by the <i>diskbuffsize</i> option. At this time, <i>largecommbuffers</i> will continue to be accepted by the Tivoli Storage Manager client in order to ease the transition to the new option. However, the value specified by <i>largecommbuffers</i> will be ignored in favor of the <i>diskbuffsize</i> setting.</p> <p>Recommendation: Discontinue the use of <i>largecommbuffers</i> because future releases of Tivoli Storage Manager might not accept this option.</p>	
<i>nfstimeout</i>	Specifies the number of seconds the server waits for a status system call on an NFS file system before it times out.	201
<i>resourceutilization</i>	Use the <i>resourceutilization</i> option in your dsm.sys file to regulate the level of resources the Tivoli Storage Manager server and client can use during processing.	225
<i>txnbytelimit</i>	Specifies the number of kilobytes the client program buffers before it sends a transaction to the server.	263

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. See “Query Systeminfo” on page 313 for more information.

Table 38. Diagnostics options

Option	Description	Page
<i>console</i>	Use the <i>console</i> option with the query systeminfo command to output system information to the console.	145
<i>filename</i>	Use the <i>filename</i> option with the query systeminfo command to specify a file name in which to store the system information.	173

Using options with 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 in this book for more information, see Table 39 on page 124.

Entering options with a command

Follow these general rules to enter 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,

```
dsmc archive -description="year end 2002" /Users/agordon/Documents/proj1.rtf
```

- 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 quotes (') or quotation marks ("). Most operating system command-line processors strip the quotes before submitting the command-line arguments to the Tivoli Storage Manager client application. In such cases, using escape characters or doubling the quotes allows the client to receive the quoted object name. In loop mode, surround such objects in either single quotes (') 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. For information about how to read the syntax diagrams, see "Reading syntax diagrams" on page x.
- Enter options before or after command parameters. For example, you can enter the option before or after a file specification:

```
dsmc selective -subdir=yes "/Users/user1/Documents/*"  
dsmc selective "/Users/user1/Documents/*" -subdir=yes
```

- When entering several options on a command, separate them with a blank space.
- Enclose the value in quotes (" ") if the option value that you enter contains a blank space. For example:

```
dsmc archive -description="Project A" "/Users/user1/Documents/*"
```

- 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 preferences file rather than overriding the current value.
- The maximum number of characters for a file name is 256. The maximum combined length of the file name and path name is 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. Long file name support is provided while running Tivoli Storage Manager on Mac OS X.

Table 39 lists client command options that you can enter only on the command line with specific commands.

Table 39. Client command options

Command option	Description	Commands	Page
<i>archmc</i>	Use the <i>archmc</i> 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	128
<i>console</i>	Use the <i>console</i> option with the query systeminfo command to output system information to the console.	query systeminfo	145
<i>deletefiles</i>	Deletes the local copy of files from your workstation after they are archived on the server.	archive	149
<i>description</i>	Assigns or specifies a description for files when performing archive, delete, retrieve, or query archive operations.	archive delete archive query archive query backupset retrieve	150
<i>detail</i>	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	151
<i>dirsonly</i>	Backs up, restores, archives, retrieves, or queries directories <i>only</i> .	archive incremental query archive query backup restore restore backupset retrieve selective	153
<i>filelist</i>	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 delete archive delete backup expire incremental query archive query backup restore retrieve selective	171
<i>filename</i>	Use the <i>filename</i> option with the query systeminfo command to specify a file name in which to store the system information.	query systeminfo	173

Table 39. Client command options (continued)

Command option	Description	Commands	Page
<i>filesonly</i>	Backs up, restores, retrieves, or queries files <i>only</i> .	archive incremental query archive query backup restore restore backupset retrieve selective	174
<i>fromdate</i>	Use the <i>fromdate</i> option with the <i>fromtime</i> 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 retrieve	175
<i>fromnode</i>	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.	query archive query backup query filespace query mgmtclass restore retrieve	176
<i>fromowner</i>	Displays file spaces for an alternate owner. Also specifies an alternate owner from which to restore or retrieve files.	query archive query backup restore retrieve	177
<i>fromtime</i>	Specifies a beginning time on the specified date. Use with the <i>fromdate</i> option. This option is ignored if the <i>fromdate</i> option is absent.	query archive query backup restore retrieve	178
<i>ifnewer</i>	Replaces existing files with the latest backup version only if the backup version is newer than the existing version.	restore restore backupset retrieve	182
<i>inactive</i>	Displays a list of active and inactive files when used with the <i>pick</i> option.	query backup restore	183
<i>incrbydate</i>	Requests an incremental backup by date.	incremental	190
<i>latest</i>	Restores the most recent backup version of a file whether it is active or inactive.	restore	192
<i>location</i>	Specifies whether Tivoli Storage Manager searches for a backup set on the during a query or restore operation.	query backupset restore backupset	194
<i>noprompt</i>	suppresses the confirmation prompt that is presented by the delete group , delete archive , expire , restore image , and set event commands.	delete archive delete backup expire	203
<i>optfile</i>	Specifies the you want to use when you start a Tivoli Storage Manager session.		205
<i>pick</i>	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 <i>inactive</i> option to view both active and inactive objects.	delete archive delete backup expire restore retrieve	210
<i>pitdate</i>	Use the <i>pitdate</i> option with the <i>pittime</i> option to establish a point in time for which you want to display or restore the latest version of your backups.	query backup restore	211

Table 39. Client command options (continued)

Command option	Description	Commands	Page
<i>pittime</i>	Use the <i>pittime</i> option with the <i>pitdate</i> option to establish a point in time for which you want to display or restore the latest version of your backups.	query backup restore	212
<i>preservepath</i>	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 retrieve	217
<i>removeoperandlimit</i>	Specifies that Tivoli Storage Manager removes the 20-operand limit for the Mac OS X platform. If you specify the <i>removeoperandlimit</i> 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	222
<i>todate</i>	Use the <i>todate</i> option with the <i>totime</i> 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 retrieve	261
<i>totime</i>	Use the <i>totime</i> option with the <i>todate</i> 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 retrieve	262
<i>type</i>	Use the <i>type</i> option with the query node command to specify the type of node to query.	query node	264
<i>v2archive</i>	Use the <i>v2archive</i> 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	266

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 40. Options that are valid on the initial command line only

Options valid on the command line

<i>commmethod</i>	<i>schedlogname</i>
<i>diskbuffsize</i>	<i>schedlogretention</i>
<i>editor</i>	<i>schedmode</i>
<i>errorlogname</i>	<i>servername</i>
<i>errorlogretention</i>	<i>sessioninitiation</i>
<i>language</i>	<i>tapeprompt</i>
<i>maxcmdretries</i>	<i>tcpbuffsize</i>
<i>nodename</i>	<i>tcpclientaddress</i>
<i>optfile</i>	<i>tcpclientport</i>
<i>password</i>	<i>tcpport</i>
<i>postschedulecmd, postnschedulecmd</i>	<i>tcpserveraddress</i>
<i>preschedulecmd, prenschedulecmd</i>	<i>tcpwindowsize</i>
<i>queryschedperiod</i>	<i>txnbytelimit</i>
<i>resourceutilization</i>	<i>virtualnodename</i>
<i>retryperiod</i>	

Client options reference

The following sections contain detailed information about each of the Tivoli Storage Manager processing options. Information for each option includes:

- A description of the option.
- A syntax diagram of the option. The option name contains uppercase and lowercase characters. The uppercase characters indicate the minimum abbreviation you can use for the option name. See “Reading syntax diagrams” on page x for an explanation of these diagrams.
- Detailed descriptions of the option parameters. If the parameter is a constant (a value that does not change), use the minimum abbreviation.
- Examples of using the option in the Preferences file.

Notes:

1. For options with a **yes** parameter, acceptable alternatives are **1**, **true**, and **on**. For options with a **no** parameter, acceptable alternatives are **0**, **false**, and **off**.
2. 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"
```

cr

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 graphical user interface (GUI). Overriding the management class using the GUI is equivalent to using the *archmc* option on the **archive** command.

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.

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
```

Asnodename

Authorized User

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.

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 /Users stored under the node name MyCluster:

```
dsmc incremental /Users -asnodename=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*. Also, *asnodename* can be used for clustered systems, although no specific cluster software is supported.

Preferences File

Place this option in the *dsm.sys* file *within* a server stanza. You can set this option on the **General** category 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 which 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 /Users/spike/test/*
04/21/07 15:25:05 ANS1651I Backed Up:
    /Users/spike/test/file.txt
04/21/07 15:25:05 ANS1652I Archived:
    /Users/spike/test/file.txt
04/21/07 15:25:05 ANS1653I Updated:
    /Users/spike/test/file.txt
04/21/07 15:25:05 ANS1654E Failed:
    /Users/spike/test/file.txt
04/21/07 15:25:05 ANS1655I Restored:
    /Users/spike/test/file.txt
04/21/07 15:25:05 ANS1656I Retrieved:
    /Users/spike/test/file.txt
04/21/07 15:25:05 ANS1657I Expired:
    /Users/spike/test/file.txt
04/21/07 15:25:05 ANS1658I Deleted:
    /Users/spike/test/file.txt
04/21/07 15:25:05 ANS1659I Skipped:
    /Users/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:
    /Users/spike/test/file.txt
04/21/07 15:25:05 ANS1661I Unchanged:
    /Users/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 will be 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.

Both the Windows and Macintosh Backup-Archive Clients will create the audit log as a Unicode file.

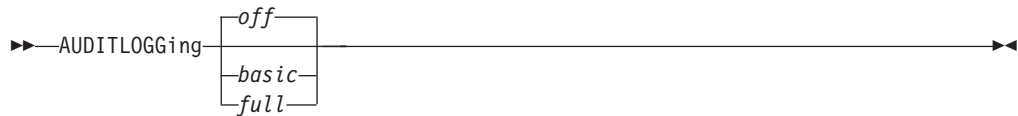
By default, the name of the audit log is `dsmaudit.log` and it is contained in the same directory as the error log, `dsmerror.log`. The name and location of the audit log can be configured using the *auditlogname* option. There are no parameters to control the size of the audit log or to prune the audit log. The *auditlogname* option cannot be set as an option in a Tivoli Storage Manager Server client options set.

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 will be disabled for the rest of the operation and the return code for the operation will be set to 12, regardless of the outcome of the operation.

Preferences File

Place this option in the `dsm.sys` or `dsm.opt` file.

Syntax



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. For more information about the audit logging facility refer to “Auditlogging” on page 131.

Preferences File

Place this option in the dsm.sys or dsm.opt file.

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 /users/* -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: MAC
Session established with server SPIKE: Windows
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 '/users/*'
Directory--> 4,096 /users/test/dir1
/file1 [Sent]
Normal File--> 32,768 /users/test/dir1
/file2 [Sent]
Normal File--> 32,768 /users/test/dir1/
file3 [Sent]
Normal File--> 32,768 /users/test/dir2/
file1 [Sent]
Normal File--> 32,768 /users/test/dir2/
file2 [Sent]
Normal File--> 32,768 /users/test/dir2/
file3 [Sent]
Successful incremental backup of '/users/test/*'
```

```

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: /users/*
07/03/07 12:05:15 ANS1651I Backed Up:
/users/test
07/03/07 12:05:15 ANS1651I Backed Up:
/users/test/dir1
07/03/07 12:05:15 ANS1651I Backed Up:
/users/test/dir2
07/03/07 12:05:15 ANS1651I Backed Up:
/users/test/file1
07/03/07 12:05:15 ANS1651I Backed Up:
/users/test/file2
07/03/07 12:05:15 ANS1651I Backed Up:
/users/test/file3
07/03/07 12:05:15 ANS1651I Backed Up:
/users/test/dir1/file1
07/03/07 12:05:15 ANS1651I Backed Up:
/users/test/dir1/file2
07/03/07 12:05:15 ANS1651I Backed Up:
/users/test/dir1/file3
07/03/07 12:05:15 ANS1651I Backed Up:
/users/test/dir2/file1
07/03/07 12:05:15 ANS1651I Backed Up:
/users/test/dir2/file2
07/03/07 12:05:15 ANS1651I Backed Up:
/users/test/dir2/file3

```

Autofsrename

Authorized User

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.

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.

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. For more information on migrating to Unicode-enabled file spaces, see “Migrating to the Unicode-enabled client” on page 4.

To restore or retrieve from a file space that is not Unicode-enabled, specify the source on the server and the destination on the client.

The server can define the *autofsrename* option and override the *autofsrename* setting on the client.

Preferences File

Place this option in the client options file (dsm.opt) file. You can set this option on the **General** category, **Rename non-Unicode files** during backup/archive drop-down list box of the Preferences editor.

Syntax



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
```

Backupsetname

Authorized User

The *backupsetname* option specifies either the name of the backup set, or the name of the file or tape device containing the backup set. This option is used in conjunction with the *location* option.

You can use *backupsetname* with the following commands:

- **query backup**
- **query backupset**
- **query filespace**
- **restore**
- **restore backupset**

Preferences File

None. You can specify this option only on the command line.

Syntax

►► —BACKUPSetname *backupsetname*
bset filename ◄◄

Parameters

backupsetname

Specifies the name of the backup set on the server. The *location* option should be set to *server*.

bset filename

Specifies the name of the file containing the beginning of the backup set. The *location* option should be set to *file*.

Examples

Command line:

```
dsmc restore /Volumes/bkSets/file.1 -backupsetname=YEAR_END_ACCOUNTING.1234
```

Changingretries

Authorized User

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.

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.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Backup** category, **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
```

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).

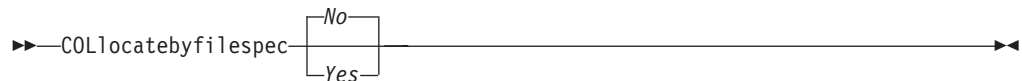
Considerations:

- 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.

Preferences File

Place this option in the client user options file (dsm.opt).

Syntax



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 will be stored on multiple tapes. Generally, the files specified in the file specification will still be 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

Authorized User

The *commmethod* option specifies the communication method you use to provide connectivity for client-server communication.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Communication** category, **Communication Method** list of the Preferences editor. TCP/IP is the only choice.

Syntax



Parameters

TCPIP

The Transmission Control Protocol/Internet Protocol (TCP/IP) communication method. This is the default.

V6Tcpi

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.

Examples

Options file:

```
commmethod tcpip
```

Use only TCP/IP Version 4

```
commmethod V6Tcpi
```

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 V6Tcpi* are specified.

Command line:

```
-comm=tcpip
```

```
-comm=V6Tcpi
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Commrestartduration

Authorized User

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.

Note: A scheduled event will continue if the client reconnects with the server before the *commrestartduration* value elapses, even if the event's startup window has elapsed.

You can use the *commrestartduration* option and the *commrestartinterval* in busy or unstable network environments to decrease connection failures.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Communication** category, **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

Authorized User

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.

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.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Communication** category, **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.

Use the *compressalways* option with the **archive**, **incremental**, and **selective** commands.

Preferences File

Place this option in the client user options file (dsm.opt). You can set this option on the **Backup** category, **Continue Compressing if Object Grows** check box of the Preferences editor.

Syntax



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.

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

Authorized User

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.

Use the *compression* option with the **archive**, **incremental**, and **selective** commands.

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 to exclude specific files or groups of files from compression processing. See “Exclude options” on page 168 for more information.
- Use the *include.compression* option in your to include files within a broad group of excluded files for compression processing. See “Include options” on page 186 for more information.

This option controls compression *only if* your administrator specifies that your client node can compress files before sending them to the server.

Preferences File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Backup** category, **Compress objects** check box of the Preferences editor.

Syntax



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.

Console

Use the *console* option with the **query systeminfo** command to output information gathered from one or more of the following items 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
- POLICY - Policy set dump.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually dmsched.log).

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, it is recommended that you 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. See “Filename” on page 173 for more information.

Syntax

►►—CONSOLE—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
query systeminfo dsmdptfile errorlog -console
```

Dateformat

The *dateformat* option specifies the format you want to use to display dates.

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.

Notes:

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**
- **restore**
- **retrieve**
- **set event**

Preferences File

Place this option in the client user options file (dsm.opt). You can set this option on the **Regional Settings** category, **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:

- 1 MM/DD/YYYY
- 2 DD-MM-YYYY
- 3 YYYY-MM-DD
- 4 DD.MM.YYYY
- 5 YYYY.MM.DD

Examples

Options file:

```
dateformat 3
```

Command line:

```
-date=3
```

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 will contact 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 dsm.opt file.

Preferences 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.

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.

Syntax

►►—DEletefiles—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc archive "/Users/dgordon/Documents/*.c" -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" /Users/van/Documents/budget.jan
```

Notes:

1. The maximum length of a description is 254 characters.
2. Enclose the value in quotes (" ") 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**

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 will be 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"
```

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"
```

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 client GUI.

Use the *detail* option with the **query backup** and **query archive** commands to display the last modification date and the of the file you specify.

- **delete filespace**
- **query archive**
- **query backup**
- **query filespace**
- **query mgmtclass**

Syntax

►►—DETail—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc query mgmtclass -detail
dsmc query filespace -detail
```

Dirmc

Authorized User

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. It is recommended that you select a management class for individual directories that retains directories at least as long as it retains the files associated with them.

Note: If you want to back up specific files to a management class see “Assigning a management class to files” on page 105 for more information.

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 you back up and does not effect 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.

The server can also define this option.

Preferences File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Backup** category, **Directory Management Class** section in the Preferences editor.

Syntax

►►—DIRMc— *mgmtclassname* —————◄◄

Parameters

mgmtclassname

Specifies the name of the management class you want to associate with directories. The client uses the management class name that you specify for all 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.

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**

Syntax

►►—Dirsonly—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc query backup -dirsonly "/Users/*"
```

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. Refer to "Standard query restore, no query restore, and restartable restore" on page 72 for more information about 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"). Refer to "Standard restore process" on page 73 for more information about the standard restore process.

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. This option is also valid with the Java GUI and with the command-line interface. The Tivoli Storage Manager client API does not support this option.

Preferences File

Place this option in the dsm.opt file.

Syntax



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

Authorized User

The *diskbuffsize* option specifies the maximum disk I/O buffer size (in kilobytes) that the client can use when reading files.

Note: The *diskbuffsize* option replaces the *largecommbuffers* option.

Optimal backup, archive 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 will require more memory and might not improve performance.

Recommendation: Use the default setting, unless otherwise directed by IBM support personnel.

Preferences File

Place this option in the dsm.sys file.

Syntax

►►—DISKBuffsize— *size* —————►◄

Parameters

size Specifies the maximum disk I/O buffer size (in kilobytes) that the client will use when reading files. The range of values is 16 through 1023; the default is 32.

Examples

Options file:

diskbuffsize 64

Command line:

Does not apply.

Diskcachelocation

Authorized User

The *diskcachelocation* option specifies the location where the disk cache database will be created if the option *memoryefficientbackup=diskcachemethod* is set during an incremental backup. You can specify the *diskcachelocation* option in your *dsm.sys* file within a server stanza or with the *include.fs* option. See “Include options” on page 186 for more information. If the *diskcachelocation* option appears in the option file, its value will be used for all file systems not represented by an *include.fs* option containing 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.

Attention: For performance reasons, do not use a remote drive for *diskcachelocation*.

The disk cache file created by subsequent disk cache incremental backups should require considerably less disk space. The actual amount of disk space required for the disk cache file created by subsequent disk cache incremental backups depends on the number of files and directories included in the backup and on the length of the longest file or directory name which is backed up. The disk cache file created by the initial disk cache incremental backup can require up to 5 gigabytes of disk space for each million files or directories being backed up.

Preferences 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 will be 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 the volumes that you want to include for incremental backup in your client domain.

Use the *domain* option in your `dsm.sys` file or the client user options file (`dsm.opt`) to define your default client domain. The server can also define this option. Tivoli Storage Manager uses your default client domain in the following situations to determine which volumes to process during an incremental backup:

- When you run an incremental backup using the **incremental** command and you do not specify which volumes to process.
- When your administrator defines a schedule to run an incremental backup for you, but does not specify which volumes to process.
- When you select the **Backup Domain** action from the Tivoli Storage Manager GUI .

If you do not use the *domain* option to specify volumes in your client options file, Tivoli Storage Manager uses the *all-local* parameter as the default.

When you use the *domain* option with the **incremental** command, Tivoli Storage Manager adds volumes that you specify to the volumes defined in your client options file. For example, if you enter the following in your client options file:

```
domain /
```

and the following on the command line:

```
dsmc incremental -domain=/Volumes/puma
```

Tivoli Storage Manager performs an incremental backup of the startup disk and `/Volumes/puma`, the `puma` volume.

If you use both a file specification and the *domain* option with the **incremental** command, Tivoli Storage Manager ignores the *domain* option and processes only those volumes that you specify in the file specification. For example, if you enter:

```
dsmc incremental /Volumes/jaguar -domain=/Volumes/puma
```

Tivoli Storage Manager performs an incremental backup for the `/Volumes/jaguar` volume only.

You can also exclude volumes by specifying the dash (-) operator before the volumes. For example, in the following option Tivoli Storage Manager will process all local volumes except for the `jaguar` volume:

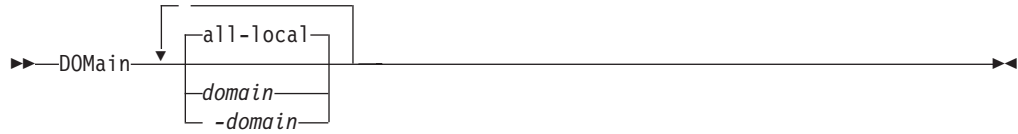
```
domain ALL-LOCAL -/Volumes/jaguar
```

Note: You cannot use the (-) operator in front of a domain keyword such as `ALL-LOCAL`.

Preferences File

Place this option in the `dsm.sys` file or the client user options file (`dsm.opt`). You can set this option on the **Backup** category, **Domain for Backup** section of the Preferences editor.

Syntax



Parameters

all-local

Backs up all local volumes. This is the default.

domain

Defines the volumes to include in your default client domain.

A volume name must be enclosed in quotes if it includes any spaces.

-domain

Defines the volumes to exclude in your default client domain.

Examples

Options file:

```
domain all-local
domain all-local -/Volumes/volume2
domain all-local -/Volumes/volume2 '-/Volumes/Macintosh HD'
```

Command line:

```
-domain="/ /Volumes/volume2"
-domain="all-local -/Volumes/volume2"
```

Editor

The *editor* option turns the command-line interface (CLI) editor and retrieve capability on or off.

Because certain terminal settings result in non-standard input, the editor might not work as defined in this document. In this case, you can disable the editor by setting the *editor* option to *no* in your dsm.sys file *or* client user options file (dsm.opt).

Preferences File

Place this option in the dsm.sys file *or* the client user options file (dsm.opt).

Syntax



Parameters

Yes Turns on the CLI editor and command retrieve capability. This is the default.

No Turns off the CLI editor and command retrieve capability.

Examples

Options file:

```
editor yes
```

Command line:

```
-editor=yes
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Enablearchiveretentionprotection

The *enablearchiveretentionprotection* option allows the client to connect to a Tivoli Storage Manager data retention server to ensure 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 will be refused.

The data retention server is specially configured for that 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 will be 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**)

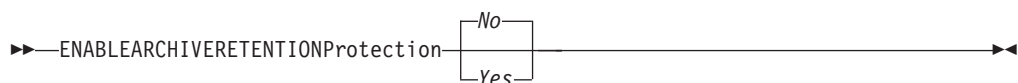
Note: **restore backupset -location=file** or **-location=tape** do not connect to any server (except the virtual one) and thus will not be blocked under any circumstances.

- 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

Preferences 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



Parameters

- No* The data retention server connection is refused. This is the default.
- Yes* The client connects to a data retention server.

Encryptiontype

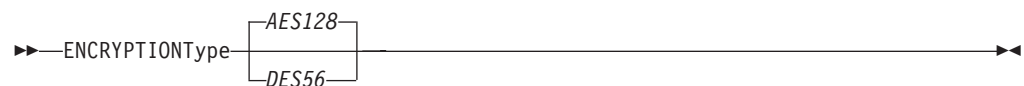
The *encryptiontype* option allows you to use AES 128-bit data encryption, providing a stronger form of data encryption than DES 56-bit data encryption. The encryption type only affects backup and archive operations. The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received. During restore and retrieve operations, the encrypted data will be decrypted with the proper encryption algorithm, regardless of the setting for this option.

Note: Files that are backed up using AES 128-bit encryption cannot be restored with an older level client. However, if you ever need to reinstall and use an older client, files backed up by the Tivoli Storage Manager Version 5.3 client, using 56-bit DES encryption, can be restored with an older client.

Preferences File

Place this option in the dsm.sys file. You can also set this option on the **Authorization** tab of the Preferences editor. The server can override this.

Syntax



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 aes128

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 will be 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.

Attention: When using the prompt option, your encryption key is not saved in the Windows Registry or the Tivoli Storage Manager password file on UNIX. If you forget the key, your data will be unrecoverable.

Supported Clients

This option is valid for all clients.

Preferences File

Place this option in the client system options file (dsm.sys). You can set this option on the **Authorization** category, **Encryption Key Password** section of the Preferences editor.

Syntax



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. This key can be up to 63 bytes in length. The key is retrieved from the password file on restore and retrieve operations.

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.
- The default for *encryptkey* is *save*, and the default for *passwordaccess* is *prompt*. However, if you specify *encryptkey=save*, you must also specify *passwordaccess generate*, when you specify include encrypt.

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 key is issued when restoring or retrieving the encrypted file. This key can be up to 63 bytes in length.

| **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

Authorized User

The *errorlogmax* option specifies the maximum size of the error log, in megabytes.

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*.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **General** category, **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

Authorized User

This option specifies the fully qualified path and file name of the file where error messages are written.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **General** category, **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 default file name is dsmerror.log and it is placed one of the following locations:

~/Library/Logs/tivoli/tsm/
/Library/Logs/tivoli/tsm/

Examples

Options file:

Command line:

-errorlogname=/Users/user1/tsmerror.log

This option is valid only on the initial command line. It is not valid in interactive mode.

Errorlogretention

Authorized User

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.

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.

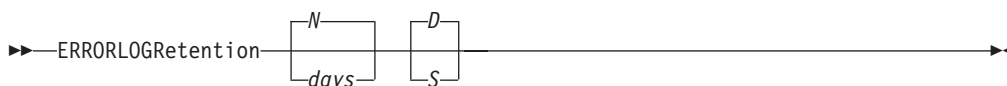
If you change from *errorlogmax* to *errorlogretention*, all existing log entries are retained and the log is pruned using the new *errorlogretention* criteria.

Restriction: You cannot specify *errorlogretention* and a non-zero *errorlogmax* value.

Preferences File

Place this option in the `dsm.sys` file. You can set this option on the **General** category, **Select Error Log** button of the Preferences editor.

Syntax



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.

Exclude options

Authorized User

The exclude options exclude objects from backup or archive services. For example, you might want to exclude all temporary files, any local caches of network files, all files that contain compiled object code that you can easily reproduce using other methods, or your operating system files.

You can exclude specific files from encryption processing during a backup.

Notes:

1. 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.
3. As with other include-exclude statements, you can use the *inclexcl* option to specify a file that can be in Unicode format, containing exclude statements with file names in Unicode. See “Inclexcl” on page 184 for more information.

Exclude any system files that could corrupt the operating system when recovered. You should also exclude the directory containing the Tivoli Storage Manager client files.

Attention: See “Excluding system files” on page 30 for a list of files that you should always exclude.

Use wildcard characters to exclude a broad range of files. See “Including and excluding groups of files” on page 31 for a list of wildcard characters that you can use. Then, if necessary, use the *include* option to make exceptions.

To exclude an entire directory called /Volumes/User1/any/test, enter the following:

```
exclude.dir /Volumes/User1/any/test
```

To exclude subdirectories that begin with "test" under the any directory, enter the following:

```
exclude.dir ../../any/test*
```

Preferences File

Place these options in the dsm.sys file. You can set these options on the **Include-Exclude** category, **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.

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 from backup processing only.

exclude.compression

Excludes files from compression processing if the *compression* option is set to *yes*. This option applies to backups and archives.

exclude.dir

Excludes a directory, its files, and all its subdirectories and their files from backup processing. For example, the statement `exclude.dir /Users/dan/test/data1` excludes the `/Users/dan/test/data1` directory, its files, and all its subdirectories and their files.

The next time that you perform an incremental backup, the previous backup versions will be inactive. If you exclude a directory that was previously included, Tivoli Storage Manager marks existing backup versions of the files and directories beneath it inactive during the next incremental backup. Use this option to exclude a portion of your data that has no underlying files to back up.

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.

Parameters

pattern

Specifies the file or group of files that you want to exclude.

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.

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.encrypt /users/home2/file1
exclude.compression /home/gordon/proj1/*
exclude.attribute.symlink ../../*
```

Command line:

Does not apply.

Controlling symbolic link and alias processing

Tivoli Storage Manager treats symbolic links and aliases 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 /Volumes/Users/user1/Documents directory, enter these statements in your dsm.sys file:

```
exclude.attribute.symlink ../../*
include.attribute.symlink /Volumes/Users/user1/Documents/../../*
```

See “Include options” on page 186 for more information about the *include.attribute.symlink* option.

Controlling compression processing

If you want to exclude specific files or groups of files from compression processing during a backup or archive operation, consider the following:

- 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. See “Compression” on page 144 for more information.

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.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 /Users/jones/proj1/*.*
exclude.compression /Users/jones/proj1/file.txt
include /Users/jones/proj1/file.txt
```

Tivoli Storage Manager examines the statements (reading from bottom to top) and determines that the /Users/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.

Filelist

Use the *filelist* option with the following commands to process a list of files:

- **archive**
- **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 quotes.
- The filelist can be an MBCS file or a Unicode file with all Unicode entries.
- Tivoli Storage Manager ignores any entry that is not valid.

The following is an example of a list of files within a filelist:

```
/Users/mike/Documents/file1
/Users/van/Documents/file2
/Users/avi/Documents/file3
/Volumes/fs1/dir2/file3
"/Volumes/fs2/Ha Ha Ha/file.txt"
/Volumes/fs3/file.txt
```

If an entry in the filelist indicates a directory, only that directory will be 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 **restore** command:

```
restore -filelist=/Users/van/Documents/file3 /Users/mike/Documents/
```

the file specification `/Users/mike/Documents/` represents the restore destination for all entries in the filelist. However, in the **selective** command:

```
selective -filelist=/Users/Documents/file3 /Users/mike/Documents/
```

the file specification `/Users/mike/Documents/` is ignored.

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 `/Users/van/proj/myfile` and `/Users/van/proj`. If the `proj` directory does not exist on the server, the `/Users/van/proj` directory is sent to the server a second time.

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:

`selective -filelist=/Users/avi/Documents/filelist.txt`

Filename

Use the *filename* option with the **query systeminfo** command to specify a file name in which to store information gathered from one or more of the following items:

- DSMOPTFILE - The contents of TSM User Preference 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
- POLICY - Policy set dump.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually dsmsched.log).

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, it is recommended that you 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. See “Console” on page 145 for more information.

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
```

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**
- **retrieve**
- **selective**

Syntax

►►—FILESOnly—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc incremental -filesonly
```

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**
- **retrieve**

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/*"
```

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 mgmtclass**
- **restore**
- **retrieve**

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/*"
```

Fromowner

The *fromowner* option specifies an alternate owner from which to restore backup versions or archived files. The owner must give access to another to use the files. For example, to restore files from the `/Users/devel/proja` directory belonging to *usermike* on system *puma*, and place the restored files in a directory you own named `/Users/id/proja`, enter:

```
dsmc restore -fromowner=usermike -fromnode=puma /Users/devel/proja/  
/Users/id/proja/
```

Note: If you specify the *fromowner* option without the *fromnode* option, the active user must be on the same node as the *fromowner* user.

See “How Tivoli Storage Manager handles long user and group names” on page 54 for information about the character limits for group names.

Use the *fromowner* option with the following commands:

- **query archive**
- **query backup**
- **restore**
- **retrieve**

Syntax

►►—FROMOwner == *owner*—◄◄

Parameters

owner

Name of an alternate owner.

Examples

Command line:

```
dsmc query archive "/Users/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**
- **retrieve**

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 "/Users/*"
```

Groups

Authorized User

The *groups* option specifies 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.

See “How Tivoli Storage Manager handles long user and group names” on page 54 for information about the character limits for group names.

Preferences File

Place this option in the dsm.sys file *within* a server stanza.

Syntax



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.

Preferences 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** category, **Return to tree window after function completed** check box of the Preferences editor.

Syntax



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.

Preferences File

Place this option in the client system options file (dsm.sys). You can set this option on the **Web Client** category → **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:

Does not apply.

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**
- **retrieve**

Syntax

►►—IFNewer—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/Users/grover/Documents/*" -sub=y -rep=y -ifnewer
```

Inactive

Use the *inactive* option with the following commands to display both active and inactive objects:

- **query backup**
- **restore**

Recommendation: 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 will be restored in an indeterminate order. This option is implicit when *pitdate* is used.

Syntax

►►—INActive—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/Users/zoe/Documents/*" -inactive -pick
```

Incl excl

Authorized User

The *incl excl* option specifies the path and file name of an include-exclude options file.

Multiple incl excl 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.

When processing occurs, the include-exclude statements within the include-exclude file are placed in the list position occupied by the *incl excl* option, in the same order, and processed accordingly.

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 Macintosh system cannot contain Japanese characters.

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 Plain Text**.
2. Enter your *include* and *exclude* statements using the appropriate include-exclude options in “Using include-exclude options” on page 28.
3. Click **File** and then click **Save As**. The **Save As** window displays.
4. From **Plain Text Encoding**, select **Unicode (UTF-16)**, specify the file and target directory, and then save the file.
5. Place an *incl excl* option specifying the include-exclude file you just created in your client options file (dsm.opt).
6. Restart the Tivoli Storage Manager client.

For more information about creating an include-exclude options file, see “Creating an include-exclude list (optional root user or authorized user task)” on page 26.

Preferences File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Include-Exclude** category of the Preferences editor.

Syntax

►► INCL EXCL *filespec* ◄◄

Parameters

filespec

Specifies the path and file name of *one* include-exclude options file.

Examples

Options file:

```
INCLExc1 /Users/user1/Documents/backup.exc1
```

Command line:

Does not apply.

Include options

Authorized User

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.
- 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 or use the *diskcachelocation* option to cause specific file systems to use different, specific locations for their diskcache.

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 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.

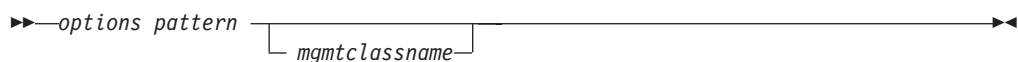
Notes:

1. The *exclude.dir* statement overrides all include statements that match the pattern.
2. The server can also define these options with the *inclexcl* option.

Preferences File

Place these options in the `dsm.sys` file. You can set these options on the **Include-Exclude** category, **Define Include-Exclude Options** section of the Preferences editor.

Syntax



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 the following example, the *archmc* management class is assigned when an archive operation is performed, because *include.backup* is used only for backup processing, it is not used for archive processing.

```
include.archive /Users/dan/Documents/* archmc
include.backup /Users/dan/Documents/*
```

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.

include.compression

Includes files for compression processing if you set the *compression* option to *yes*. This option applies to backups and archives.

include.encrypt

Includes the specified files for encryption processing. By default, Tivoli Storage Manager does not perform encryption processing.

Attention: 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.fs

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 will be 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 following *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
/kalafsl snapshotproviderfs=JFS2
```

Parameters

pattern

Specifies the objects to include for backup or archive processing or to assign a specific management class.

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.

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 <file of dir spec> <mgmtcname>
include /myvfs/mygroupname ManagementClass
```

Examples

Options file:

```
include /Users/proj/text/devel.*
include /Users/proj/text/* textfiles
include * managall
include.archive /Users/proj/text/* myarchiveclass
include.backup /Users/proj/text/* mybackupclass
include.compression /Users/proj/text/devel.*
include.encrypt /Users/proj/gordon/*
include.attribute.symlink /Users/spike/.../*
```

Command line:

Does not apply.

Controlling symbolic link and alias processing

Tivoli Storage Manager treats symbolic links and aliases 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 /Users/user1/Documents directory, enter these statements in your dsm.sys file:

```
exclude.attribute.symlink /.../*
include.attribute.symlink /Users/user1/Documents/.../*
```

See “Exclude options” on page 168 for more information about the *exclude.attribute.symlink* option.

Compression and encryption processing

If you want to include specific files or groups of files for compression and encryption processing during a backup or archive operation, consider the following:

- 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. See “Compression” on page 144 for more information.
- Tivoli Storage Manager processes *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 /Users/jones/proj1/file.txt
include.compression /Users/jones/proj1/file.txt
include.encrypt /Users/jones/proj1/file.txt
```

Tivoli Storage Manager examines the *exclude /Users/jones/proj1/file.txt* statement first and determines that */Users/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*.

- As with other include-exclude statements, you can use the *inclexcl* option to specify a file that is in Unicode format, which contains *include.compression* and *include.encrypt* statements specifying Unicode files. See “Inclexcl” on page 184 for more information.

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. Files added at the client after the last incremental backup, but with a modification date earlier than the last incremental, are not backed up.

Attention: 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 will back up these files again.

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, 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.

Syntax

►►—INCRbydate—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc incremental -incrbydate
```

Language

The *language* option specifies the national language in which to present client messages.

You can use US English (ENU) with all clients. Currently, the Macintosh client only supports US English (ENU).

Preferences File

Place this option in the client options file (dsm.opt). You can set this option on the **Regional Settings** category, **Language** drop-down list of the Preferences editor.

Syntax

►►—LANGUage— *language* —————►◄

Parameters

language
Specifies the language you want to use.

Examples

Options file:
language enu

Command line:
Does not apply.

Latest

Use the *latest* option with the following commands to restore the most recent backup version of a file, even if the backup is inactive:

- **restore**

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.

Syntax

►►—LATEST—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/Users/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. You can place this option in the `dsm.opt` or `dsm.sys` file.

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.

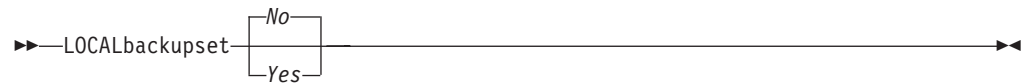
If you set the *localbackupset* option to *no* (the default), the GUI attempts initial logon with the server and enables all GUI functions.

Note: The **restore backupset** command supports restore of local backup sets on a standalone workstation *without* using the *localbackupset* option. See “Restore Backupset” on page 320 for more information.

Preferences File

Place this option in the `dsm.sys` file.

Syntax



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`

Location

The *location* option specifies where Tivoli Storage Manager searches for a backup set during a query or restore operation. You can use this option to locate backup sets on the server or locally. Backup sets that are generated on the server can be used locally specifying the *location* option and the file name of the file containing the start of the backup set.

Use the *location* option with the **query** and **restore** commands.

Syntax



Parameters

server

Specifies that Tivoli Storage Manager searches for the backup set on the server. This is the default.

file Specifies that Tivoli Storage Manager searches for the backup set on a local file.

Examples

Command line:

```
restore backupset mybackupsetname -loc=server
restore backupset /Users/budget/backupsetfile.name
-lc=file
```

Mailprog

Authorized User

The *mailprog* option passes the information that is specified in the option value to the operating system to execute.

Preferences File

Place this option in the dsm.sys file *within* a server stanza.

Syntax

►►MAILprog *command*◄◄

Parameters

command

Indicates how you want to invoke a mail program.

Examples

Options file:

mailprog /usr/bin/mail myusername

Command line:

Does not apply.

Managedservices

Authorized User

The *managedservices* option specifies whether the Tivoli Storage Manager client acceptor daemon manages the scheduler.

See “Configuring the client scheduler” on page 21 for instructions to set up the client acceptor daemon to manage the scheduler.

Attention: You cannot use the *dsmcad* for scheduling when you set the *sessioninitiation* option to *serveronly*. Refer to “Sessioninitiation” on page 242 for more information.

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.

Notes:

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 will connect and disconnect to 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. Set the *passwordaccess* option to *generate* in your *dsm.sys* file and generate a password, so Tivoli Storage Manager can manage your password automatically. See “Passwordaccess” on page 207 for more information.

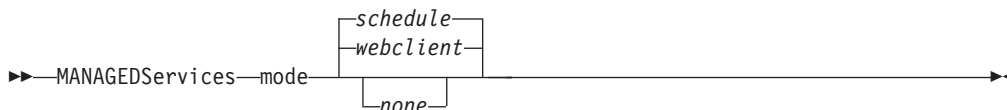
Using the client acceptor daemon to manage the scheduler service can provide the following benefits:

- Scheduled events can occur without a user being logged into the system.
- 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.
- By default, if you do not specify the *managedservices* option, the client acceptor daemon manages the scheduler program.
- To use the client acceptor daemon, you must specify this option in the client system options file.

Preferences 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** category of the Preferences editor.

Syntax



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.

schedule

Specifies that the client acceptor daemon manages the scheduler. This is the default.

none

Specifies that the client acceptor daemon not manage the Web client or schedules. Set *manageservices none* to enable the **dsmc schedule** command.

Examples

Options file:

The following are examples of how you might specify the *manageservices* option in your client system options file (dsm.sys).

Task Specify that the client acceptor daemon manages the Web client *only*.

```
manageservices webclient
```

Task Specify that the client acceptor daemon manages the scheduler *only*.

```
manageservices schedule
```

Task Specify that the client acceptor daemon manages both the Web client and the scheduler.

```
manageservices schedule webclient
```

Note: The order in which these values are specified is not important.

Command line:

Does not apply.

Maxcmdretries

Authorized User

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.

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.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Scheduler** category, **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. See “Include options” on page 186 for more details. If the *memoryefficientbackup* option appears in the option file, its value will be used for all file systems not represented by an *include.fs* option containing the *memoryefficientbackup* parameter. If a directory-level incremental is being performed, and *memoryefficientbackup=diskcachemethod* is present, *memoryefficientbackup=yes* will be used.

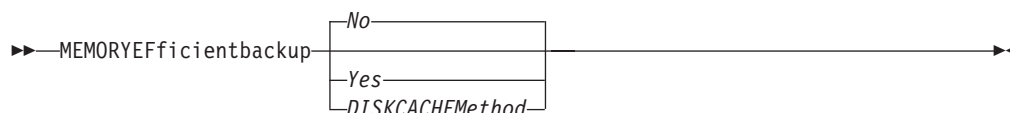
Use *memoryefficientbackup=diskcachemethod* for any file space that has too many files for Tivoli Storage Manager to complete the incremental backup with either the default setting, *memoryefficientbackup=no*, or with *memoryefficientbackup=yes*. The disk cache file created by the initial disk cache incremental backup can require up to 5 gigabytes of disk space for each million files or directories being backed up.

The disk cache file created by subsequent disk cache incremental backups should require considerably less disk space. The actual amount of disk space required for the disk cache file created by subsequent disk cache incremental backups depends on the number of files and directories included in the backup and on the length of the longest file or directory name which is backed up. Since the disk cache file can become very large, ensure that large file support is enabled on the file system that is being used for the disk cache file.

Preferences File

Place this option in the client user options file (dsm.opt), the client system options file (dsm.sys) within a server stanza, or on the initial command line. You can also set this option on the **Performance Tuning** category, **Use memory-saving algorithm** checkbox of the Preferences editor.

Syntax



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
```

Nfstimeout

Authorized User

The *nfsttimeout* 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 will be 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.

Preferences 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.

Examples

Options file:

```
nfsttimeout 10
```

Command line:

```
-nfsttimeout=10
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Nodename

Authorized User

Use the *nodename* option in your dsm.sys file to identify your workstation to the server. You can use different node names to identify multiple operating systems on your workstation.

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. See “Virtualnodename” on page 268 for more information.

When connecting to a server, the client must identify 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 user options file (dsm.opt), 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.

Preferences File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **General** category, **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 will permit the node name to default to the host name of the workstation

Examples

Options file:

nodename cougar

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**
- **expire**

Syntax

►►—NOPrompt—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc delete archive -noprompt "/Users/van/Documents/*"
```

Numberformat

The *numberformat* option specifies the format you want to use to display numbers.

Use this option if you want to change the default number format for the language of the message repository you are using.

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.

You can use the *numberformat* option with the following commands:

- **delete archive**
- **delete backup**
- **expire**
- **query archive**
- **query backup**
- **restore**
- **retrieve**
- **set event**

Preferences File

Place this option in the client user options file (dsm.opt). You can set this option on the **Regional Settings** category, **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.

1 1,000.00

2 1,000,00

3 1 000,00

4 1 000.00

5 1.000,00

6 1'000,00

Examples

Options file:

num 4

Command line:

-numberformat=4

Optfile

The *optfile* option specifies the Preferences file you want to use when you start a Tivoli Storage Manager session.

Syntax

►►—OPTFILE == *file_name*—————►◄

Parameters

file_name

Specifies an alternate preferences 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 /Library/Preferences/Tivoli Storage Manager/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="/Users/tsm/Prefs Alpha"
```

client acceptor daemon:

```
dsmcad -optfile=/Users/admin/Library/Preferences/tsm/mytsm.opt
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Password

Authorized User

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.

Notes:

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 will be 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*.

Preferences 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

Authorized User

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.

If a password is required, you can choose to:

- 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.

For information on where the password is stored, see “Passworddir” on page 209.

When the *passwordaccess* option is set to *generate* and you specify the *password* option, the *password* option is ignored.

- When using the Web client.

Preferences File

Place this option in the `dsm.sys` file *within* a server stanza. You can set this option on the **Authorization** category, **Password Access** section of the Preferences editor.

Syntax



Parameters

prompt

You are prompted for your workstation 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.

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.

Passworddir

Authorized User

The *passworddir* option specifies the directory location in which to store an encrypted password file. The default directory is `"/Library/Preferences/Tivoli Storage Manager"`. Regardless of where it is stored, the password file created by Tivoli Storage Manager is always named `TSM.PWD`.

Preferences File

Place this option in the `(dsm.sys)`. You can set this option on the **Authorization** category, **Location** section of the Preferences editor.

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.

For non-root users, if the Tivoli Storage Manager executable is owned by a non-root user and the owner execution permission bit is set to `s`, the default directory is the directory specified by the `DSM_DIR` (or `DSMI_DIR` for the API) environment variable. If `DSM_DIR` is not set, the default directory is the client installation directory.

Examples

Options file:

```
passworddir "/Users/user1/Library/Preferences/Tivoli Storage Manager/"
```

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.

Use the *pick* option with the following commands:

- **delete archive**
- **delete backup**
- **expire**
- **restore**
- **retrieve**

Syntax

►►—Pick—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/Users/van/Documents/*" -pick -inactive
```

Pitdate

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. Files that were backed up *on or before* the date and time you specified, and which were not deleted *before* the date and time you specified, 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**
- **restore**

When *pitdate* is used, the *inactive* and *latest* options are implicit.

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
```

Pitttime

Use the *pitttime* 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. 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 *pitttime* option with the following commands:

- **delete backup**
- **query backup**
- **restore**

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*, *pitttime*, and *totime* options.

Examples

Command line:

```
dsmc query backup -pitt=06:00:00 -pitd=08/01/2003  
"/Volumes/proj5/myproj/*"
```

Postschedulecmd/Postnschedulecmd

Authorized User

The *postschedulecmd* 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.

If you do not want to wait, specify *postnschedulecmd*.

Notes:

1. If the *postschedulecmd* command does not complete with return code 0, the client will report 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).

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Scheduler** category, **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 executed 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 quotes.

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.

Otherwise, update the command string to contain the full path of the file, such as:

"/Volumes/La Pomme/Scripting/Database Script"

If the *postschedulecmd* schedule command is an AppleScript, you must use the *osascript* command to run the script. For example, if "*Database Script*" is an AppleScript, enter this command:

```
postschedulecmd osascript "/Volumes/La Pomme/Scripting/  
Database Script"
```

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"
```

Command line:

```
-postschedulecmd="/Volumes/La Pomme/Scripting/postsched.sh"
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Preschedulecmd/Prenschedulecmd

Authorized User

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*.

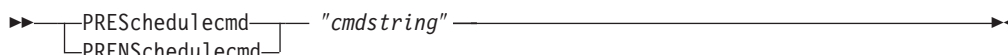
Notes:

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 will report that the scheduled event failed, and the return code will be 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).

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Scheduler** category, **Schedule Command** button of the Preferences editor.

Syntax



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 quotes.

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.

The command string should contain the full path of the file, such as:

```
"/Volumes/La Pomme/Scripting/Database Script"
```

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"
```

If the *preschedulecmd* 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:

```
preschedulecmd "/Volumes/La Pomme/Scripting/presched.sh"
```

Examples

Options file:

```
preschedulecmd "<your database product's quiesce command>  
database"  
preschedulecmd "/Volumes/La Pomme/Scripting/presched.sh"
```

The command string is a valid command for quiescing your database.

Command line:

```
-preschedulecmd="/Volumes/La Pomme/Scripting/presched.sh"
```

This option is valid only on the initial command line. It is not valid in interactive mode.

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**
- **retrieve**

Syntax



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 *=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:

```
/Volumes/fs/h1/m1/file.a
/Volumes/fs/h1/m1/file.b
/Volumes/fs/h1/m1/l1/file.x
/Volumes/fs/h1/m1/l1/file.y
```

This command:

```
dsmc res /Volumes/fs/h1/m1/ /Volumes/devel/ann/ -preser=complete
```

Restores these directories and files:

/Volumes/devel/ann/h1/m1/file.a
/Volumes/devel/ann/h1/m1/file.b

This command:

```
dsmc res /Volumes/fs/h1/m1/ /Volumes/devel/ann/ -preser=nobase
```

Restores these directories and files:

/Volumes/devel/ann/file.a
/Volumes/devel/ann/file.b

This command:

```
dsmc res backupset /Volumes/fs/h1/m1/ /Volumes/devel/ann/ -su=yes  
-preser=nobase -loc=file
```

Restores these directories and files:

/Volumes/devel/ann/file.a
/Volumes/devel/ann/file.b
/Volumes/devel/ann/file.x
/Volumes/devel/ann/file.y

This command:

```
dsmc res /Volumes/fs/h1/m1/ /Volumes/devel/ann/ -preser=subtree
```

Restores these directories and files:

/Volumes/devel/ann/m1/file.a
/Volumes/devel/ann/m1/file.b

This command:

```
dsmc res /Volumes/fs/h1/m1/ /Volumes/devel/ann/ -preser=none
```

Restores these directories and files:

/Volumes/devel/ann/file.a
/Volumes/devel/ann/file.b

This command:

```
dsmc res /Volumes/fs/h1/m1/ /Volumes/devel/ann/ -su=yes -preser=complete
```

Restores these directories and files:

/Volumes/devel/ann/h1/m1/file.a
/Volumes/devel/ann/h1/m1/file.b
/Volumes/devel/ann/h1/m1/l1/file.x
/Volumes/devel/ann/h1/m1/l1/file.y

This command:

```
dsmc res /Volumes/fs/h1/m1/ /Volumes/devel/ann/ -su=yes -preser=nobase
```

Restores these directories and files:

/Volumes/devel/ann/file.a
/Volumes/devel/ann/file.b
/Volumes/devel/ann/l1/file.x
/Volumes/devel/ann/l1/file.y

This command:

```
dsmc res /Volumes/fs/h1/m1/ /Volumes/devel/ann/ -su=yes -preser=subtree
```

Restores these directories and files:

/Volumes/devel/ann/m1/file.a
/Volumes/devel/ann/m1/file.b
/Volumes/devel/ann/m1/l1/file.x
/Volumes/devel/ann/m1/l1/file.y

This command:

```
dsmc res /Volumes/fs/h1/m1/ /Volumes/devel/ann/ -su=yes -preser=none
```

Restores these directories and files:

```
/Volumes/devel/ann/file.a  
/Volumes/devel/ann/file.b  
/Volumes/devel/ann/file.x  
/Volumes/devel/ann/file.y
```

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. 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.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Scheduler** category, **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.

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.

Note: The server can also define the **quiet** option, overriding the client setting.

Preferences File

Place this option in the client user options file (dsm.opt). You can set this option on the **Command Line** category, **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 for the Mac OS X platform. 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 for the Mac OS X platform.

Notes:

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 for the Mac OS X platform is removed.
3. Because it adversely affects performance to allow the shell to expand wild cards, it is recommended that you 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.

Syntax

►►—REMOVEOPerandlimit—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
-removeoperandlimit
```


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 will be the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

Resourceutilization

Authorized User

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.

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. The client is also restricted to a single session if a user who is not an authorized user invokes a Macintosh 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:

```
inc /Volumes/filespaceA /Volumes/filespaceB
```

and you specify *resourceutilization=5*, the client might start a second session to query files on file space B. 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). See “Collocatebyfilespec” on page 139 for more information.

Regulating restore sessions

When you request a restore, the default is to use a maximum of one session. Additional restore sessions are based on the following factors:

- 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

Notes:

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 will be used for the restore. If you increase the *resourceutilization* setting to 5, then 5 sessions will be 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.

Considerations

The following factors can affect the throughput of multiple sessions:

- The server's ability to handle multiple client sessions. Is there sufficient memory, multiple storage volumes, and CPU cycles to increase backup throughput?
- The client's ability to drive multiple sessions (sufficient CPU, 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.

Note: The server can also define this option.

Preferences File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **General** category, **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 8

Command line:

`-resourceutilization=7`

This option is valid only on the initial command line. It is not valid in interactive mode.

Retryperiod

Authorized User

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.

Your administrator can also set this option. If your administrator specifies a value for this option, that value overrides the value in your file after your client node successfully contacts the server.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Scheduler** category, **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.

Preferences File

Place this option in the client user options file (dsm.opt) *or* the dsm.sys file. You can set this option on the **Web Client** category of the Preferences editor.

Syntax



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

Authorized User

The *schedcmddisabled* option specifies whether to disable the scheduling of commands by the server *action=command* option on the **define schedule** server command.

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. See “Query Schedule” on page 311 for more information.

Note: The server can also define this option.

Preferences File

Place this option in the dsm.sys file *within* a server stanza.

Syntax



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.

Schedlogmax

Authorized User

The *schedlogmax* option specifies the maximum size of the schedule 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.

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 pruned log, *dsmsched.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*.

Preferences File

Place this option in the *dsm.sys* file. You can set this option on the **Scheduler** category, **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

Authorized User

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 TSM Schedule Log file is created in the same directory as the dsmerror.log file. See “Errorlogname” on page 165 for more information on placement of the TSM Schedule Log file.

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

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Scheduler** category, **Schedule Log** button of the Preferences editor.

Attention: Set the SCHEDLOGNAME option 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 default folder. The default directories are:

~/Library/Logs/tivoli/tsm
/Library/Logs/tivoli/tsm

The **TSM Schedule Log** file *cannot* be a symbolic link.

Examples

Options file:

SCHEDLOGN /Users/user1/Library/Logs/schedlog.jan

Command line:

-schedlogname=/Users/user1/Library/Logs/schedlog.jan

This option is valid only on the initial command line. It is not valid in interactive mode.

Schedlogretention

Authorized User

The *schedlogretention* option specifies the number of days to keep entries in the schedule 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.

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.

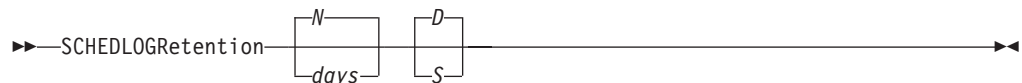
If you change from *schedlogmax* to *schedlogretention*, all existing log entries are retained and the log is pruned using the new *schedlogretention* criteria.

Restriction: You cannot specify a non-zero *schedlogmax* value *and* enable *schedlogretention*.

Preferences File

Place this option in the `dsm.sys` file. You can set this option on the **Scheduler** category, **Schedule Log** button of the Preferences editor.

Syntax



Parameters

N or *days*

Specifies how long to wait before pruning the schedule 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.

Tivoli Storage Manager copies the entries pruned from the log to **TSM Sched Pruned 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

Authorized User

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.

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 will not be processed.

If you specify *prompted* mode, you should consider supplying values for the *tcpclientaddress* and *tcpclientport* options in your 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).

Notes:

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.

Preferences File

Place this option in the `dsm.sys` file. You can set this option on the **Scheduler** category, **Schedule Mode** section of the Preferences editor.

Syntax



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 *querschedperiod* option.

PRompted

The client scheduler waits for the server to contact your client node when scheduled work needs to be done.

Note: If you use the **dsmc schedule** command and both *SCHEDMODE PRompt* 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.

Schedrestretrdisabled

Authorized User

The *schedrestretrdisabled* option specifies whether to disable the execution of restore or retrieve schedule operations.

Preferences File

Place this option in the dsm.sys file within a server stanza for the scheduler. You can set this option on the **Scheduler** category, **Schedule Command** button of the Preferences editor.

Syntax



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* and you use one of the following commands.

You can use the *scrolllines* option with the following commands only:

- **delete** **filesystem**
- **query** **archive**
- **query** **backup**
- **query** **backupset**
- **query** **filesystem**
- **query** **options**

Note: The server can also define this option.

Preferences File

Place this option in the dsm.opt file. You can set this option on the **Command Line** category, **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:

-scroll=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 will be 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:

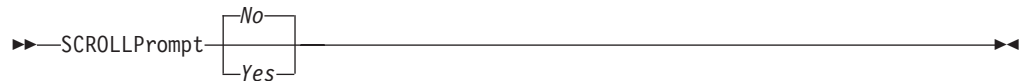
- **delete** **filespace**
- **query** **archive**
- **query** **backup**
- **query** **backupset**
- **query** **filespace**
- **query** **options**

Note: The server can also define this option.

Preferences File

Place this option in the dsm.opt file. You can set this option on the **Command Line** category, **Pause after displaying the following number of lines** field of the Preferences editor.

Syntax



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:

 -scrllp=yes

Command line:

 -scrllp=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 will be 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      server1
  NODename      node1
  COMMMethod    TCPip
  TCPPort       1500
  TCPSeveraddress node.domain.company.com
  PASSWORDAccess generate
  GRoups        qadev qadev1 tools23
  USERS         ashton stewart kaitlin
  incl excl     /Users/user1/Documents/backup1.excl
SErvername      server2
  NODename      node2
  COMMMethod    TCPip
  TCPPort       1500
  TCPSeveraddress almvmd.almaden.ibm.com
  PASSWORDAccess prompt
  GRoups        qadev qadev1 tools23
  USERS         ashton stewart kaitlin
  incl excl     /Users/user1/Documents/backup2.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: If the Tivoli Storage Manager server name changes or Tivoli Storage Manager clients are directed to a different Tivoli Storage Manager server, all clients will need to have a new password initialized for the new server name.

Preferences File

Place this option in the client user options file (`dsm.opt`) and the `dsm.sys` file.

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_aservername yakima
```

Command line:

`-se=server_b`

This option is valid only on the initial command line. It is not valid in interactive mode.

Sessioninitiation

Authorized User

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.

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 and GUI client will still attempt to initiate sessions.

Attention: 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 will initiate 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.

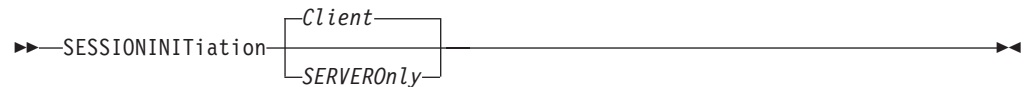
Notes:

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.

Preferences File

Place this option in the `dsm.sys` file *within* a server stanza. You can set this option on the **Scheduler** category, **Session Initiation** field of the Preferences editor.

Syntax



Parameters

Client

Specifies that the client will initiate 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 (see “*Tcpclientport*” on page 255. Except for client acceptor daemon-managed schedulers, the command-line client and GUI client will still attempt 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.

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 will skip ACL processing completely. No ACL data will be backed up or restored. *skipacl=yes* 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) will continue to be backed up and restored as before.

If a prior version of the file being restored has ACL data backed up, ACLs will be restored even when the *skipacl* option is set.

Preferences File

Place this option in the client user options (dsm.opt) file.

Syntax



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 will be 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 will not be backed up.

Preferences File

Place this option in the client user options (dsm.opt) file.

Syntax



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

Srvprepostscheddisabled

Authorized User

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* option to disable the execution of any unwanted Operating System command by a Tivoli Storage Manager administrator on a Tivoli Storage Manager client node.

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.

Preferences 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



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 Manager administrator, the command defined by the administrator will override 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 Manager 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* option.

Examples

Options file:

```
srvprepostscheddisabled yes
```

Command line:

```
Does not apply.
```

Subdir

The *subdir* option specifies whether you want to include subdirectories of named directories for processing on the following commands:

- **archive**
- **delete archive**
- **delete backup**
- **incremental**
- **query archive**
- **query backup**
- **restore**
- **restore backupset**
- **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.

Note: The server can also define this option.

Preferences File

Place this option in the client user options file (dsm.opt).

Syntax



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. For more information, see “Preservepath” on page 217.

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
```

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 will be the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

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. See “Standard query restore, no query restore, and restartable restore” on page 72 for a discussion of standard and no query restore operations.

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.

Preferences File

Place this option in the client user options file (dsm.opt). You can set this option on the **General** category, **Prompt before mounting tapes** checkbox of the Preferences editor.

Syntax



Parameters

- No** You are not prompted for your choice. The server waits for the appropriate tape to mount. This is the default.
- 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

Authorized User

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.

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 will be allowed on the same port as client sessions.

If *tcpadminport* is set on the server, then administrative sessions will be 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*.

Preferences File

Place this option in the `dsm.sys` file *within* a server stanza. You can set this option on the **Communication** category, **Admin Port** field of the Preferences editor. See the *IBM Tivoli Storage Manager Administrator's Reference* for your operating system, for more information.

Syntax

►—TCPADMINPort —————►
└──admin_port_address──┘

Parameters

admin_port_address

Specifies the port number of the server. The default value is the value of the *tcppport* option.

Examples

Options file:

tcpadminport 1502

Tcpbuffsize

Authorized User

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.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Communication** category, **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

Authorized User

The *tpccadaddress* 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 (see “Commmethod” on page 140).

Preferences 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.

Tcpclientaddress

Authorized User

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 will use this address when it begins the server prompted scheduled operation.

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.

Preferences File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Scheduler** category, **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

Command line:

```
-tcpclientaddress=128.33.10.249  
-tcpclientaddress=khoyt.mycompany.mydomain.com  
-tcpclientaddress=2002:92b:111:221:128:33:10:249
```

This option is valid only on the initial command line. It is not valid in interactive mode.

Tcpclientport

Authorized User

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. Use this option only if you specify the *prompted* parameter with the *shedmode* 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.

Preferences File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Scheduler** category, **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 1000 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

Authorized User

The *tcpnodelay* option specifies whether the client disables the delay of sending successive small packets on the network, per transaction. Change the value from the default of *yes* only under one of these 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.

Preferences File

Place this option in the client system options file (dsm.sys) *within* a server stanza.

Syntax



Parameters

- | | |
|-----|--|
| No | Specifies that the server does not allow successive small packets to be sent immediately over the network. Setting this option to <i>no</i> can degrade performance. |
| Yes | Specifies that the server or client allows successive small packets to be sent immediately over the network. The default is <i>yes</i> . |

Examples

Options file:
tcptodelay yes

Command line:
Does not apply.

Tcpport

Authorized User

The *tcpport* option specifies a TCP/IP port address for a Tivoli Storage Manager server. You can obtain this address from your administrator.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Communication** category, **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 1000 through 32767; the default is 1500.

Examples

Options file:

tcpp 1501

Tcpserveraddress

Authorized User

The *tcpserveraddress* option specifies the TCP/IP address for a Tivoli Storage Manager server. You can obtain this port address from your administrator.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **Communication** category, **Server Address** field of the Preferences editor.

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.xyzcompany.vas.com

Command line:

-tcpserveraddress=2002:92b:111:221:128:33:10:249Does not apply.

Tcpwindow size

Authorized User

Use the *tcpwindow size* 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.

Preferences File

Place this option in the `dsm.sys` file. You can set this option on the **Communication** category, **Window Size** field of the Preferences editor.

Syntax

►►—TCPWindow size— *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 1 through 128; the default is 63.

If you specify a value less than 1, the TCP window size defaults to 1. If you specify a value greater than 128, the TCP window size defaults to 128.

Notes:

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:

```
tcpwindow size 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.

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**
- **restore**
- **retrieve**
- **set event**

Preferences File

Place this option in the client options file (dsm.opt). You can set this option on the **Regional Settings** category, **Time Format** field of the Preferences editor.

Syntax

►►—TIMEformat— *format_number* —————►►

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.

- 1 23:00:00
- 2 23,00,00
- 3 23.00.00
- 4 12:00:00 A/P

Examples

Options file:

timeformat 4

Command line:

-time=3

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. For example, you might request a list of files that were backed up before 11:59 PM on June 30, 2002.

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**
- **retrieve**

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
```

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. For example, you might request a list of files that were backed up before 11:59 PM on June 30, 2003. 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**
- **retrieve**

Syntax

►► —TTime =- *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
```

Txnbytelimit

Authorized User

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. Because the client program can transfer more than one file or directory between the client and server before it commits the data to server storage, a transaction can contain more than one file or directory. This is called a *transaction group*.

This option permits you to control the amount of data sent between the client and server before the server commits the data and changes to the server database, thus changing the speed with which the client performs work. The amount of data sent applies when files are batched together during backup or when receiving files from the server during a restore procedure.

The server administrator can limit the number of files or directories contained within a transaction group using the *txngroupmax* option on the server, so the actual size of a transaction can be less than your *txnbytelimit*. Once this number is reached, the client sends the files to the server *even if* the transaction byte limit is not reached.

Preferences File

Place this option in the dsm.sys file. You can set this option on the **General** category **Transaction Buffer Size** field of the Preferences editor.

Syntax

►►—TXNBytelimit— *number* —————►◄

Parameters

number

Specifies the number of kilobytes the client program can buffer together in a transaction before it sends data to the server. The range of values is 300 through 2097152 (2 GB); the default is 25600.

Examples

Options file:

```
txnb 25600
```

Command line:

```
-txnb=25600
```

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.

Syntax



Parameters

- any*
Specifies all nodes registered at the server. This is the default.
- 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=client`

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.

See “How Tivoli Storage Manager handles long user and group names” on page 54 for information about the character limits for group names.

Preferences File

Place this option in the dsm.sys file *within* a server stanza.

Syntax



Parameters

username

Names a user that you want to authorize to request Tivoli Storage Manager services.

Examples

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.

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.

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.

Note: The server can also define this option.

Preferences File

Place this option in the client user options file (dsm.opt). You can set this option on the **Command Line** category, **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.

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 user options file, or with a command:

- You must specify the name you specified with the *nodename* option in your client system options file ().
- 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 identify itself to the server. This login identification is determined in the following ways:

- If the *nodename* and *virtualnodename* options are not specified, you are prompted for a node name.
- 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 specified by the *nodename* option.

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 will only work 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).

Preferences 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 scheduler program outside a firewall by specifying the TCP/IP port number used by the Tivoli Storage Manager Client Acceptor daemon.

A value for the Client Acceptor daemon is required. If you do not specify this option, the default value, zero (0), is used for the client acceptor daemon. This causes TCP/IP to randomly assign a free port number for the Client Acceptor daemon.

Preferences File

Place this option in the dsm.sys file *within* a server stanza. You can set this option on the **Scheduler** category, **Client Acceptor Port** field of the Preferences editor.

Syntax

►►—WEBPorts— *cadport* —————►◄

Parameters

cadport

Specifies the *required* Tivoli Storage Manager Client Acceptor 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:

webports 2123

Command line:

Does not apply.

Chapter 10. 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 chapter describes how to start or end a client command session and how to enter commands. Table 41 shows a list of tasks related to entering commands.

Table 41. Entering commands

Task	Page
Starting and ending a client command session	272
Entering client commands	274
Remembering previous commands	276
Using wildcard characters	277

Table 42 provides an alphabetical list of the commands, a brief description, and where to locate more information.

Table 42. Commands

Command	Description	Page
archive	Archives files from a workstation to Tivoli Storage Manager storage.	279
cancel restore	Displays a list of restartable restore sessions from which you can select one to cancel.	281
delete access	Deletes authorization rules for files that are stored on the server.	282
delete archive	Deletes archived files from Tivoli Storage Manager server storage.	283
delete backup	Deletes active and inactive backup files from Tivoli Storage Manager server storage.	285
delete filespace	Deletes file spaces in Tivoli Storage Manager server storage.	288
expire	Inactivates backup objects that you specify in the file specification or with the <i>filelist</i> option.	289
help	Displays a Table of Contents of help topics for the command-line client.	290
incremental	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.	291
loop	Starts an interactive command session.	295
macro	Executes commands within a macro file that you specify.	296
preview	Generates a file that lists objects to be backed up or archived according to the include-exclude list, prior to sending any data to the server.	297
query access	Displays a list of current authorization rules.	298
query archive	Displays a list of archived files.	299

Table 42. Commands (continued)

Command	Description	Page
query backup	Displays a list of backup versions.	301
query backupset	Queries a backup set from a local file or the Tivoli Storage Manager server.	303
query filespace	Displays a list of file spaces in Tivoli Storage Manager storage. You can also specify a single file space name to query.	305
query inclexcl	Displays a list of include-exclude statements in the order in which they are processed during backup and archive operations.	306
query mgmtclass	Displays information about available management classes.	307
query node	Displays all the nodes for which an administrative user ID has authority to perform operations.	308
query options	Displays all or part of your options and their current settings.	309
query restore	Displays a list of your restartable restore sessions in the server database.	310
query schedule	Displays information about scheduled events for your node.	311
query session	Displays information about your session, including the current node name, when the session was established, server information, and server connection information.	312
query systeminfo	Gathers Tivoli Storage Manager system information and outputs this information to a file or the console.	313
restart restore	Displays a list of restartable restore sessions from which you can one to restart.	315
restore	Restores copies of backup versions of your files from a Tivoli Storage Manager server.	316
restore backupset	Restores a backup set from the Tivoli Storage Manager server or a local file.	320
retrieve	Retrieves copies of archived files from the Tivoli Storage Manager server.	323
schedule	Starts the client scheduler on the workstation.	326
selective	Backs up selected files.	328
set access	Authorizes another user to access your backup or archive data.	330
set event	Allows you to specify the circumstances for when archived data is deleted.	332
set password	Changes the Tivoli Storage Manager password for your workstation.	334

Starting and ending 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. Refer to Chapter 9, “Using processing options,” on page 111 for the option that you are using to see if it is valid only on the initial command line.

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.

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.

Entering client commands

A client command can include one or more of these components:

- Command name
- Options
- Parameters

The sections 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

There are two groups of options that you can use with commands:

- **Client options:** The group of options that are set in your client user options file (dsm.opt). To override an option in the client user options file (dsm.opt), enter the option with a command. For detailed information about client options, see “Client options reference” on page 127.
- **Client command options:** Use this group of options with specific commands on the command line *only*. For detailed information about client command options, see “Client options reference” on page 127.

Options handling in interactive mode

In interactive mode, options you enter on the initial command line will override 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 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:

```
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 and file name 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*. For example, to restore the /project/budget.fin file to /newproj/newbudg.fin, enter:

```
dsmc restore /project/budget.fin /newproj/newbudg.fin
```

Enter parameters in the order indicated in the command syntax diagram.

File specification syntax

Use the following syntax rules when entering file specification parameters, such as *filespec*, *sourcefilespec*, and *destinationfilespec*:

- The command-line client accepts a forward slash (/) delimiter. For example, to back up all files in the /Volumes/fs1/devel directory with an extension of .c, enter:

```
dsmc sel "/Volumes/fs1/devel/*.c"
```

- 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 and Tivoli Storage Manager builds the fully qualified file specification. For example, if the current working directory is /Users/me, then the *destinationfilespec* would be /Users/me/mydir in the following command:

```
dsmc restore "/Volumes/fs/dir1/*" 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:

```
dsmc sel /Users
```

- 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 /Users/mydir/ /Volumes/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 /Users/mydir /Volumes/away/yourdir
```

- 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. For example, to allow access to all files in all directories *in* and *below* the /Users directory, enter:

```
set access backup /Users/* * *
set access backup /Users/*/* * *
```

Do not use wildcards for the directory path name, for example:

```
/Users/j*asler/file1.c
```

- The maximum number of characters for a file name is 256. The maximum combined length of the file name and path name is 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.
- The maximum number of file specifications per command:
 - The Query commands can accept only one file specification.
 - The **restore** and **retrieve** commands can accept a *sourcefilespec* and a *destinationfilespec*.
 - Unless the *removeoperandlimit* option is used, there is a 20-operand limit on these commands. This is to prevent inadvertent use of wildcards, which get expanded by the shell command processor. Allowing the command processor to expand wildcards results in very inefficient processing by the Tivoli Storage Manager client. You can 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. See “Filelist” on page 171 for more information.

Notes:

1. You can overcome these limitations by using 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. See “Filelist” on page 171 for more information.
2. You can also use the *removeoperandlimit* option to specify that Tivoli Storage Manager removes the 20-operand limit for the Mac OS X platform. 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. See “Removeoperandlimit” on page 222.

Remembering previous commands

If you set the *editor* option to *yes* in your client options file (dsm.opt), Tivoli Storage Manager permits you to recall and edit as many as 20 previously entered commands by using the Up arrow and Down arrow keys. If you set the *editor* option to *no*, the feature to recall previous commands is not active. If the editor and command retrieve functions are not working on a specific workstation setting, you should turn off the *editor* option. For more information regarding the *editor* option, see “Editor” on page 159.

Pressing the Up arrow key displays the previous command in memory. Pressing the Down arrow key displays the next command in memory. Table 43 lists other functions you can perform when you recall commands.

Table 43. Command recall and edit functions

Function	Press
Display the previous command in memory.	Up arrow
Display the next command in memory.	Down arrow
Move to the left.	Left arrow
Move to the right.	Right arrow
Move five spaces to the right.	Tab right
Move to the beginning of the previous word	Ctrl-left arrow or CTRL-L
Move to the beginning of the next word.	Ctrl-right arrow or CTRL-R
Delete a character to the right of the cursor.	Del

Table 43. Command recall and edit functions (continued)

Function	Press
Delete a character to the left of the cursor.	Delete
Finish or execute the command.	Return
Quit the program.	F3 or Esc
End the program.	CTRL-C

Using wildcard characters

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 directories. You cannot specify a directory whose name contains an asterisk (*) or a question mark (?). Tivoli Storage Manager recognizes these characters only as 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. Valid wildcard characters that you can use include:

- * Asterisk. Matches zero or more characters.
- ? Question mark. Matches any single character at the present position.

Table 44 shows examples of each wildcard.

Table 44. Wildcard characters

Pattern	Matches	Does not match
Asterisk (*)		
ab*	ab, abb, abxxx	a, b, aa, bb
ab*rs	abrs, abtrs, abrrs	ars, aabrs, abrrs
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, abkkrs

Attention: 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, you must enclose values containing wildcards in quotation marks. For example:

```
dsmc selective "/Users/me/*.c"
```

Entering commands

Follow the general rules below when you enter commands:

- 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.

Client commands reference

The following sections contain detailed information about each of the Tivoli Storage Manager commands. Information for each command includes:

- A description of the command.
- A syntax diagram of the command. The command name contains uppercase and lowercase characters. The uppercase characters indicate the minimum abbreviation you can use for the command name. See “Reading syntax diagrams” on page x for an explanation of these diagrams.
- 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.

See “Supported file systems” on page 53 for supported file systems, ACL, and Extended Attribute support.

Removing operand limits

You can use the *removeoperandlimit* option to specify that Tivoli Storage Manager removes the 20-operand limit for the Mac OS X platform. 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. See “Removeoperandlimit” on page 222.

Migrating to Unicode-enabled file spaces

See “Autofsrename” on page 135 for information about using the Unicode-enabled client.

Syntax



Parameters

options

Table 45. Archive command: Related options

Option	Where to use	Page
<i>archmc</i>	Command line only.	128
<i>autofsrename</i>	Client options file (dsm.opt) only.	135
<i>changingretries</i>	Client system options file (dsm.sys) or command line.	138
<i>compressalways</i>	Client user options file (dsm.opt) or command line.	143
<i>compression</i>	dsm.sys file <i>within</i> a server stanza or command line.	144
<i>deletefiles</i>	Command line only.	149
<i>description</i>	Command line only.	150
<i>dirsonly</i>	Command line only.	153
<i>encryptiontype</i>	System options file (dsm.sys) within a server stanza.	161
<i>encryptkey</i>	System options file (dsm.sys) within a server stanza.	162
<i>filelist</i>	Command line only.	171
<i>filesonly</i>	Command line only.	174
<i>removeoperandlimit</i>	Command line only.	222

Table 45. Archive command: Related options (continued)

<i>subdir</i>	Client user options file (dsm.opt) or command line.	247
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	249
<i>v2archive</i>	Command line only.	266

filespec

Specifies 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. See “Maximum file size for backup and restore” on page 54 for the maximum file size for archive processing.

Examples

Task Archive a single file named budget.jan in the /Users/mike/Documents directory.

Command: archive /Users/mike/Documents/budget.jan

Task Archive a single file named budget.jan in the /Volumes/proj4 directory.

Command: archive /Volumes/proj4/budget.jan

Task Archive all files in the /Users/user1/Documents directory with a file extension of .txt.

Command: archive /Users/user1/Documents/*.txt

Task Archive all files in the directory tree headed by the /Users directory.

Command: archive -subdir=yes /Users/*

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 when:

- 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 server's sequential volumes.

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 specified by that rule.

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	/Users/dthomas/proja /list/
2	Archive	NODE3	LUIE	/Users/budg/depta/
3	Backup	NODE4	USER2	/Users/exp/deptc/
4	Archive	NODE5	USER2S	/Users/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 , 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.

Attention: When you delete archived files, *you cannot retrieve them*. Verify that the files are obsolete *before* you delete them.

Syntax

►► Delete ARchive options filespec
"filespec" ►►

Parameters

options

Table 46. Delete Archive command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	146
<i>description</i>	Command line only.	150
<i>filelist</i>	Command line only.	171
<i>noprompt</i>	Command line only.	203
<i>numberformat</i>	Client user options file (dsm.opt) or command line.	204
<i>pick</i>	Command line only.	210
<i>subdir</i>	Client user options file (dsm.opt) or command line.	247
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	249
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	260

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.

Examples

Task Delete a file named budget.jan.

Command: delete archive /Users/user2/Documents/budget.jan

Task Delete a file named budget.jan in the /Volumes/proj4 directory.

Command: delete archive /Volumes/proj4/budget.jan

Task Delete all files archived from the /Users/user1/Documents directory with a file extension of .txt.

Command: delete archive /Users/user1/Documents/*.txt

Task Delete files archived from the /Users/user1/Documents 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 /Users/user1/Documents/*. * -pick

Task Delete selected files from the group of files archived with the description "Monthly Budgets 1999".

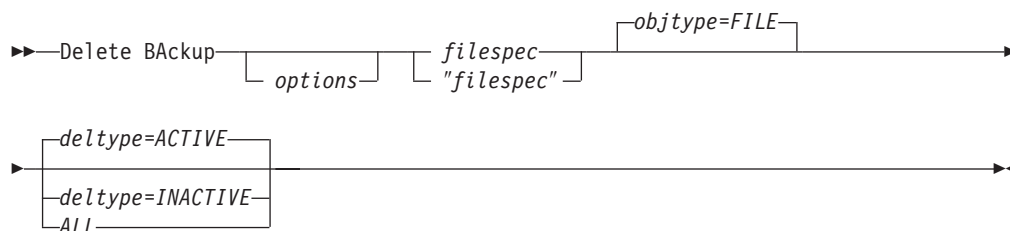
Command: delete ar -description="Monthly Budgets 1999" -pick

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.

Attention: 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 will prompt 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.

Syntax



Parameters

options

Table 47. Delete Backup command: Related options

Option	Where to use	Page
<i>description</i>	Command line only.	150
<i>filelist</i>	Command line only.	171
<i>fromdate</i>	Command line, and in GUI find function.	171
<i>fromtime</i>	Command line, and in GUI find function.	171
<i>noprompt</i>	Command line only.	203
<i>pick</i>	Command line only.	210
<i>pitdate</i>	Command line, and in GUI find function.	171
<i>pittime</i>	Command line, and in GUI find function.	171
<i>subdir</i>	Client user options file (dsm.opt) or command line.	247
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	249
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	260
<i>todate</i>	Command line, and in GUI find function.	171
<i>totime</i>	Command line, and in GUI find function.	171

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 will be 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.

filespec

Specifies the path and file name that you want to delete from storage.

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.

Examples

Task Delete all active and inactive file objects named `budget.jan` in directory `/Users/user2/proj1`.

Commands:

```
delete backup /Users/user2/proj1/budget.jan -deltype=inactive
delete backup /Users/user2/proj1/budget.jan -deltype=active
```

Note: To delete active and inactive file objects, you must first delete the inactive objects, followed by the active objects.

Task Delete a file named `budget.jan` in the `/Volumes/proj4` directory.

Command: `delete backup /Volumes/proj4/budget.jan`

Task Delete all inactive files with names ending in `.txt` backed up from the `/Users/user2/proj1` directory and its subdirectories.

Command: `delete backup /Users/user2/proj1/*.txt -deltype=inactive -subdir=yes`

Task Delete selected active files backed up from the `/Users/user2/` 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 /Users/user2/project/* -pick`

Task Delete all active and inactive versions of files and subdirectories in the `/Users/user2/myproject` directory. Then delete all active and inactive versions of the `/user/myproject` directory.

Command: `delete backup /Users/user2/myproject/* -deltype=all`

Note: The backup versions of directory object `/Users/user2/myproject` are not deleted.

Delete Filespace

System Administrator

The **delete filesystem** 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. 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 filesystem** command, a list of your file spaces is displayed. From this list, select the file space that you want to delete.

Your 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.

Attention: 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.

Syntax

►►—Delete Filespace—┐ options ┘◄◄

Parameters

options

Table 48. Delete Filespace command: Related options

Option	Where to use	Page
<i>detail</i>	Command line only.	151
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	238
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	239

Examples

Task Delete a file space.

Command: delete filesystem

Expire

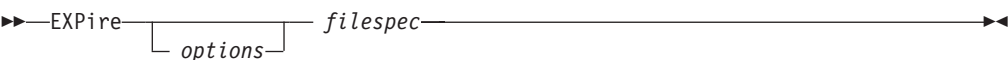
The **expire** command inactivates the backup objects you specify in the file specification or with the *filelist* option.

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 will be displayed on the command line, if you specify the proper query with a wildcard character for the directory.

Syntax



Parameters

options

Table 49. Expire command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	146
<i>filelist</i>	Command line only.	171
<i>noprompt</i>	Command line only.	203
<i>numberformat</i>	Client user options file (dsm.opt) or command line.	204
<i>pick</i>	Command line only.	210
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	260

Note: If you specify *filelist*, then *pick* is ignored.

filespec

Specifies a path and a filename 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.

Examples

Task Inactivate the letter1.txt file in the home directory.

Command: dsmc expire /Users/home/letter1.txt

Task Inactivate all files in the /Users/user1/Documents/ directory.

Command: expire /Users/user1/Documents/*

Task Inactivate all files named in the /Users/avi/Documents/filelist.txt file.

Command: expire -filelist=/Users/avi/Documents/filelist.txt

Help

The **help** command displays a Table of Contents of help topics for the command-line client. The topics include help for the following:

- Summary of Changes
- Using Commands
- Using Processing Options
- Glossary
- Messages

Enter the number of the topic that you want to view. If there is more than one screen of topics, scroll backward or forward through the list. To exit, type **q** and press Enter.

Note: If you use the **help** command on the initial command line, no server contact is made and no password is needed.

Syntax

►►—Help—◄◄

Parameters

There are no parameters for this command.

Examples

Task Display a list of help topics.

Command: help

Incremental

The **incremental** command 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. See “Supported file systems” on page 53 for supported file systems, ACL, and Extended Attribute support.

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. See “Domain” on page 157 for information on how to change which objects are included in the default domain.

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.

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.

For more information on management classes, see Chapter 8, “Understanding storage management policies,” on page 99.

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.

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. For more information, see “Selective” on page 328.

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.

Removing operand limits

See “Removeoperandlimit” on page 222 for information about operand limits.

Migrating to Unicode-enabled file spaces

See “Autofsrename” on page 135 for information about using the Unicode-enabled client.

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.

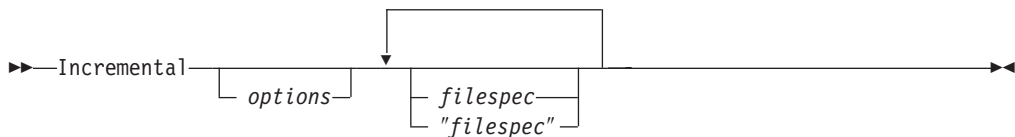
Unlike a full incremental, an incremental-by-date does not maintain current server storage of *all* your workstation files because:

- 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, 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.

Syntax



Parameters

options

Table 50. Incremental command: Related options

Option	Where to use	Page
<i>changingretries</i>	dsm.sys file or command line.	138
<i>compressalways</i>	Client user options file (dsm.opt) or command line.	143
<i>compression</i>	dsm.sys file <i>within</i> a server stanza or command line.	144
<i>dironly</i>	Command line only.	153
<i>domain</i>	dsm.sys file or the client user options file (dsm.opt) or command line only.	157
<i>encryptiontype</i>	System options file (dsm.sys) within a server stanza.	161
<i>encryptkey</i>	System options file (dsm.sys) within a server stanza.	162
<i>filelist</i>	Command line only.	171
<i>filesonly</i>	Command line only.	174
<i>incrbydate</i>	Command line only.	190
<i>memoryefficientbackup</i>	Client user options file (dsm.opt), client system options file (dsm.sys), server, or command line.	199
<i>removeoperandlimit</i>	Command line only.	222
<i>subdir</i>	Client user options file (dsm.opt) or command line.	247
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	249

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. See “Maximum file size for backup and restore” on page 54 for the maximum file size for backup processing.

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.

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-by-date backup for the /Users file system.

Command: Incremental -incrbydate /Users

Task Run an incremental backup of all files in the /Volumes/proj6 directory that begin with the string abc.

Command: Incremental -subdir=yes /Volumes/proj6/abc*

Task Run an incremental backup of the abc file in the /Users/user1/Documents directory.

Command: Incremental -subdir=yes /Users/user1/Documents/abc

Task Run an incremental backup of the directory object /Users/user2/Documents, but not any of the files in the /Users/user2/Documents directory.

Command: Incremental /Users/user2/Documents

Task Run an incremental backup of the directory object /Users/user2/Documents, all of the files in the /Users/user2/Documents directory, and all files and subdirectories under /Users/user2/Documents.

Command: Incremental -subdir=yes /Users/user2/Documents/

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 will be 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`, will only be 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 will override 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.*

See Chapter 9, “Using processing options,” on page 111 for options that you cannot use in interactive mode.

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.

- Enter `quit`
- Press the Escape key, and then enter `QQ`
- If other commands fail, enter the following keystroke combination: `"QQ"`.

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.

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:
 /Users/user3/Documents/proja
 /Users/user3/Documents/projb
 /Users/user3/Documents/projc

Command: macro backabc.mac

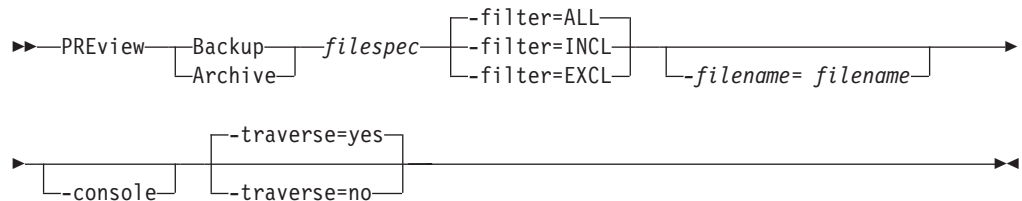
where backabc.mac contains the following statements:

```
Selective /Users/user3/Documents/proja/  
Selective /Users/user3/Documents/projb/  
Selective /Users/user3/Documents/projc/
```

Preview

The **preview** command simulates a backup or archive command without sending data to the server. The **preview** command generates a tab delimited 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, the file is excluded will be listed, along with the source for the pattern.

Syntax



Parameters

Backup|Archive

Indicates whether to preview output from a selective backup or archive operation.

filespec

Specifies the path and file name that you want to back up. Use wildcard characters to select a group of files or all of 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-delimited 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.

Attention: Specifying **-traverse** does not preview directories excluded using the *exclude.dir* option.

Query Access

The **query access** command displays a list of users to whom 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 information includes:

- 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.

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
- Archive description

If you use the *detail* option with the **query archive** command, the client displays the following additional information:

- Last modification date

Syntax

```
➤—Query Archive—┐ ┐ filespec ┐ ┐  
                  └ options ┘ └ "filespec" ┘ ┐ ┐
```

Parameters

options

Table 51. Query Archive command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	146
<i>description</i>	Command line only.	150
<i>detail</i>	Command line only.	151
<i>dirsonly</i>	Command line only.	153
<i>filelist</i>	Command line only.	171
<i>filesonly</i>	Command line only.	174
<i>fromdate</i>	Command line only.	175
<i>frommode</i>	Command line only.	176
<i>fromowner</i>	Command line only.	177
<i>fromtime</i>	Command line only.	178
<i>numberformat</i>	Client user options file (dsm.opt) or command line.	204
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	238
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	239
<i>subdir</i>	Client user options file (dsm.opt) or command line.	247
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	260
<i>to date</i>	Command line only.	261
<i>to time</i>	Command line only.	262

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.

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 `/Users/user3/Documents` directory and all of its subdirectories.
Command: `query archive /Users/user3/Documents/*`
- 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 of each file.
Command: `q ar -detail *`
- Task** Display a list of archived files in the `/Users/user1/Documents` directory whose first four characters of the file name begin with `proj`.
Command: `q ar /Users/user1/Documents/proj*`

Query Backup

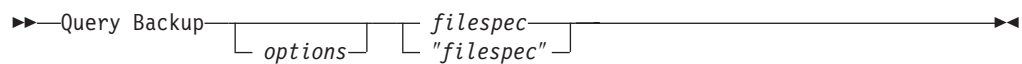
The **query backup** command displays a list of backup versions of your files stored on the Tivoli Storage Manager server, or inside a backupset when a backupset is specified. File information includes the following:

- File specification
- File size
- Backup date
- Whether the file is active or inactive
- The management class 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

Syntax



Parameters

options

Table 52. Query Backup command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	137
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	146
<i>detail</i>	Command line only.	151
<i>dirsonly</i>	Command line only.	153
<i>filelist</i>	Command line only.	171
<i>filesonly</i>	Command line only.	174
<i>fromdate</i>	Command line only.	175
<i>fromnode</i>	Command line only.	176
<i>fromowner</i>	Command line only.	177
<i>fromtime</i>	Command line only.	178
<i>inactive</i>	Command line only.	183
<i>location</i>	Command line only.	194
<i>numberformat</i>	Client user options file (dsm.opt) or command line.	204
<i>pitdate</i>	Command line only.	211
<i>pittime</i>	Command line only.	212
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	238
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	239
<i>subdir</i>	Client user options file (dsm.opt) or command line.	247
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	260
<i>todate</i>	Command line only.	261

Table 52. Query Backup command: Related options (continued)

<i>totime</i>	Command line only.	262
---------------	--------------------	-----

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.

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 *

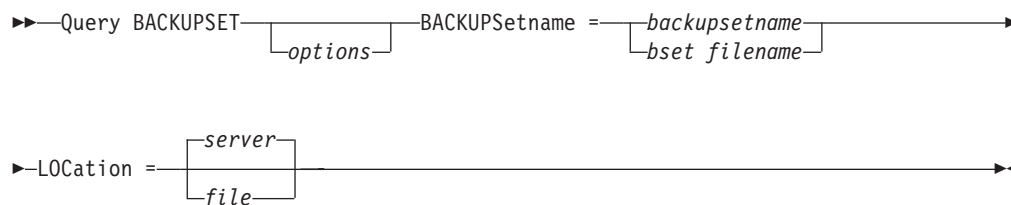
Task Display a list of files that were backed up from the /Users/user3/Documents directory with file names that begin with proj.

Command: q b /Users/user3/Documents/proj*

Query Backupset

The **query backupset** command queries a backup set from a local file or the Tivoli Storage Manager server. This command displays the backup set name, generation date, retention, and description.

Syntax



Parameters

options

Table 53. Query Backupset command: Related options

Option	Where to use	Page
<i>location</i>	Command line only.	194
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	238
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	239

BACKUPSetname=

Specifies the name of the backup set on the server you want to query when **-location=server** is in effect. 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 on the screen. 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.

backupsetname

Specifies the name of the backup set.

bset filename

Specifies the name of the file containing the beginning of the backup set. The *location* option should be set to *file*.

LOCation=

Specifies where Tivoli Storage Manager searches for a backup set during a query or restore operation. Use this option to locate backup sets on the server or local files. Backup sets that are generated on the server can be used locally specifying the *location* option and either the file name of the file containing the start of the backup set, or the tape device where the tape containing the start of the backup set is located.

server Specifies that Tivoli Storage Manager searches for the backup set on the server. This is the default.

file Specifies that Tivoli Storage Manager searches for the backup set on a local file.

Examples

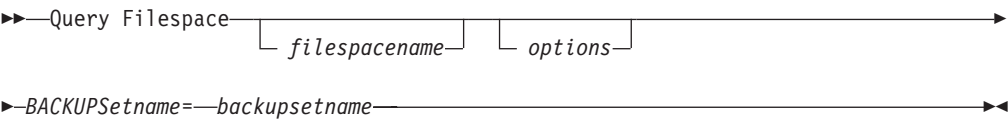
- Task** Query a backup set called `monthly_financial_data` on the server.
- Command:** `query backupset backupsetname=monthly_financial_data -loc=server`
- Task** Query the backup set in the `weekly_budget_data.ost` file in the budget directory.
- Command:** `dsmc query backupset -backupsetname=/Users/user4/Documents/budget/weekly_budget_data.ost -loc=file`

Query Filespace

The **query filesystem** command displays a list of file spaces for a node that are stored on the Tivoli Storage Manager server, or inside a backupset when a backupset 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 file system at your workstation from which you back up or archive files.

A Unicode file space name might not display correctly if the server is unable to display the Unicode name. In this case, use the file space identifier (fsID) to identify these file spaces on the server. Use the **query filesystem** command with the *detail* option to determine the fsID of a file space.

Syntax



Parameters

filesystemname
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.

options

Table 54. Query Filespace command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	137
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	146
<i>detail</i>	Command line only.	151
<i>fromnode</i>	Command line only.	176
<i>fromowner</i>	Command line only.	177
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	238
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	239
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	260

Examples

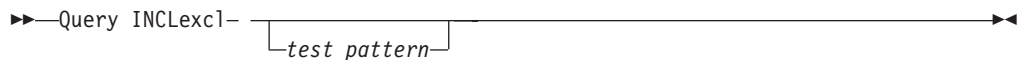
- Task** Display your file spaces.
Command: query filesystem
- Task** Display your file spaces. Use the *dateformat* and *timeformat* options to reformat the dates and times.
Command: query filesystem -date=5 -time=4
- Task** Display the /Users file space.
Command: query filesystem /Users

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.

Syntax



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.

Examples

Task Display a list of include-exclude statements.

Command: query inclexcl

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 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 from the command line.

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. Valid values are *client*, *server*, and *any*. The default is *any*. See “Type” on page 264 for more information.

Syntax

►►—Query Node—►►
└─ options ─┘

Parameters

options

Table 56. Query Node command: Related options

Option	Where to use	Page
<i>type</i>	Command line only.	264
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	238
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	239

Examples

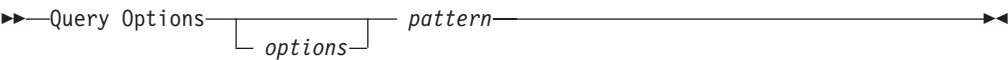
Task Display all client nodes that are backup-archive clients.

Command: query node -type=client

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.

Syntax



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.

options

Table 57. Query Options command: Related options

Option	Where to use	Page
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	238
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	239

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 server's sequential volumes. To unlock the file space, either restart the restore and allow it to complete (**restart 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.

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.

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.

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:

TSM Server Connection Information

Server Name.....: FIJI_0918GA
Server Type.....: AIX-RS/6000
Server Version.....: Ver. 5, Rel. 1, Lev. 0.0
Last Access Date.....: 09/04/1999 15:09:40
Delete Backup Files.....: "No"
Delete Archive Files.....: "Yes"

Node Name.....: EPSILON3
User Name.....: thompson

Query Systeminfo

Use the **query systeminfo** command to gather information on one or more of the following items and output this information to a file or the console:

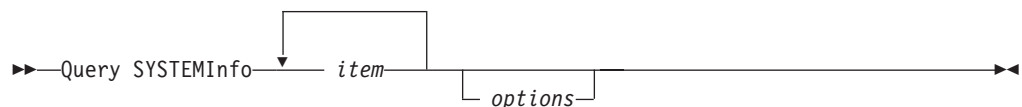
- DSMOPTFILE - The contents of dsm.opt file.
- DSMSYSFILE - The contents of the dsm.sys file.
- ENV - Environment variables.
- ERRORLOG - The Tivoli Storage Manager error log file.
- FILE - Attributes for the file name that you specify.
- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- OPTIONS - Compiled options.
- OSINFO - Name and version of the client operating system
- POLICY - Policy set dump.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually TSM Schedule Log).
- ENCRYPT - Available encryption methods.

Notes:

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. See “Filename” on page 173 for more information.
2. Use the *console* option if you want to output the information to the console. See “Console” on page 145 for more information.

Note: 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. In this case, it is recommended that you 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.

Syntax



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.

options

Table 58. Query Systeminfo command: Related options

Option	Where to use	Page
<i>console</i>	Command line only.	145
<i>filename</i>	Command line only.	173

Examples

Task Gather and store the contents of the TSM User Preference file and the Tivoli Storage Manager error log file in the tsminfo.txt file.

Command: query systeminfo dsmdptfile errorlog
-filename=tsminfo.txt

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 server's sequential volumes.
- You cannot back up files affected by the restartable restore.

Options from the failed session supersede new or changed options for the restarted session.

Syntax

►►—REStArt Restore—◄◄

Parameters

There are no parameters for this command.

Examples

Task Restart a restore.

Command: restart restore

Restore

The **restore** command obtains copies of backup versions of your files from a Tivoli Storage Manager server, or inside a backupset when a backupset is specified. 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. For more information, see “Preservepath” on page 217.

See “Supported file systems” on page 53 for supported file systems, ACL, and Extended Attribute support.

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.

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
```

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 "/Users/user1/*" -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.

Restoring from file spaces that are not Unicode-enabled

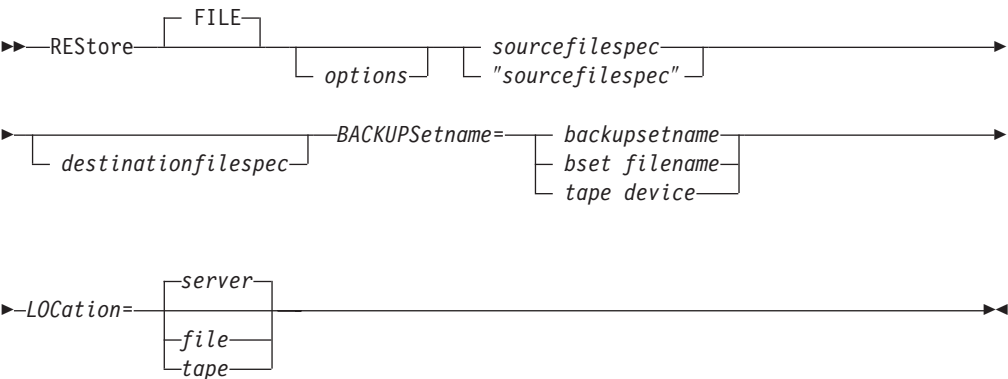
If you want to restore from file spaces that are not Unicode-enabled, you must specify the source on the server and a destination on the client, prior to installing the Unicode-enabled client..

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/
```

Syntax



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**.

options

Table 59. Restore command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	137
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	146
<i>dirsonly</i>	Command line only.	153
<i>filelist</i>	Command line only.	171
<i>filesonly</i>	Command line only.	174
<i>fromdate</i>	Command line only.	175

Table 59. Restore command: Related options (continued)

<i>fromnode</i>	Command line only.	176
<i>fromowner</i>	Command line only.	177
<i>fromtime</i>	Command line only.	178
<i>ifnewer</i>	Command line only.	182
<i>inactive</i>	Command line only.	183
<i>latest</i>	Command line only.	192
<i>location</i>	Command line only.	194
<i>numberformat</i>	Client user options file (dsm.opt) or command line.	204
<i>pick</i>	Command line only.	210
<i>pitdate</i>	Command line only.	211
<i>pittime</i>	Command line only.	212
<i>preservepath</i>	Command line only.	217
<i>replace</i>	Client user options file (dsm.opt) or command line.	223
<i>subdir</i>	Client user options file (dsm.opt) or command line.	247
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	249
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	260
<i>to date</i>	Command line only.	261
<i>totime</i>	Command line only.	262

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.

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 will create it.

Examples

Task Restore a single file named budget in the /Users/user1/Documents directory.

Command: restore /Users/user1/Documents/budget

Task Restore a single file named budget in the /Volumes/proj5 directory.

Command: restore /Volumes/proj5/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 /Users/user3/Documents directory.
- Command:** restore /Users/user3/Documents/*.c
- Task** Restore files in the /Users/user4/Documents directory. Use the *pick* and *inactive* options to select active and inactive backup versions.
- Command:** /Users/user4/Documents/* -pi -ina
- Task** Restore all files from the /Volumes/fs2/projecta directory that end with the character .c to the /Volumes/newfs/projectn/projecta directory. If the projectn or the projectn/projecta directory does not exist, it is created.
- Command:** restore /Volumes/fs2/projecta/*.c /Volumes/newfs/projectn/projecta
- Task** Restore all files in the /Users/user1/Documents directory to their state as of 1:00 PM on August 17, 2002.
- Command:** restore -pitd=8/17/2002 -pitt=13:00:00 /Users/user1/Documents/
- Task** Restore a file from the renamed file space Jaguar_OLD to its original location. Enter both the source and destination as follows:
- Command:** res Jaguar_OLD/user5/Documents/myresume.doc /Users/user5/Documents/
- Task** Restore files specified in the filelist to a different location.
- Command:** res -filelist=/Users/user1/documents/restorelist.txt /Users/user1/NewRestoreLocation/

Restore Backupset

The **restore backupset** command restores a backup set from the Tivoli Storage Manager server or a local file.

Use the *backupsetname* and *location* options with the **restore backupset** command to specify where Tivoli Storage Manager searches for a backup set during the restore operation. See “Location” on page 194 for more information.

If you are restoring a file space from a backup set to a system that did not perform the original backup, you might need to:

- Specify a destination
- Use the syntax below to specify the source file

```
dsmc restore backupset -backupsetname=<backupsetname>
                        -location=<server|tape|file>
                        "/fsname/*" /destfs/ -subdir=yes
```

- You must set the *subdir* option to *yes*.
- The *sourcefilespec* must be the virtual file space name , followed with a terminating directory delimiter. For example:

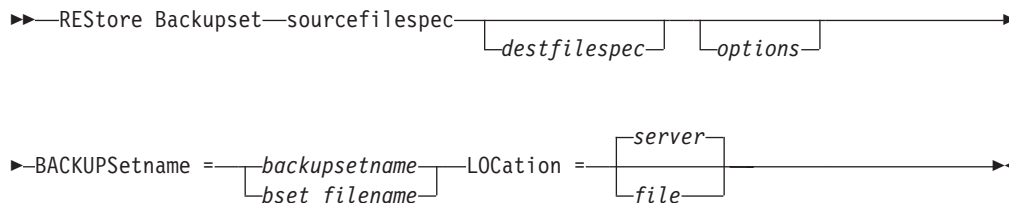
```
restore backupset {myGrpFs}\* c:\myGroupRestDir
  -backupsetname=groupData.1234 -location=server -subdir=yes
restore backupset /virtfs/* /Users/devel/projectb/
  -backupsetname=mybackupset -loc=server -subdir=yes
```

The entire group, or groups in the virtual file space will be restored. You cannot restore a partial group by specifying a qualified source file space.

Considerations:

- You must be a root user or an authorized user to restore an entire backup set from the server, otherwise only files you own are restored.
- 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 and you set the *location* option to *file*, no attempt is made to contact the server.
- 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.

Syntax



Parameters

options

Table 60. Restore Backupset command: Related options

Option	Where to use	Page
<i>dironly</i>	Command line only.	153
<i>filesonly</i>	Command line only.	174
<i>ifnewer</i>	Command line only.	182
<i>location</i>	Command line only.	194
<i>preservepath</i>	Command line only.	217
<i>quiet</i>	Client user options file (dsm.opt) or command line.	221
<i>replace</i>	Client user options file (dsm.opt) or command line.	223
<i>subdir</i>	Client user options file (dsm.opt) or command line.	247

sourcefilespec

Specifies that you want to restore all Windows XP system objects from the backup set. You cannot restore individual system objects from the backup set. This parameter is for the Windows XP clients only.

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.

BACKUPSetname=

Specifies the name of the backup set on the server from which to perform a restore operation. You cannot use wildcard characters to specify the backup set name. Tivoli Storage Manager restores the backup set from the server using the LAN.

backupsetname

Specifies the name of the backup set on the server from which to perform a restore operation. The *location* option should be set to *server*.

bset filename

Specifies the name of the file containing the beginning of the backup set. The *location* option should be set to *file*.

LOCation=

Specifies where Tivoli Storage Manager searches for a backup set during a query or restore operation. Use this option to locate backup sets on the server or local files. Backup sets that are generated on the server can be used locally specifying the *location* option and either the file name of the file containing the start of the backup set, or the tape device where the tape containing the start of the backup set is located.

server Specifies that Tivoli Storage Manager restores the backup set on the server. This is the default.

file Specifies that Tivoli Storage Manager restores the backup set on a local file.

Examples

Task Restore a backup set called `monthly_financial_data` from the server.

Command: `dsmc restore backupset
-backupsetname=monthly_financial_data -loc=server`

Task Restore the backup set contained in the `weekly_budget_data.ost` file in the `/Users/mike/Documents` directory.

Command: `dsmc restore backupset -backupsetname=/Users/mike/
Documents/weekly_budget_data.ost loc=file`

Task Restore the backup set contained in local file `/Users/user3/Documents/bset01.file`.

Command: `dsmc restore backupset /Users/user3/Documents/bset01.file
loc=file`

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. For more information, see “Client options reference” on page 127.

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.

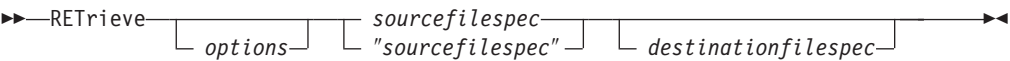
Retrieving from file spaces that are not Unicode-enabled

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. 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/
```

Syntax



Parameters

options

Table 61. Retrieve command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	146
<i>description</i>	Command line only.	150
<i>dirsonly</i>	Command line only.	153
<i>filelist</i>	Command line only.	171
<i>filesonly</i>	Command line only	174
<i>fromdate</i>	Command line only	175

Table 61. Retrieve command: Related options (continued)

<i>fromnode</i>	Command line only.	176
<i>fromowner</i>	Command line only	177
<i>fromtime</i>	Command line only	178
<i>ifnewer</i>	Command line only	182
<i>pick</i>	Command line only.	210
<i>preservepath</i>	Command line only.	217
<i>replace</i>	Client user options file (dsm.opt) or command line.	223
<i>subdir</i>	Client user options file (dsm.opt) or command line.	247
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	249
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	260
<i>to date</i>	Command line only.	261
<i>to time</i>	Command line only.	262

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. See “Maximum file size for backup and restore” on page 54 for the maximum file size for retrieve processing.

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.

Examples

- Task** Retrieve a single file named budget.
- Command:** `dsmc retrieve /Users/user5/Documents/budget`
- Task** Retrieve all files with an extension of .c from the /Volumes/proj5 directory.
- Command:** `dsmc retrieve "/Volumes/proj5/*.c"`
- Task** Retrieve all files in the /Users/devel/Documents directory.
- Command:** `dsmc retrieve "/Users/devel/Documents/*"`
- Task** Retrieve all files with a file extension of .c from the /Users/devel/projecta directory to the /Users/newdevel/projectn/projecta directory. If the /projectn or the /projectn/projecta directory does not exist, it is created.
- Command:** `retrieve /Users/devel/projecta/*.c /Users/newdevel/projectn/projecta/`
- Task** Retrieve files in the /Users/user5/Documents directory. Use the *pick* option.
- Command:** `ret /Users/user5/Documents/* -pick`
- Task** Retrieve all files archived from the /Users/user1/proj directory with the description "2002 survey results."
- Command:** `retrieve /Users/user1/proj/* -desc="2002 survey results"`
- 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/`

Schedule

System Administrator

The **schedule** command starts the client scheduler on your workstation. The client scheduler must be running before scheduled work can start.

Notes:

1. To use the **schedule** command, specify *managedservices none* in the dsm.sys file.
2. This command is valid only on the initial command line. It is not valid in interactive mode or in a macro file.

If the *shedmode* option is set to polling, the client scheduler contacts the server for scheduled events at the hourly interval you specified with the *querschedperiod* option in your client user options file (dsm.opt). If your administrator sets the *querschedperiod* 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 *shedmode* 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. See “Sessioninitiation” on page 242 for more information.

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.

See “Migrating to the Tivoli Storage Manager Version 5.5 Macintosh client from releases prior to 5.3.2” on page 2 for updates to the scheduler paths.

Note: You *cannot* enter this command in interactive mode.

Syntax

►►—Schedule—┐
 └ options ┘◄◄

Parameters

options

Table 62. Schedule command: Related options

Option	Where to use	Page
<i>maxcmdretries</i>	or command line.	198
<i>password</i>	client user options file (dsm.opt)	206
<i>querschedperiod</i>	or command line.	220
<i>retryperiod</i>	or command line.	228
<i>schedlogname</i>	or command line.	232
<i>tcpclientport</i>	or command line.	255

Examples

Task Start the client scheduler.

Command: `dsmc sch -password=notell`

Task Run the schedule command in the background and keep the client scheduler running, even if you log off your system. Ensure the *passwordaccess* option is set to *generate*.

Command: `nohup dsmc schedule 2> /dev/null &`

When you run the **schedule** command, all messages regarding scheduled work are sent to the TSM Schedule 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` TSM Schedule Log file will reside in the default directory:

`/Library/Logs/tivoli/tsm/`

Attention: 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.

Selective

The **selective** command backs up files that you specify. If these files become damaged or lost, you can replace them with backup versions from the server. When you run a selective backup, all the files are candidates for back up unless you exclude them from backup, or they do not meet management class requirements for serialization. See “Supported file systems” on page 53 for supported file systems, ACL, and Extended Attribute support.

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:

```
selective /Users/user1/Documents/dir1/bogus.txt
```

still backs up Documents/dir1 even if the file bogus.txt does not exist.

If the **selective** command is retried because of a communication failure or session loss, the transfer statistics will display 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.

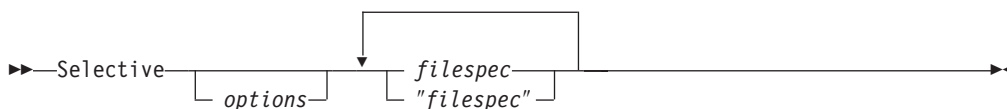
Removing operand limits

You can use the *removeoperandlimit* option to specify that Tivoli Storage Manager removes the 20-operand limit for the Mac OS X platform. 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. See “Removeoperandlimit” on page 222.

Migrating to Unicode-enabled file spaces

See “Autofsrename” on page 135 for information about using the Unicode-enabled client.

Syntax



Parameters

options

Table 63. Selective command: Related options

Option	Where to use	Page
<i>changingretries</i>	Client system options file or command line.	138
<i>compressalways</i>	Client user options file (dsm.opt) or command line.	143
<i>compression</i>	Client user options file (dsm.opt) or command line.	144
<i>dironly</i>	Command line only.	153
<i>filelist</i>	Command line only.	171
<i>filesonly</i>	Command line only.	174
<i>removeoperandlimit</i>	Command line only.	222
<i>subdir</i>	Client user options file (dsm.opt) or command line.	247
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	249

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.

Examples

Task Back up the proja file in the /Users/user1/Documents directory.

Command: selective /Users/user1/Documents/proja

Task Back up all files in the /Users/user2/Documents directory whose file names begin with proj.

Command: selective /Users/user2/Documents/proj*

Task Back up all files in the /Users/user3/Documents directory whose file names begin with proj. Back up the single file named budget in the /Users/user3/Documents/proja directory.

Command: selective /Users/user3/Documents/proj*
/Users/user3/Documents/proja/budget

Task Back up the /Users file system.

Command: selective /Users/ -subdir=yes

Set Access

The **set access** command gives users at other nodes access to your backup versions or archived copies. You can give another user access to a specific file, multiple files, 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.

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.

Syntax

►►—SET Access—
 └─ Archive ─┘ filespec— node ─┘
 └─ Backup ─┘ └─ user ─┘

Parameters

Archive

Permits access to archived files.

Backup

Permits access to backup versions of files.

filespec

Specifies the path, file or directory to which you are giving access to another node or user. Use wildcard characters to specify a group of files, or all files in a directory; all objects in a directory branch; or all objects in a volume. Use a single asterisk "*" for the file spec to give access to all files 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 `/Users/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 `/Users/test`, use file spec `/Users/test/*/*`.

Attention: Use of the form `/*/*` alone will not give access to objects in the named directory; only those in directories below the named directory will be 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: `/Users/sub1/subsub1`.
- The `/Users` directory contains the `h1.txt` and `h2.txt` files.
- The `/Users/sub1` directory contains file `s1.htm`.
- The `/Users/sub1/sub2` directory contains the `ss1.cpp` file.

To allow access to all files in the `/Users/sub1/sub2` directory, enter:

```
set access backup /Users/sub1/sub2/* * *
```

To allow access to only those files in the `/Users` directory, enter:

```
set access backup /Users/* * *
```

To allow access to all files in all directories *in* and *below* the `/Users` directory, enter:

```
set access backup /Users/* * *  
set access backup /Users/*/* * *
```

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.

Examples

Task Give the user at `node_2` authority to restore the budget file from the `/Users/user1/documents` directory.

Command: `set access backup /Users/user1/documents/budget node_2`

Task Give `node_3` authority to retrieve all files in the `/Users/user2/documents/proja` directory whose file names end with `.`

Command: `set access archive /Users/user2/documents/proja/*.c node_3`

Task Give `node_3` the authority to retrieve all files in the `/Users/user3/proja` directory.

Command: `set access archive /Users/user3/proja/ node_3`

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 will be rejected with an error message indicating that the current server does not support event-based policy.

Syntax

```
► SET Event TYPE= Hold Release Activateretention <filespec> ►
► --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* (See “Dateformat” on page 146)
- *Numberformat* (See “Numberformat” on page 204)
- *Noprompt* (See “Noprompt” on page 203)
- *Subdir* (See “Subdir” on page 247)
- *Timeformat* (See “Timeformat” on page 260)

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 will show 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

Set Password

System Administrator

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.

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

Syntax

►►—SET Password—┐
└ *oldpw newpw* ┘◄◄

Parameters

oldpw
Specifies the current password for your workstation.

newpw
Specifies the new password for your workstation.

Examples

The following is an example of using the **set password** command.

Task Change your password from osecret to nsecret.

Command: set password osecret nsecret

Appendix. Using the Tivoli Storage Manager central scheduler

This section discusses how to use the Tivoli Storage Manager scheduler.

Overview of the Tivoli Storage Manager scheduler

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 will take 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 will be executed for a node. The following conditions must be met before automatic scheduled events will 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 (see "Setting the client scheduler process to run as a background task and start automatically at boot time" on page 339 for more information).

When creating a schedule definition on the Tivoli Storage Manager server, schedule actions that you can take include incremental, selective, archive, restore, retrieve, 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 drives defined by the client domain option. A schedule definition using the *command* action allows an operating system command or shell script to be executed.

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 will 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 will 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.

Handling 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 quotes around all of the specifications. Examples:

```
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 will ensure that /Users/user1/Documents/Some file.doc is treated as a single file name, as opposed to two separate files (/Users/user1/Documents/Some, and file.doc).

Note: The examples above apply to the Tivoli Storage Manager Version 5.3.2 client only. For those nodes using previous versions of the clients, the following examples apply. Assume that 'Macintosh' is the name of the startup volume.

```
objects="Macintosh:Users:user1:Documents:Some file.doc"
objects=
    "Macintosh:Users:user1:Documents:Some file.doc"
    "Macintosh:Users:user5:Documents:Another file.txt"
    Macintosh:Users:user3:Documents:noblanks.txt'
```

```
objects="Macintosh:Users:user1:My Directory With Blank Spaces:'"
```

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*

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 schedule's defined start time. 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 will process the new schedule first.

Understanding scheduler processing options

There are several processing options which impact the behavior of the client scheduler. On the Tivoli Storage Manager client, you can define most of these options in the client user options file (dsm.opt) or dsm.sys file. However, some of these options can be set globally on the Tivoli Storage Manager server for all Tivoli Storage Manager clients. The **Managing Throughput of Scheduled Operations** section of the *Tivoli Storage Manager Administrator's Guide* provides detailed information on all topics described in the section.

Option	Client defined	Server defined	Server global override
<i>manageservices</i>	X		
<i>maxcmdretries</i>	X		set maxcmdretries command
<i>maxschedsessions</i>		X	
<i>postschedulecmd</i> , <i>postnschedulecmd</i>	X		
<i>preschedulecmd</i> , <i>prenschedulecmd</i>	X		
<i>querschedperiod</i>	X		set querschedperiod command
<i>randomize</i>		X	
<i>retryperiod</i>	X		set retryperiod command
<i>schedcmddisabled</i>	X		
<i>schedlogname</i>	X		
<i>schedlogretention</i>	X		
<i>schedmode</i>	X		set schedmodes command
<i>sessioninitiation</i>	X	X (update node command)	
<i>tcpclientaddress</i>	X	X (also defined on server when <i>sessioninit=serveronly</i> as part of the node definition)	
<i>tcpclientport</i>	X	X (also defined on server when <i>sessioninit=serveronly</i> 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 will execute.

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*.

Handling return codes from *preschedulecmd* and *postschedulecmd* scripts

Beginning with Tivoli Storage Manager Version 5.1, the scheduler exhibits the following behavior when the *preschedulecmd* and *postschedulecmd* options are used:

- 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 will run. The administrative **query event** command with *format=detailed* option will show 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 will show 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 prevails. Therefore, if the scheduled operation completes with return code 0 or 4 and the *postschedulecmd* command fails, the administrative **query event** command will show that the event completed with return code 8. If the scheduled operation completes with return code 12, that return code prevails, and **query event** will show 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 command might exit with return code 3, if 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
```

See the following references for more information:

- “Postschedulecmd/Postnschedulecmd” on page 213
- “Preschedulecmd/Prenschedulecmd” on page 215
- “Client return codes” on page 96

Using the client acceptor to manage 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. These methods are compared as follows:

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.

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

Setting the client scheduler process to run as a background task and start automatically at boot time

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 boot time; not the scheduler process. For the traditional method, the scheduler process should be set up to start automatically at boot 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. 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).

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.

Changing processing options used by the scheduler service

When you configure the Tivoli Storage Manager central-scheduling services (the scheduler, the client acceptor, or the remote client agent), some of the processing options that you specify are defined in the Windows registry. The following options can also be specified in the client options file (`dsm.opt`).

- `nodename`
- `httpport`
- `tcpserveraddress`
- `tcpport`
- `webports`

When the client scheduler runs as a foreground process using the `dsmc sched` command, the options in the client options file are used. However, when the scheduler runs as a Windows service, the options in the registry are used instead. If you are using the scheduler service and change an option in the `dsm.opt` file, you must update the corresponding value in the registry as well.

Managing 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 will be missed if the first schedule does not complete before the end of the startup window of the second schedule. Suppose that most of the drives on your client system must be backed up daily, and that one drive 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 Mac OS X 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 will be 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.

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
```

Restarting the scheduler process on a remote system

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

The SystemStarter utility can be used to start, restart, or 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
```

It is also possible to stop the client acceptor daemon by finding the PID of the `dsmcad` application and sending it a signal with the `kill` command:

```
sudo ps -x | grep dsmcad
sudo kill <dsmcad PID>
```

Troubleshooting the Tivoli Storage Manager scheduler

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

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

You can limit query results to exception cases with:

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

Query results include a status field which gives a summary of the result for a specific event. By using the *format=detailed* option on the **query event** command, complete information for events displays, including the return code passed back by the Tivoli Storage Manager client. Table 64 summarizes the meaning of the event status codes which are likely to exist for a scheduled event that has already taken place:

Table 64. Scheduled client event status codes

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

Investigating abnormal scheduled events

If a particular client node displays multiple consecutive events with a result of *missed*, the client scheduler process is not running, has terminated, or is hanging. If a scheduled event is missed, but other consecutive scheduled events for that node show a result of *completed*, investigate the server activity log and the client schedule log to determine the cause. Scheduled events with a result of *failed*, have

encountered a processing error originating either from the Tivoli Storage Manager client or server. The first place to check is the server activity log during the time period in which the scheduled event was processed. The activity log shows Tivoli Storage Manager server processing errors and some client errors which have been remotely logged to the Tivoli Storage Manager server. If the explanation cannot be found in the activity log, check the client schedule log.

Checking the server activity log

When checking the server activity log, narrow the query results down to the time period in which the scheduled event occurred. Begin the event log query at a time shortly before the start window of the scheduled event. For example, if you are investigating the following suspect event:

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

You could use one of the following queries:

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

Inspecting the Client Schedule Log

The Tivoli Storage Manager client maintains a detailed log of all scheduled activities. If queries of the server activity log do not explain a failed scheduled event, check the Tivoli Storage Manager client schedule log. Access to the client system is required for inspecting the schedule log. The schedule log file typically resides in the same directory as the Tivoli Storage Manager client error log (dsmerror.log), in a file named TSM Schedule Log. The default location for this log is /Library/Logs/tivoli/tsm. The location of the log file can be specified using the *shedlogname* option, so you might need to refer to the options file to see if the *shedlogname* option was used to relocate the log file. When you locate the schedule log, it is easy to search through the file to find the time period corresponding the start date and time of the scheduled event. Here are some tips on what to look for:

- If you are investigating a *missed* event, check the details of the previous event, including the time at which the previous event completed.
- If you are investigating a *failed* event, look for error messages which explain the failure, such as an exceeded Tivoli Storage Manager server session limit.
- If an explanation remains unclear, check the client error log file (usually named dsmerror.log.)

Understanding the scheduler log

The scheduler log file provides information about which schedule runs next and which files are processed. The default file name, **TSM Scheduler Log**, is created in the application folder. Use the *shedlogname* option to change the location and name of this file.

The schedule log continues to grow in size unless you prune it using the *shedlogretention* option or specify a maximum size using the *shedlogmax* option.

When the scheduler queries the server for the next schedule, it sends the schedule information to the log file. The following information is logged:

Table 65. Schedule log information

Category	Description
Schedule name	Specifies the name the administrator assigns to the schedule.
Action	Specifies the type of task to perform on the files listed in the Objects field: <ul style="list-style-type: none"> • Incremental. Incrementally backs up the files. • Selective. Backs up the files. • Archive. Archives the files. • Restore. Restores the files. • Retrieve. Retrieves the files. • Command. Executes a client operating system command described in the Objects field. • Macro. Executes a macro whose file name is described in the Objects field.
Objects	Names the volumes or files to process. If the action is a backup, archive, retrieve, or restore operation, the objects are file spaces, directories, or logical volumes. If this field is empty for an incremental backup, then the default is used. If the action is to run a command or macro, the object is the name of the command or macro to run.
Options	Overrides options that you set for scheduling in your Preferences file. See “Specifying scheduling options” on page 93 for more information about scheduling options.
Server window start	Specifies the date and time when the next schedule will run. If the server uses randomization for schedules, the schedule then runs <i>after</i> the specified time.

Output from scheduled commands is sent to the schedule log file. After scheduled work is performed, check the log to ensure the work completed successfully.

When a scheduled command is processed the schedule log contains the following entry:

Scheduled event *eventname* completed successfully

This is merely an indication that Tivoli Storage Manager successfully issued the scheduled command associated with the *eventname*. No attempt is made to determine the success or failure of the command. You should assess the success or failure of the command by evaluating the return code from the scheduled command in the schedule log. The schedule log entry for the command’s return code is prefaced with the following text:

Finished command. Return code is:

You can modify the scheduling options in your Preferences file if you do not like the current values and the options are not over-ridden by server settings for your schedule. You can determine this by comparing a ‘dsmc q options’ output to your local options files. For more information about scheduling options, see “Scheduling options” on page 118.

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 give you any license to these patents. You can send license inquiries, in writing, to:

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

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

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, 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 document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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 is for planning purposes only. The information herein is subject to change before the products described become available.

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.

If you are viewing this information in softcopy form, the photographs and color illustrations not appear.

Trademarks

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX	Redbooks
C Set ++	RISC System/6000
DFS	S/390
FlashCopy	SP
HACMP	System/390
i5/OS	System i
IBM	System p
IBMLink	Tivoli
Lotus	Tivoli Enterprise
Multiprise	Tivoli Enterprise Console
OpenPower	TotalStorage
OS/390	VisualAge
OS/400	WebSphere
POWER	xSeries
Power PC	z/OS

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.

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

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.

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

Glossary

The terms in this glossary are defined as they pertain to the IBM Tivoli Storage Manager library. If you do not find the term you need, refer to the IBM Software Glossary on the Web at this address: <http://www.ibm.com/ibm/terminology/>. You can also refer to IBM Dictionary of Computing, New York: McGraw-Hill, 1994.

This glossary include terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright (ANSI). You can purchase copies from the American National Standards Institute, 11 West 42nd Street, New York, New York 10036.
- The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC2/SC1).

A

absolute mode. 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 *mode*. Contrast with *modified mode*.

access control list (ACL).

1. In computer security, a collection of all access rights for one object.
2. 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 a list that is associated with a file that identifies the users who can access the file and that identifies the users' access rights to that file.

ACL. See *access control list*.

active policy set. The activated policy set that contains the policy rules currently in use by all client nodes assigned to the policy domain. See also *policy domain* and *policy set*.

active version. The most recent backup copy of a file stored by Tivoli Storage Manager. 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 workstation. Contrast with *inactive version*.

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*.

administrator. A user who has been registered to the server. Administrators can be authorized to one or more of the following administrative privilege classes: system, policy, storage, operator, or analyst. Administrators can use the administrative commands and queries allowed by their privileges.

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 data transfer rate. Dividing the total number of bytes transferred by the elapsed processing time calculates the data transfer rate.

archive. To copy one or more files to a storage pool for long-term storage. Archived files can include descriptive information, and you can retrieve them by archive date, file name, or by description. Contrast with *retrieve*.

archive copy. A file or group of files that have been archived to the Tivoli Storage Manager server.

archive copy group. A policy object containing attributes that control the generation, destination, and expiration of archived files. An archive copy group belongs to a management class.

archive retention grace period. The number of days that Tivoli Storage Manager retains an archived copy when the server is unable to rebind the file to an appropriate management class.

authentication. The process of checking and authorizing a user's password before allowing that user to access the Tivoli Storage Manager server. An administrator with system privilege can enable or disable authentication.

authorization rule. A specification allowing another user to either restore or retrieve a user's files from Tivoli Storage Manager 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. An authorized user is any user running with a real user ID of 0 (root) or a user who owns an executable whose owner execution permission bit is set to **s**. In the following example, the user **tivoli** is an authorized user while running **dsmc** since the **dsmc** owner execution permission bit is set to **s**:

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

A system administrator becomes an authorized user by using the TSM Tools for Administrators.

B

back up. To copy information to another location to ensure against loss of data. In IBM Tivoli Storage Manager, you can back up user files, the IBM Tivoli Storage Manager database, and storage pools. Contrast with *restore*. See also *incremental backup* and *selective backup*.

backup-archive client. A program that runs on a file server, PC, or workstation 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 controlling the generation, destination, and expiration of backup versions of files. The backup copy group belongs to a management class.

backup retention grace period. The number of days Tivoli Storage Manager retains a backup version when the server is unable to rebind the file to an appropriate management class.

backup set. A collection of active files in your file spaces that reside on the Tivoli Storage Manager server. The Tivoli Storage Manager administrator creates the backup set and copies it onto portable media device that is supported by the Tivoli Storage Manager server and client.

backup version. A file that a user 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*.

binding. The process of associating a file with a management class name. See also *rebinding*.

boot. To prepare a computer system for operation by loading an operating system.

C

CAD. Client acceptor daemon

central schedule. A function that allows 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*.

client. A program running on a PC, workstation, file server, LAN server, or mainframe that requests services of another program, called the server. The following types of clients can obtain services from a Tivoli Storage Manager server: administrative client, application client, API client, backup-archive client, and HSM client (also known as Tivoli Storage Manager for Space Management). See *administrative client* and *backup-archive client*.

client acceptor. The program that starts the client acceptor process is called **dsmcad**. A system administrator can use the client acceptor to launch the Tivoli Storage Manager scheduler during boot up.

client domain. The set of drives, file systems, or volumes that a user selects to back up or archive 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 options set. Client option sets allow the Tivoli Storage Manager administrator to specify additional options that not be included in the client options file, (*dsm.opt*). Client option sets are used in conjunction with client option files on client nodes.

client-polling scheduling mode. A client and server communication technique where the client node queries the server for scheduled work. Contrast with *server-prompted scheduling mode*.

client/server. A communications network architecture in which one or more programs (clients) request computing or data services from another program (the server).

client system options file. An editable file that contains communication, authorization, central scheduling, backup, archive, and space management options. A root user on your workstation sets the options in a client system options file. The file name is *dsm.sys* located in your Tivoli Storage Manager installation directory.

client user options file. A user-editable file containing options that identify the Tivoli Storage Manager server to contact, specify backup, archive, restore, retrieve, and space management options, and set date, time, and number formats. The file name is **dsm.opt** located in your Tivoli Storage Manager installation directory.

closed registration. A registration process in which a Tivoli Storage Manager administrator must 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.

command-click (Macintosh). Click an item while pressing the Command key. In file and folder lists, this deselects individual items embedded in a group of selected items.

command-line interface. A type of user interface where commands are specified on the command line. Contrast with *graphical user interface*.

communication method. The method by which a client and server exchange information. For Tivoli Storage Manager backup-archive clients, the method can be TCP/IP. See *Transmission Control Protocol/Internet Protocol*.

communication protocol. A set of defined interfaces that permits computers to communicate with each other.

copy group. A policy object that contains attributes that control backup and archive file:

- Generation
- Destination
- Expiration.

Backup and archive copy groups belong to management classes. See *frequency*, *destination*, *mode*, *retention*, *serialization*, and *version*.

D

default management class. A management class assigned to a policy set. This class is used to govern backed up or archived files when a user does not explicitly associate a file with a specific management class through the include-exclude list.

destination. A copy group attribute that specifies the storage pool in which to back up or archive a file. At installation, Tivoli Storage Manager provides two storage destinations named **backuppool** and **archivepool**.

domain. See *policy domain* or *client domain*.

drag. Move the mouse while holding down the mouse button, thus moving the selected object.

drag-and-drop. Move (drag) an object on top of another object and release the mouse button, thus relocating the object.

dynamic. A copy group serialization value that specifies Tivoli Storage Manager accept the first attempt to back up or archive an object, regardless of any changes made during backup or archive processing. See *serialization*. Contrast with *shared static* and *static*.

E

error log. A text file written on disk that contains Tivoli Storage Manager processing error messages. The Tivoli Storage Manager server detects and saves these errors.

exabyte (EB). (1) For processor storage, real and virtual storage, and channel volume, 1 152 921 504 606 846 976 bytes. (2) For disk storage capacity and communications volume, 1 000 000 000 000 000 000 bytes.

exclude. To identify files in an include-exclude list that you do not want to include in a specific client operation, such as backup or archive.

expiration. The process in which files are identified for deletion because their expiration date or retention period is passed. Backups or archives are marked for deletion based on the criteria defined in the backup or archive copy group.

expiring file. A migrated or premigrated file that is marked for expiration and removal from Tivoli Storage Manager 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. It expires and is removed from Tivoli Storage Manager storage after the number of days specified with the *migfileexpiration* option have elapsed.

F

file server. A dedicated computer and its peripheral storage devices connected to a local area network that stores both programs and files shared by users on the network.

file space. A logical space on the Tivoli Storage Manager server that contains a group of files. In Tivoli Storage Manager, users can restore, retrieve, or delete file spaces from Tivoli Storage Manager storage. A file space for systems:

- **Windows**— file spaces for removable media are identified by volume label. Fixed drive file spaces are identified by Universal Naming Convention (UNC) name.
- **UNIX** — Logical space that contains a group of files backed up or archived from the same file system, or part of a file system defined with the *virtualmountpoint* option in the client system options file.

frequency. A copy group attribute that specifies the minimum interval, in days, between incremental backups.

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. An archive copy of a file that might not accurately reflect what is currently in the file because Tivoli Storage Manager archived the file while the file was being modified.

G

generate password. Processing that 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 *options file*.

gigabyte (GB). (1) One billion (10⁹) bytes. (2) When referring to memory capacity, 1 073 741 824 in decimal notation.

globally unique identifier (GUID). A 16-byte code that identifies an interface to an object across all computers and networks. The identifier is unique because it contains a time stamp and a code based on the network address that is hard-wired on the host computer's LAN interface card.

graphical user interface (GUI). A graphical user interface offers pictorial rather than text-based access to a computer. A graphical user interface includes:

- A combination of graphics and icons
- Use of a mouse or pointing device
- Menu bars, dropdown lists, and overlapping windows

Contrast with *command-line interface*. See *windowed interface*.

GUI. Graphical user interface.

GUID. See *globally unique identifier*.

H

hierarchical storage management client. A program that runs on a workstation or file server to provide space management services. The hierarchical storage management client automatically migrates eligible files to Tivoli Storage Manager storage to maintain specific levels of free space on local file systems. Automatic recalls are made for migrated files when they are accessed. Users are also permitted to migrate and recall specific files.

HSM. Hierarchical Storage Management.

I

inactive version. A copy of a backup file in Tivoli Storage Manager storage that either is not the most recent version, or the corresponding original file was deleted from the client file system. Inactive backup versions are eligible for expiration according to the management class assigned to the file.

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 *include-exclude list*.

include-exclude list. A list of include and exclude 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. The include-exclude list is defined in one or more include-exclude files or in the *dsm.sys* file. See *client system options file*.

incremental backup. A function that permits user to back up new or changed files or directories from a client domain or from specified directories or files. These directories or files are not excluded in the include-exclude list and meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class assigned to each file. Contrast with *selective backup*.

inode. A data structure that describes the individual files in an operating system. There is one inode for each file. The number of inodes in a file system, and therefore the maximum number of files a file system can contain, is set when the file system is created. Hard linked files share the same inode.

inode number. A number that specifies a particular inode in a file system.

IPL. Initial Program Load. See *boot* and *reboot*.

L

LAN. Local area network.

Local Area Network (LAN). A variable-sized communications network placed in one location. LAN connects servers, PCs, workstations, a network operating system, access methods, and communications software and links.

logical unit number (LUN). A logical unit number (LUN) is a unique identifier used on a SCSI bus that enables it to differentiate between up to eight separate devices (each of which is a logical unit). Each LUN is a unique number that identifies a specific logical unit, which be a hard disk, tape drive, or other device which understands the SCSI protocol.

M

management class. A policy object that is a named collection of copy groups. A management class is associated with a file to specify how the server should manage backup versions or archive copies of workstation files. See *binding* and *copy group*.

mode. A copy group attribute that specifies whether a backup file should be created for a file that was not modified since the last time the file was backed up. See *absolute* and *modified*.

modified. A backup copy group attribute indicating a file is considered for backup only if the file has been changed since the last backup. A file is considered changed if the date, size, owner, or permissions have changed. See *absolute* and *mode*.

N

network data transfer rate. The data transfer rate calculated by dividing the total number of bytes transferred by the data transfer time. For example, the time spent transferring data over the network.

node. See *client node*.

node name. A unique name used to identify a workstation, file server, or PC to the server.

non-Unicode file space. Non-Unicode file spaces and file names are limited to the character set of the current locale when the files were backed up.

O

open registration. A registration process in which users can register their own workstations or PCs as client nodes with the server. Contrast with *closed registration*.

owner. The owner of backup-archive files sent from a multi-user client node, such as UNIX, Linux, or Mac OS X.

P

pattern-matching character. See *wildcard character*.

plug-in. A self-contained software component that modifies (adds or changes) function in a particular software system. When you add a plug-in to a software system, the foundation of the original software system remains intact.

policy domain. A policy object that contains one or more policy sets. Client nodes are associated with a policy domain. See *policy set*, *management class*, and *copy group*.

policy set. A policy object that contains a group of management class definitions that exist for a policy domain. At any one time, there can be many policy sets within a policy domain, but only one policy set can be active. See *active policy set* and *management class*.

progress indicator. A control used to inform a user about the progress of a process.

R

rebinding. The process of associating a backup with a new management class name. For example, rebinding occurs when the management class associated with a file is deleted. See also *binding*.

reboot. To restart the operating system.

registration. The process of identifying a client node or administrator to the server by specifying a user ID, password, and contact information. For client nodes, a policy domain, compression status, and deletion privileges are also specified.

restore. A function that permits users to copy a version of a backup file from the storage pool to a workstation or file server. The backup copy in the storage pool is not affected. Contrast with *backup*.

retention. The amount of time, in days, that inactive backed up or archived files are retained in the storage pool before they are deleted. The following copy group attributes define retention: retain extra versions, retain only version, retain version.

retrieve. A function permitting users to copy an archived file from the storage pool to the workstation or file server. The archive copy in the storage pool is not affected. Contrast with *archive*.

S

scheduling mode. The type of scheduling operation for the client-server node. Tivoli Storage Manager supports two scheduling modes: client-polling and server-prompted.

scroll. Move through a list of items in a window by operating the scrollbars with the mouse cursor.

select. Choose an item from a list or group of items.

selective backup. A function permitting users to back up specified files. These files are not excluded in the include-exclude list and meet the requirement for serialization in the backup copy group of the management class assigned to each file. Contrast with *incremental backup*.

serialization. A copy group attribute that specifies whether a file can be modified during a backup or archive operation. See *static*, *dynamic*, *shared static*, and *shared dynamic*.

server. A program running on a mainframe, workstation, or file server that provides shared services such as backup and archive to other various (often remote) programs (called clients).

server-prompted scheduling mode. A client-server communication technique where the server contacts the client node when tasks need to be done. Contrast with *client-polling scheduling mode*.

session. A period of time in which a user can communicate with a server to perform backup, archive, restore, or retrieve requests.

shared dynamic. A Tivoli Storage Manager copy group serialization mode. This mode specifies if a file changes during backup or archive and continues to change after a number of retries. The last retry commits the file to the Tivoli Storage Manager server whether or not the file changed during backup or archive. Contrast with *dynamic*, *shared static*, and *static*.

shared static. A copy group serialization value specifying 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 *serialization*. Contrast with *dynamic*, *shared dynamic*, and *static*.

share point. A drive or directory on Windows XP and .NET whose files are available for shared access across a network. The share point name is part of a UNC name. See *Universal Naming Convention (UNC)* name.

shift-click. Click an item while pressing the Shift key.

space management. The process of keeping sufficient free storage space available on a local file system for new data and making the most efficient and economical use of distributed storage resources.

stabilized file space. A file space that exists on the server but not on the client. This situation can arise in at least two instances:

1. A drive is removed from a client workstation
2. A file space is renamed on the server

Stabilized file spaces remain on the server until deleted by the user or administrator. Files and directories can be restored and retrieved from a stabilized file space. However, it is not possible to back up or archive data to a stabilized file space.

stanza. On UNIX, Linux, and Mac OS X, a stanza is a group of lines in a file that together have a common function or define a part of the system. The Tivoli Storage Manager Client System Options file (dsm.sys) contains a stanza for each server to which the client can connect. Each stanza begins with the *servername* option and ends at the next *servername* option or the end of file. Each stanza must include communications options.

static. A copy group serialization value specifying that a file must not be modified during a backup or archive operation. If the file is in use during the first attempt, Tivoli Storage Manager will not back up or archive the file. See *serialization*. Contrast with *dynamic*, *shared dynamic*, and *shared static*.

storage pool. A named set of storage volumes used as the destination of backup, archive, or migrated copies.

T

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 being operated on.

TCP/IP. Transmission Control Protocol/Internet Protocol.

timeout. A time event involving:

- An event that happens at the end of a predetermined period of time that began at the happening of another specified event.
- A time interval allotted for certain operations to happen. For example, response to polling or addressing before system operation is interrupted and must be restarted.
- A terminal feature that logs off a user if an entry is not made within a specified period of time.

Tivoli Storage Manager. A client-server licensed program product that provides storage management and data access services to customers in a multivendor computer environment.

Transmission Control Protocol/Internet Protocol (TCP/IP). A standard set of communication protocols that supports peer-to-peer connectivity of functions for both local and wide-area networks.

Trusted Communications Agent (TCA). A program that can handle the sign-on password protocol when password access is generated. The main process (for example, *dsmc*) makes a run time decision based on the password access option setting, the user ID, and the executables' access privileges to run this program. The file that contains this program (*dsmtca*) must have the 's' bit set in its mode field and the owner must be root.

V

version. Storage management policy allow back-level copies of backed up objects to be kept at the server whenever an object is newly backed up. The most recent backed up copy is called the "active" version. Earlier copies are "inactive" versions. The following backup copy group attributes define version criteria: versions data exists, and versions data deleted.

W

wildcard character. An asterisk (*) or question mark (?) character used to represent multiple (*) or single (?) characters when searching for various combinations of characters in alphanumeric and symbolic names.

windowed interface. A type of user interface that is either a graphical user interface or a text-based interface. The text-based interface maintains a close affinity to the graphical user interface, including action bars and their associated pull-down menus and windows. See *graphical user interface*.

workstation. A programmable high-level workstation (usually on a network) with its own processing hardware such as a high-performance personal computer. In a local area network, a personal computer that acts as a single user or client. A workstation can also be used as a server.

world wide name. A unique 48 or 64 bit number assigned by a recognized naming authority (often with block assignment to a manufacturer) that identifies a connection or a set of connections to the network. Abbreviated WWN. A WWN is assigned for the life of a connection (device). Most networking technologies (e.g., Ethernet, FDDI, etc.) use a world wide name convention.

Index

Special characters

? 69

* 69

Numerics

128-bit AES encryption support 52

A

absolute mode 104

access

permissions, archive 89

accessibility

disability xiii

keyboard xiii

shortcut keys xiii

active backup versions

displaying 44, 71, 301

restoring 71

administrative client

allowing secure sessions within a private network 251

archive

a list of files 87

assign description to 150

assigning description on command line 87

authorize users on your workstation to request services

from server 265

binding management class to 106, 128

binding management classes to files 87

client node proxy 87

command 279

compressing files 144

configuring client node proxy support 87

copy group 100

copy group attributes 101

copy mode 104

delete file spaces 82

delete files after 149

delete individual archives from server file space 89, 283

display details 151

display the last modification date and 299

encrypting data during 52

encryptiontype 161

exclude objects from 168

files only 174, 266

grace period retention 100

hard links 89

how managed 91, 99

include files for 186

information, query 299

list of files 171

managing file growth during compression 143

maximum file size for 85

more than one file specification 87

only files; not directories 87

overriding management class during 106

primary tasks 85

process directories only (not files) 153

purpose 47

archive (*continued*)

query user access 298

regulate resources server and client use during

processing 225

removeoperandlimit 222

retrieving using client GUI 90

retrieving using command line 91

revoke access 282

running 85

shared data on multiple clients under a single node

name 87, 129

specify end date and time to search for 261, 262

specifying number of attempts to archive open files 138

specifying whether to include subdirectories for 247

specifying whether to wait for tape to mount if it is

required 249

starting a Web client session 42

subdirectories 87

summary of options 114

suppress confirmation prompt before deleting 203

testing include-exclude list 34

using commands 89

using the client command line 86

archive with client node proxy

agent node 87

target node 87

archmc option 128

asnodename option 129

Attention

agent nodes and platform type 88

DSM_LOG environment variable 232, 327

DSM_LOG environment variable name 19

dual-boot system problems 50

incremental-by-date backup dates 190

securing local backup sets on portable media 78

sessioninitiation option and serveronly 242

Using the asterisk (*) and not the question mark (?) for
multibyte code pages 69, 277

when Tivoli Storage Manager cannot write to the log
file 19

auditlogging option 131

auditlogname option 133

authentication

Tivoli Storage Manager client 39

authorization

summary of options 120

authorized user

definition ix

tasks 13

authorizing

changing authorizations 80

groups to use Tivoli Storage Manager 179

Tivoli Storage Manager users 265

users to restore or retrieve files 80

autofsrename option 135

automating backup services

displaying scheduled work 93, 95

options for 93

overview 93, 335

process commands after back up 213

process commands before back up 215

automating backup services (*continued*)
starting client scheduler 22

B

backup

- advanced considerations 66
- authorize users on your workstation to request services from server 265
- binding management class to files 106
- client node proxy 63
- comparison: incremental, incremental-by-date 58
- compressing files 144
- configuring support for client node proxy backups 63
- copy group 57, 100
- copy mode 103
- delete file spaces 82
- delete individual backups from server file space 61
- delete individual files or images from server file space 285
- display details 151
- displaying processing status 62
- encrypting data during 52
- encryptiontype 161
- estimating processing time 59
- exclude objects from 168
- excluding domains 60, 157
- excluding system objects 169
- files only 174
- files, assigning management class 105
- fuzzy 103
- grace period retention 100
- hard links 67
- how Tivoli Storage Manager manages 52
- inactivate a list of 289
- include files for 186
- include-exclude processing for 28
- incremental
 - client domain 157
 - folders, overview 57
 - new and changed files with modification date later than last backup 190
 - process a list of files 171
 - using command line 60
- incremental-by-date 57
 - folders, processing overview 57
 - using command line 60
- incremental-by-date, using client command line 60
- incremental-by-date, using client Java GUI 59
- incremental, specify memory-conserving algorithm for 199
- incremental, using client command line 60
- incremental, using client Java GUI 59
- managing 99
- managing file growth during compression 143
- maximum file size 54
- migration of file spaces to Unicode 50
- multi-session, send files contiguously to the server 139
- new or changed files 56
- open files 68
- overview 47, 56
- planning 47
- pre-backup considerations 48
- primary tasks 47
- process directories only (not files) 153
- purpose 47
- query information 301

backup (*continued*)

- query user access 298
- regulate resources server and client use during processing 225
- removeoperandlimit 222
- restore 71
- revoke access 282
- saving encryption key password 162
- selective
 - back up list of files 60
 - overview 58
 - using command line 60
- selective, using client command line 60
- selective, using client Java GUI 59
- shared data on multiple clients under a single node name 129
- skip acl processing 244
- skip acl update checking 245
- specify end date and time to search for 261, 262
- specifying number of attempts to backup open files 138
- specifying one server session per file specification 139
- specifying whether to include subdirectories for 247
- specifying whether to wait for tape to mount if it is required 249
- starting a Web client session 42
- subdirectories 60
- summary of options 114
- testing include-exclude list 34
- using include-exclude list 51
- using multiple sessions 62
- volume name 48

backup copy group

- attributes 101

backup set

- enabling GUI for local restore of 193
- specify location for restore or query 194

backup with client node proxy

- agent node 63
- overview 63
- target node 63

backupset

- enabling GUI for local restore 77
- restore 77

backupsetname option 137

Bash shell variables

- setting 20

batch mode 273

- starting a session 41

bottom-up processing

- include-exclude list 34

Bourne and Korn shell

- pointing to client user options file 16

Bourne shell variables

- setting 20

C

C shell variables

- setting 20

cancel restore command 281

central scheduling

- summary of options 118

changingretries option 138

classic (standard) restore 72

client

- registering with server 25
- setting password 26

- client (*continued*)
 - specify size for the TCP/IP sliding window for client node 259
 - specifying client TCP/IP address other than the one for first server contact 254
 - specifying client TCP/IP port number other than the one for first server contact 255
- client acceptor daemon
 - manage scheduler 196
 - use to manage scheduler 196
- Client Acceptor daemon
 - configuring to manage scheduler 21
- client command options
 - overview 274
- client components 7
- client node proxy archive
 - overview 87
- client node proxy backup
 - scheduling 65
- client node proxy support 63, 87
- client options
 - display current settings 18, 309
 - exclude
 - exclude.archive 28
 - exclude.attribute.symblink 28
 - exclude.backup 28
 - exclude.compression 28
 - exclude.dir 28
 - exclude.file 28
 - exclude.file.backup 28
 - guitreeviewafterbackup 180
 - order of processing (precedence) 123
 - overriding using command line 123
 - overview 274
 - setting in a file 18
 - using with commands 123
- client options file
 - creating and modifying 15
 - required options for 15
- client options reference 127
- client scheduler
 - displaying scheduled work 93, 95
 - options for 93
 - overview 93, 335
 - starting 22, 326
 - starting automatically 43
- client system options file
 - copying and modifying 15
 - example of 15
 - minimum required statements 15
 - setting options 18
 - specifying include-exclude options 26
- client user options
 - setting options in 18
- client user options file
 - customizing 17
 - overriding using commands 123
 - setting options 18
- client-server communication
 - establishing 15
 - identify your workstation to the server 202
 - specify maximum disk I/O buffer size client uses when reading or writing files 155
 - specify size for the TCP/IP sliding window for client node 259
 - specifying client TCP/IP address other than the one for first server contact 254

- client-server communication (*continued*)
 - specifying client TCP/IP port number other than the one for first server contact 255
 - specifying default server 148
 - specifying method 140
 - specifying name of server to contact for services 240
 - specifying number of kilobytes client buffers before sending transaction to server 263
 - specifying reconnection attempts after failure 141
 - specifying reconnection interval after failure 142
 - specifying size of internal TCP/IP communication buffer 252
 - specifying TCP/IP address for dsmdcad 253
 - specifying TCP/IP address of Tivoli Storage Manager server 258
 - specifying TCP/IP port address of Tivoli Storage Manager server 257
 - specifying whether to send small transactions to server without buffering them first 256
- closed registration
 - permissions 26
 - using 26
- collocatebyfilespec option 139
- command line
 - archiving files 89
 - assigning description to archive 87
 - display current settings for client options 18, 309
 - displaying processing status 62
 - enabling command history capabilities 159
 - ending a session 45
 - entering commands 274
 - general rules when entering options with commands 123
 - options 124
 - overriding management class during archive 106
 - overview of parameters 274
 - performing large restore operations 72
 - performing point-in-time restore 77
 - recall commands 276
 - retrieving archived data 91
 - return codes for operations 96
 - rules for entering commands 277
 - specifying file specification 275
 - specifying options file during session 205
 - starting a session 41
 - using wildcard characters 277
- command parameters
 - overview 274
- command processing, summary of options 120
- command session
 - ending 272
 - starting 272
- commands
 - archive 279
 - batch mode 273
 - cancel restore 281
 - delete access 282
 - delete archive 283
 - delete backup 285
 - delete filespace 288
 - entering 274
 - entering on command line 274
 - expire 289
 - general rules when entering options with 123
 - help 290
 - incremental 291
 - interactive (loop) mode 273
 - loop 295

- commands (*continued*)
 - macro 296
 - maximum file specifications permitted 276
 - overview of parameters 274
 - preview 297
 - query access 298
 - query archive 299
 - query backup 301
 - query backupset 303
 - query filespace 305
 - query inlxl 306
 - query mgmtclass 307
 - query node 308
 - query options 309
 - query restore 310
 - query schedule 311
 - query session 312
 - query systeminfo 313
 - recall previous 276
 - restart restore 315
 - restore 316
 - restore backupset 320
 - retrieve 323
 - rules for entering 277
 - schedule 326
 - scheduled, enabling or disabling 96
 - selective 328
 - set access 330
 - set event 332
 - set password 334
 - specifying file specification 275
 - using 271
 - using in executables 96
 - using options with 123
 - using wildcard characters 277
- commmethod option 140
- commrestartduration option 141
- commrestartinterval option 142
- communication method
 - TCP/IP 7
- communication methods
 - summary 111
- communications
 - establishing through firewall 23
- compressalways option 143
- compression
 - disabling processing 170, 188
 - enabling processing 170, 188
 - include-exclude statements 170, 188
- compression and encryption processing
 - back up 188
 - exclude from backup 188
 - exclude options 188
- compression option 144
- compression processing
 - exclude from backup 170, 188
 - exclude options 170, 188
 - excluding files from 168
 - include files for 186
- configuring
 - the client scheduler 21
 - the Web client 20
- configuring Tivoli Storage Manager
 - optional tasks 13
 - required tasks 13
- console option 145
- copy destination attribute 104

- copy frequency attribute 102
- copy group name attribute 102
- copy groups 99
 - archive 100
 - backup 100
- copy mode attribute
 - absolute 103
 - modified 103
- copy serialization attribute 103
- copy type attribute 102
- customer support
 - contacting xiii

D

- date format
 - specifying 146
- dateformat option 146
- default client user options file
 - creating and modifying 16
 - example of 16
- default domain
 - excluding domains from backup 60, 157
- default management class 99
- default policy domain 99
- defaultserver option 148
- delete access command 282
- delete archive command 283
- delete backup command 285
- delete file space 82
- deleted file systems 68
- deletefiles option 149
- deleting
 - authorizations 80
 - individual archives from server file space 89, 283
 - individual backups from server file space 61, 285
- description option 150
- desktop
 - overlapping objects 82
- detail option 151
- diagnostics
 - options 122
- directories
 - assigning management class for 152
 - excluding 28
 - excluding from backup processing 168
 - specifying on command line 275
- dirmc option 152
- dironly option 153
- disability
 - accessibility xiii
 - keyboard xiii
 - shortcut keys xiii
- disablenqr option 154
- disaster recovery 82
- disk space requirements 7
- disk, restoring 82
- diskbuffsize option 155
- diskcachelocation option 156
- displaying
 - archive information 299
 - online help 45
 - restartable restore sessions 310
 - scheduled events 311
 - session information 312
- domain
 - back up using the GUI 59

- domain (*continued*)
 - include for incremental backup 157
 - specifying drives in the default 59
- domain option 157
- DSM_CONFIG
 - pointing to client user options file 16, 19
- DSM_DIR
 - pointing to executable files 19
 - pointing to resource files 19
 - pointing to system options (dsm.sys) file 19
- dsm_log 19
- DSM_LOG
 - set to point to dsmerror.log and dsmsched.log 19
- dsm.opt file
 - creating 16
 - creating and modifying 15
 - customizing 17
 - example of 16
 - required options for 15
 - specifying a drive specification using wildcards 32
- dsm.opt.smp file 16
- dsm.smp file
 - copying to dsm.opt 15
 - location 15
- dsm.sys file
 - creating 15
 - example of 15
 - setting options in 18
- dsm.sys.smp file 15
- dsmc executable file
 - set DSM_DIR to point to 19
- dsmerror.log 19
 - set DSM_LOG to point to 19
- dsmsched.log
 - set DSM_LOG to point to 19
- dual boot systems
 - duplicate names 49
- dynamic and shared serialization 103

E

- editor option 159
- enablearchiveretentionprotection option 160
- encryption
 - multiple clients under a single node name 129
 - of file data 52
 - saving encryption key password 162
- encryption processing
 - encryption methods available 52
 - include files for 186
 - query systeminfo command 313
- encryptiontype option 161
- encryptkey option
 - encryptkey=generate 162
 - encryptkey=prompt
 - encryptkey=save 162
- enhanced query schedule 93
- enhanced query schedule command 311
- environment prerequisites 6
- environment variables
 - setting 19
- error log
 - controlling the size 164
 - pruning 166
 - specifying path and file name 165
- error processing, summary of options 121
- errorlogmax option 164

- errorlogname option 165
- errorlogretention option 166
- estimate function 59
- event logging
 - scheduler 95
- event-based policy retention protection
 - archive 108
 - backup 108
- exclude options 168
 - exclude.archive 28
 - exclude.attribute.symblink 28
 - exclude.backup 28
 - exclude.compression 28
 - exclude.dir 28
 - exclude.file 28
 - exclude.file.backup 28
 - processing 34
 - testing 34
 - wildcard characters 31, 32
- excluding files
 - system files 30
 - using wildcard characters 32
 - wildcard characters 31
- executable file
 - return codes from 96
- expire command 289
- extended permissions
 - archive 89

F

- file space
 - determining fsID 151
 - excluding 28
- file spaces
 - names 48
- file specification
 - maximum allowed on commands 276
- file systems
 - deleted 68
 - excluding from backup processing 168
 - supported 53
- filelist option 171
- filename option 173
- files
 - archive a list of 87, 171
 - archive using commands 89
 - archived, overriding management class 106
 - archives, how managed 91
 - archiving 85, 279
 - archiving more than one file specification 87
 - archiving using the client command line 86
 - assigning management classes 67
 - backing up hard-linked 67
 - backing up open 68
 - backup, how managed 52
 - binding management classes to 106
 - compressing during archive or backup 144
 - definition of changed 56
 - delete after archive 149
 - delete individual archives from server file space 89, 283
 - delete individual backups from server file space 61, 285
 - encryption 52
 - excluding groups 31, 32
 - include-exclude
 - creating in Unicode format 184
 - including groups 31, 32

files (*continued*)

- managing growth during compression 143
- maximum file size for archive and retrieve 85
- maximum file size for backup and restore 54
- performing large restore operations 72
- processing include-exclude 34
- query archive information 299
- query backup information 301
- query user access 298
- renaming file spaces that are not Unicode to
Unicode-enabled 135, 291, 328
- restore, point-in-time 76
- restoring
 - to another workstation 81
- restoring hard-linked 67
- retrieve using client GUI 90
- retrieve using commands 91
- retrieving
 - to another workstation 81
- revoke access 282
- sorting list of 44
- filesonly option 174
- filespace
 - delete 82, 288
- filtering
 - files for restore 74
- firewall
 - establishing communications through 23, 181, 257
 - specify whether server or client initiates sessions
through 242
- folders
 - incremental backup processing overview 57
 - processing during incremental-by-date 57
- format and language
 - summary of options 120
- fromdate option 175
- fromnode option 176
- fromowner option 177
- fromtime option 178
- full incremental
 - comparing with incremental-by-date 58
 - definition 56
 - description 56
 - when to use 58
- fuzzy backup 103

G

- getting started
 - changing your password 39
 - client scheduler 39
 - command-line session 39
 - displaying online help 39
 - ending a session 39
 - GUI session 39
 - sorting file lists 39
- graphical user interface
 - changing password 43
 - delete individual files or images from server file
space 285
 - displaying active and inactive backup versions 44, 71
 - displaying online help 45
 - displaying processing status 62
 - enabling for local backup set restore 193
 - enabling local backupset 77
 - ending a session 45
 - retrieving archived data 90

graphical user interface (*continued*)

- starting a session 40
- using to back up objects 59
- groups option 179
- GUI
 - ending a session 45
 - overriding management class during archive 106
 - performing point-in-time restore 77
 - starting a session 40
- guitreeviewafterbackup option 180

H

- hard links
 - archive and retrieve 89
 - backing up 67
 - restoring 67
- hard mounts, NFS 67
- hardware requirements 7
- help
 - displaying online 45
 - Internet resources 45
 - online forum 45
 - service and technical support 45
- help command 290
- HFS
 - file systems supported 53
- HFS+
 - access control list (ACL) 53
 - ACL (see access control list) 53
 - file systems supported 53
- HFSX
 - file systems supported 53
- httpport option 181

I

- icons
 - overlapping 82
- ifnewer option 182
- image backup
 - deleting 285
- inactive backup versions
 - displaying 44, 71, 301
 - restoring 71
- inactive option 183
- inclexcl option 184
- include option
 - management class 105
 - processing 34
 - wildcard characters 31, 32
- include-exclude list
 - creating 26
 - overview 51
 - query order of processing 306
 - size restriction 31, 35
 - testing 34
- include-exclude options file
 - specifying path and file name of 184
 - to manage archives 91
 - Unicode-enabled file spaces 184
- include-exclude processing
 - options for 28
 - overview 28
- incrbydate option 190

- incremental
 - incremental-by-date 57
 - of folders
 - processing overview 57
- incremental backup
 - back up new and changed files with modification date later than last backup 190
 - by date 60
 - description 56
 - new and changed files 56
 - of folders
 - processing overview 57
 - overview 56
 - specify memory-conserving algorithm for 199
 - using client Java GUI 59
 - using command line 60
 - using the client command line 60
- incremental command 291
- incremental-by-date
 - back up using client Java GUI 59
 - back up using the client command line 60
 - comparing with incremental 58
 - description 57
 - overview 56
 - using command line 60
 - when to use 58
- installation requirements 6
- installing
 - system administrator task 1
 - Tivoli Storage Manager 1
- interactive mode 273
- interactive session
 - ending 295
 - starting 41, 295
 - using 295
- ISO 9600
 - file systems supported 53

J

- Java GUI
 - configuration restrictions 40

K

- keyboard
 - accessibility xiii
 - disability xiii
 - shortcut keys xiii

L

- language option 191
- last access date
 - specifying whether to update during backup or archive 57
- latest option 192
- local backupset
 - enabling GUI for local restore 77
- localbackupset option 193
- location option 194
- log
 - DSM_LOG environment variable 165, 232
 - error log, controlling the size 164
 - error log, pruning 164
 - error log, specifying path and file name 164

- log (*continued*)
 - errorlogname option 165
 - errorlogretention option 165
 - schedlogname option 232, 326
 - schedlogretention option 232, 326
 - specifying path and file name 165, 232, 326
- loop command 295

M

- Macintosh client
 - communication method 7
 - components 7
 - disk space requirements 7
 - environment prerequisites 6
 - hardware requirements 7
 - installation requirements 6
 - installing 7
 - memory requirements 7
 - operating system requirements 7
 - software requirements 7
- macro command 296
- mailprog option 195
- manageservices option 196
- management class
 - assigning 67
- management classes
 - assigning to directories 106, 152
 - assigning to files 105
 - binding archive files to 87
 - binding to files 106
 - default 99
 - displaying 101
 - displaying information about 307
 - how Tivoli Storage Manager uses 52
 - overriding during archive processing 106
 - overriding the default 105
 - processing 105
 - questions to consider 104
 - selecting for files 104
 - specifying with include option 105
 - using management class, example 105
- maxcmdretries option 198
- memory requirements 7
- memoryefficientbackup option 199
- messages
 - displaying on screen 267
 - specifying language type 191
 - stop displaying 221
- migrating from earlier versions of the product 1
- migrating to Unicode 4
- migration 4
- modes
 - batch 273
 - interactive (loop) 273
- modified mode 104
- multiple sessions
 - backup objects, using 62

N

- NFS
 - hard mounts 67
 - soft mounts 67
- nfstimeout option 67, 201
- no query restore 72

- node
 - specifying type to query 264
- node name 15
- node option, summary 112
- nodename option 81, 202
- non-authorized user
 - tasks 13
- noprompt option 203
- numberformat
 - specifying 204
- numberformat option 204

O

- online help
 - displaying 45
 - online forum 45
 - service and technical support 45
- open registration
 - permissions 26
 - using 26
- operating system requirements 7
- optfile option 205
- options
 - archive, summary 114
 - archmc 128
 - asnodename 129
 - auditlogging 131
 - auditlogname 133
 - authorization, summary 120
 - autofsrename 135
 - backup, summary 114
 - backupsetname 137
 - central scheduling, summary 118
 - changingretries 138
 - collocatebyfilespec 139
 - command processing, summary 120
 - commmethod 140
 - commrestartduration 141
 - commrestartinterval 142
 - communication, summary 111
 - compressalways 143
 - compression 144
 - console 145
 - dateformat 146
 - defaultserver 148
 - deletefiles 149
 - description 150
 - detail 151
 - diagnostics 122
 - dirmc 152
 - dironly 153
 - disablenqr 154
 - diskbuffsize 155
 - diskcachelocation 156
 - domain 157
 - editor 159
 - enablearchiveretentionprotection 160
 - encryptiontype 161
 - encryptkey
 - encryptkey=generate 162
 - encryptkey=prompt 162
 - encryptkey=save 162
 - errorlogmax 164
 - errorlogname 165
 - errorlogretention 166

- options (*continued*)
 - exclude
 - exclude.archive 28, 168
 - exclude.attribute.symblink 28
 - exclude.backup 28, 168
 - exclude.compression 28
 - exclude.dir 28, 168
 - exclude.file 28, 168
 - exclude.file.backup 28, 168
 - filelist 171
 - filename 173
 - filesonly 174
 - format and language, summary 120
 - fromdate 175
 - fromnode 176
 - fromowner 177
 - fromtime 178
 - general rules when entering with commands 123
 - groups 179
 - guitreeviewafterbackup 180
 - httpport 181
 - ifnewer 182
 - inactive 183
 - inclexcl 184
 - include.archive 186
 - include.backup 186
 - include.compression 186
 - include.encrypt 186
 - include.file 186
 - incrbydate 190
 - language 191
 - latest 192
 - localbackupset 193
 - location 194
 - mailprog 195
 - managedservices 196
 - maxcmdretries 198
 - memoryefficientbackup 199
 - nfstimeout 201
 - nodename 202
 - noprompt 203
 - numberformat 204
 - optfile 205
 - order of processing (precedence) 123
 - password 206
 - passwordaccess 207
 - passworddir 209
 - pick 210
 - pitdate 211
 - pittime 212
 - postnschedulecmd 213
 - postschedulecmd 213
 - prenschedulecmd 215
 - preschedulecmd 215
 - preservepath 217
 - querschedperiod 220
 - quiet 221
 - removeoperandlimit 222
 - replace 223
 - resourceutilization 225
 - restore and retrieve, summary 117
 - retryperiod 228
 - revokeremoteaccess 229
 - schedcmddisabled 230
 - schedcmduser (server defined only) 96
 - schedlogmax 231
 - schedlogname 232

- options (*continued*)
 - schedlogretention 233
 - schedmode 235
 - schedrestretrdisabled 237
 - scrolllines 238
 - scrollprompt 239
 - servername 240
 - sessioninitiation 242
 - skipacl 244
 - skipaclupdatecheck 245
 - specifying in commands 123
 - srvprepostscheddisabled 246
 - subdir 247
 - tapeprompt 249
 - tcpadminport 251
 - tcpbuffsize 252
 - tcpcadaddress 253
 - tcpclientaddress 254
 - tcpclientport 255
 - tcpnodelay 256
 - tcpport 257
 - tcpserveraddress 258
 - tcpwindowsize 259
 - timeformat 260
 - to date 261
 - to time 262
 - transaction processing, summary 121
 - txnbytelimit 263
 - type 264
 - users 265
 - v2archive 266
 - verbose 267
 - virtualnodename 268
 - webports 269
- overlapping desktop objects 82

P

- parameters
 - yes and no, alternatives 127
- partial incremental
 - definition 56
 - incremental by date, running 60
- password
 - changing 43, 334
 - number of characters 43
 - setting 206
 - setting for client 26
 - specifying directory location to store encrypted password
 - file 209
 - specifying whether to generate automatically or set as user
 - prompt 207
 - using 41
 - valid characters 43
- password option 206
- passwordaccess option 207
- passworddir option 209
- performance
 - transaction options 121
 - transaction processing 263
- permissions
 - access, saving standard and extended 89
- pick option 210
- pitdate 211
- pittime option 212
- point-in-time restore 76
- policies, storage management 99

- policy domains
 - default policy domain 99
 - standard policy domain 99
- policy sets
 - active policy set 99
- portable media
 - restoring backupsets 77
- postnschedulecmd option 213
- Preferences editor
 - excluding domains from back up 59
- prenschedulecmd option 215
- preschedulecmd option 215
- preserve path option 217
- preview command 297
- previewing
 - excluded files 297
 - included files 297
- processing aliases 169, 188
- processing options
 - authorization 120
 - backup and archive 114
 - central scheduling 118
 - communication 111
 - diagnostics 122
 - error processing 121
 - format and language 120
 - node option 112
 - overview 111
 - restore and retrieve 117
 - setting 18
 - specifying in commands 123
 - transaction processing 121
 - using 111
- processing symbolic links 169, 188
- processing time
 - estimating 59
- proxied session restrictions 64, 89

Q

- query
 - backups, establish point-in-time 211, 212
 - based on date and time of backup, archive 175, 178
 - display active and inactive objects 183
 - files for another node 176
 - include-exclude list 306
 - nodes to which client has proxy authority 87
 - nodes to which client has proxy node authority 63
 - process directories only (not files) 153
 - specify end date and time to search for backups or
 - archives 261, 262
 - specify location of backup set for 194
 - specifying amount of information that displays on
 - screen 238
 - specifying description for 150
 - specifying scrolling preferences after displaying
 - information on screen 239
 - specifying whether to include subdirectories for 247
 - system information 313
- query access command 298
- query archive command 299
- query backup command 301
- query backupset command 303
- query filespace command 305
- query inclexcl command 306
- query mgmtclass command 307
- query node command 308

- query options command 309
- query restore command 310
- query schedule
 - enhanced 93
- query schedule command 311
- query schedule command, enhanced 311
- query session command 312
- query systeminfo command 313
 - encryption processing 313
- querschedperiod option 220
- quiet option 221

R

- rebinding files to a different management class 107
- recall commands
 - limitations 276
- recommendations
 - default diskbuffsize setting 155
 - discontinue use of largecommbuffers xxi, 122
 - group include-exclude options in client options file 27
 - include-exclude list 31
 - long user and group names 54
 - long user names 265
 - minimum include-exclude list 30
 - path names for setting environment variables 17
 - set up SSL in the Integrated Solutions Console (ISC)
 - environment 25
 - specifying diskbuffsize xxi
 - target nodes and traditional nodes 64
 - to simplify retrieve 48
 - traditional scheduler system resources 339
- registering
 - client with server 25
 - using closed registration 26
 - using open registration 26
- removeoperandlimit option 222
- replace option 223
- resourceutilization option 225
- restart restore command 315
- restartable restore 72
- restartable restore sessions, display 310
- restore
 - active version 71
 - advanced considerations 80
 - another user's files 177
 - authorizing another user 80
 - backup set
 - supported tape devices 320
 - backups, establish point-in-time 211, 212
 - backupsets
 - from local device 74
 - from server 74
 - overview 77
 - based on date and time of backup 175, 178
 - classic (also known as standard) 72
 - create list of backup versions to 210
 - desktop objects overlapping 82
 - disk 82
 - display active and inactive objects 183
 - estimating processing time 59
 - files belonging to another node 81
 - files for another node 176
 - files only 174
 - filtering files 74
 - from file spaces that are not Unicode-enabled 317
 - restore (*continued*)
 - from portable media
 - overview 77
 - GUI, displaying active and inactive versions 44
 - hard links 67
 - image, suppress confirmation prompt 203
 - inactive version 71
 - list of files 171
 - local backup set using the GUI 193
 - local backupset using the GUI 77
 - maximum file size 54
 - most recent backup version 192
 - no query 72
 - overview 71
 - performing large operations 72
 - point-in-time 76
 - point-in-time, using command line 77
 - point-in-time, using GUI 77
 - primary tasks 71
 - process directories only (not files) 153
 - regulate resources server and client use during
 - processing 225
 - replace existing file with latest backup 182
 - restartable 72
 - running 74
 - searching files 74
 - sorting file list 44
 - specify end date and time to search for backups 261, 262
 - specify location of backup set for 194
 - specifying whether to include subdirectories for 247
 - specifying whether to prompt before overwriting existing
 - files 223
 - specifying whether to wait for tape to mount if it is
 - required 249
 - standard (also known as classic) 72
 - starting a Web client session 42
 - summary of options 117
 - to another workstation 81
 - to different workstation 268
 - restore backupset command 320
 - restore command 316
 - performing large operations 72
 - restoring
 - directories 74
 - files 74
 - files and directories 74
 - restrictions
 - asnodename option 129
 - within a proxied session 64, 89
 - retain extra versions attribute 102
 - retain only versions attribute 103
 - retain versions attribute 104
 - retention grace period
 - archive 100, 107
 - backup 100, 107
 - retrieve
 - archive copies 90
 - authorizing another user 80
 - based on date and time of archive 175, 178
 - files belonging to another node 81
 - files for another node 176
 - files only 174
 - hard links 89
 - list of files 171
 - maximum file size for 85
 - primary tasks 85
 - process directories only (not files) 153

- retrieve (*continued*)
 - replace existing file with latest archive if existing file is newer 182
 - sorting file list 44
 - specify end date and time to search for archives 261, 262
 - specifying description for 150
 - specifying whether to include subdirectories for 247
 - specifying whether to prompt before overwriting existing files 223
 - specifying whether to wait for tape to mount if it is required 249
 - starting a Web client session 42
 - summary of options 117
 - to another workstation 81
 - to different workstation 268
 - using client GUI 90
 - using command line 91
- retrieve command 323
- retryperiod option 228
- return codes for operations 96
- revokeremoteaccess option 229
- root user
 - tasks 13
- root user tasks
 - creating default client user options file 16
 - setting up 15

S

- schedcmddisabled option 230
- schedcmduser option (server defined only) 96
- schedlogmax option 231
- schedlogname option 232
- schedlogretention option 233
- schedmode option 235
- schedrestretrdisabled option 237
- schedule command 326
- schedule log
 - controlling the size 231
 - specifying number of days to keep entries and whether to save pruned entries 233
 - specifying path and file name to store schedule log information 232
- scheduled (automated) backups
 - closing files before back up 68
 - displaying scheduled work 93, 95
 - options for 93
 - overview 93, 335
 - process commands after back up 213
 - process commands before back up 215
 - restart applications after back up 68
 - starting 22
- scheduled commands
 - enabling-disabling 96
- scheduled events, displaying 311
- scheduled services
 - defining schedules for UID other than zero 96
 - disabling scheduled commands 230
- scheduler
 - configuring 21
 - displaying scheduled work 93, 95
 - event logging 95
 - options for 93
 - overview 93, 335
 - resolving memory retention after scheduled backups 196
 - specify whether server or client initiates sessions through firewall 242
- scheduler (*continued*)
 - specifying number of hours between contacts to server for scheduled work 220
 - specifying number of minutes between attempts to process scheduled commands 228
 - specifying polling mode or prompted mode 235
 - specifying whether to disable execution of restore or retrieve operations 237
 - starting 22
 - troubleshooting 342
- scheduler information
 - completed work 343
 - scheduled work 343
- scheduler log
 - log file 343
 - schedlogname option 343
 - schedlogretention option 343
 - understanding the 343
- Scheduler Service Configuration Utility 21
- scheduler wizard 21
- scheduling
 - client node proxy backup 65
- scrolllines option 238
- scrollprompt option 239
- searching
 - files for restore 74
- selective backup
 - overview 56, 58, 60
 - using client Java GUI 59
 - using command line 60
 - using the client command line 60
- selective command 328
- serialization
 - copy serialization
 - dynamic 103
 - shared static 103
 - static 103
- server
 - establishing communications through firewall 23
 - establishing communications with 15
 - identify to begin a stanza containing options for 240
 - specifying name of server to contact for services 240
 - specifying TCP/IP address of Tivoli Storage Manager server 258
 - specifying TCP/IP port address for 257
- server and node options
 - summary 112
- servername option 240
- service and technical support 45
- session information, displaying 312
- sessioninitiation option 242
- set access command 330
- set event command 332
- set password command 334
- setting
 - user privileges 43
- setting up
 - required root user tasks 15
- shared dynamic serialization 103, 138
- shared static serialization 103, 138
- shortcut keys
 - accessibility xiii
 - disability xiii
 - keyboard xiii
- skipacl option 244
- skipaclupdatecheck option 245
- soft mounts, NFS 67

- software requirements 7
- srvprepostscheddisabled option 246
- standard (classic) restore 72
- standard management class
 - copy destination 104
 - copy frequency 102
 - copy group name 102
 - copy mode
 - absolute 103
 - modified 103
 - copy serialization 103
 - copy type 102
 - default values 101
 - retain extra versions 102
 - retain only version 103
 - retain versions 104
 - versions data deleted
 - active versions 102
 - inactive versions 102
 - versions data exists 102
- standard policy domain 99
- starting
 - automatically 43
- starting a session
 - batch mode 41
 - interactive mode 41
- static serialization 103
- static, shared serialization 103
- storage
 - displaying restartable restore sessions 310
- storage management policies 99
 - assigning management classes to files 67
 - copy groups 99
 - default management class 99
 - display on backup-archive client or Web client GUI 67
 - include-exclude list 100
 - management classes 99
 - policy domains
 - default 99
 - standard 99
 - policy sets
 - active policy set 99
- subdir option 247
- subdirectories
 - archive 87
 - include in backup 60
- summary of changes for version 5.3.2 xvii
- support
 - gathering system information for 145, 173, 313
- supported file systems
 - HFS 53
 - HFS+ 53
 - HFSX 53
 - ISO 9600 53
 - UDF 53
 - UFS 53
- Swing-enabled browser
 - necessary to run Web client 42
- symbolic links and aliases
 - backing up 169, 188
 - exclude from backup 169, 188
 - exclude options 169, 188
 - processing 169, 188
- system administrator
 - task, installing Tivoli Storage Manager 1
- system files
 - excluding, recommended 30

- system information
 - gathering 145, 173, 313
- system objects
 - exclude from backup processing 169
- system state
 - display active and inactive objects 183

T

- T shell variables
 - setting 20
- tapeprompt option 249
- tasks
 - closed registration 25
 - open registration 25
- TCP/IP communication method
 - options 112
- tcpadminport option 251
- tcpbuffsize option 252
- tcpcadaddress option 253
- tcpclientaddress option 254
- tcpclientport option 255
- tcpnodelay option 256
- tcpserveraddress option 258
- tcpwindowsize option 259
- time format
 - specifying 260
- timeformat option 260
- Tivoli products and documentation
 - feedback survey xiii
 - ordering xiii
- Tivoli Storage Manager
 - client components 7
 - communication method 7
 - configuring 13
 - disk space requirements 7
 - environment prerequisites 6
 - hardware requirements 7
 - installation requirements 6
 - installing, system administrator task 1
 - memory requirements 7
 - migrating from earlier versions of the product 1
 - online forum 45
 - password 41
 - software requirements 7
- Tivoli Storage Manager client
 - authentication 39
- Tivoli Storage Manager password
 - using 40
- today option 261
- totime option 262
- transaction processing 263
 - summary of options 121
- txnbytelimit option 263
- type option 264

U

- UDF
 - file systems supported 53
- UFS
 - file systems supported 53
- Unicode 4
 - migrating file spaces to 50
 - renaming file spaces that are not Unicode to Unicode-enabled 135, 291, 328

- Unicode (*continued*)
 - restore from file spaces that are not Unicode-enabled 317
- Unicode migration 4
- UNIX
 - saving standard access permissions 89
- user privileges
 - setting 43
- users option 265

V

- v2archive option 266
- verbose option 267
- versions data
 - deleted attribute 102
 - exists attribute 102
- virtualnodename option 268
- volume label
 - duplicate names 49
- volume name 48
- volume naming precautions
 - dual boot systems 49
 - UNIX mount point 49

W

- Web client
 - enable to run in a Swing-enabled browser 42
 - establishing communications through firewall 23, 181
 - restrict administrator from accessing client running Web client 229
 - specifying TCP/IP port address for 181
 - starting 42
 - supported browsers 42
- webports option 269
- wildcard characters
 - guidelines 277
 - include or exclude groups of files 31
 - specifying a drive specification in dsm.opt 32
 - to include or exclude groups of files 32
 - using with commands 69
 - using with file specifications 69



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

Printed in USA

SC32-0143-01



Spine information:

IBM Tivoli Storage Manager for Macintosh **Version 5.5**

Backup-Archive Clients Installation and User's Guide

