



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 523.

First Edition (January 2007)

This edition applies to Version 5.4 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), IBM Tivoli Storage Manager for Space Management (5608-HSM), 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 xvi.

© 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 ix

About this book xi

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

Summary of changes for Tivoli Storage Manager xix

Technical changes for Version 5.4 - January 2007	xix
Technical changes for Version 5.3 - March 2006	xx
Technical changes for Version 5.3 - October 2005	xxi
Technical changes for Version 5.3 - December 2004	xxiii
Technical changes for Version 5.2.2 - December 2003	xxvii
Technical changes for Version 5.2 - April 2003	xxviii

Chapter 1. Installing Tivoli Storage Manager 1

Migrating from earlier versions	1
Upgrade path for clients and servers	1
Additional migration information	2
Client environment requirements	3
AIX client environment	3
HP-UX PA-RISC client environment	4
HP-UX Itanium 2 client environment	5
Linux x86/x86_64 client environment	6
Linux on POWER client environment	7
Linux on IBM eServer zSeries client environment	9
z/OS UNIX System Services client environment	10
Solaris client environment	10
NDMP support requirements (Extended Edition only)	12
Client installation media on product CD-ROMs	12
Installing the backup-archive client	13
Installing the AIX client	14
Uninstalling the AIX client	17
Installing the HP-UX PA-RISC client	18
Uninstalling the HP-UX PA-RISC client	21
Installing the HP-UX Itanium 2 client	22
Uninstalling the HP-UX Itanium 2 client	24
Installing the Linux x86/x86_64 client	25
Uninstalling the Linux x86/x86_64 client	27
Installing the Linux for zSeries client	28
Uninstalling the Linux for zSeries client	30
Installing the Linux on POWER clients	31
Uninstalling the Linux on POWER clients	33

Installing the Solaris clients	34
Uninstalling Tivoli Storage Manager on Sun Solaris clients	40

Chapter 2. Configuring Tivoli Storage Manager 41

Root and authorized user tasks	41
Creating and modifying the client system options file (required root user or authorized user task)	43
Creating a default client user options file (optional root user or authorized user task)	44
Setting options in an options file	45
Setting environment variables	46
Setting language environment variables	46
Setting processing environment variables	48
Setting Bourne and Korn shell variables	49
Setting C shell variables	50
Setting API environment variables	50
Configuring the Web client	50
Configuring the client scheduler	51
Starting the client scheduler	52
Configuring Tivoli Storage Manager client/server communication across a firewall	54
Configuring the journal daemon	56
Registering your workstation with a server (required root user or authorized user task)	57
Using closed registration	57
Using open registration	57
Associating your client node with a host system (optional)	58
Creating an include-exclude list (optional root user or authorized user task)	59
Using include-exclude options	60
Controlling symbolic link processing	64
Controlling compression and encryption processing	64
Testing an include-exclude list with the Preview command	64
Processing include and exclude options	65

Chapter 3. Getting started 69

Tivoli Storage Manager client authentication	69
Starting a GUI session	70
Using your Tivoli Storage Manager password	70
Setup Wizard	70
Starting a command line session	71
Using batch mode	71
Using interactive mode	71
Starting: Additional considerations	72
Starting a Web client session	72
Setting user privileges	74
Starting the client scheduler automatically	75
Changing your password	75
Sorting file lists using the Tivoli Storage Manager Java GUI	76

Displaying online help	76
Ending a session.	77
Online forums	77
Other sources of online help.	78
Chapter 4. Backing up your data	79
Planning your backups	79
When to back up and when to archive files.	80
Pre-backup considerations	81
LAN-free data movement.	81
Configure memory-constrained systems to run incremental backups	82
Configure systems with large numbers of files to run incremental backups	82
Using an include-exclude options list to control processing.	83
Encrypting data during backup or archive operations.	84
File system and ACL support	84
Maximum file size for operations	86
How Tivoli Storage Manager handles long user and group names	86
Performing an incremental, selective, or incremental-by-date backup	87
Full and partial incremental backup	87
Incremental-by-date backup	90
Comparing full incremental, partial incremental, and incremental-by-date incremental-by-date, and journal-based backups.	90
Selective backup.	91
Solaris global zone and local zones backup.	91
Saving access permissions	91
Setting a virtual mount point	92
Backing up data using the Java GUI	92
Backing up data using the command line	93
Deleting backup data	95
Deleting file spaces.	95
Group backup: Backing up files from one or more file spaces	96
Backup with client node proxy support	97
Scheduling with client node proxy support.	98
Associating a local snapshot with a server file space	101
Performing an image backup	102
Before you perform an image backup	102
Utilizing image backup to perform file system incremental backup	105
Performing an image backup using the GUI	106
Performing an image backup using the command line	108
Backing up NAS file systems	108
Backing up NAS file systems using the Web client GUI	109
Backing up NAS file systems using the command line	110
Backing up NFS file systems	111
Backing up the WebSphere Application Server (WAS)	112
Displaying backup processing status.	112
Backup: Additional considerations	114
Understanding how files are stored	114

How special file systems are handled	115
Understanding how files are managed	116
Understanding how symbolic links are handled	116
Understanding how hard links are handled	117
Understanding how sparse files are handled	118
Understanding how NFS hard and soft mounts are handled	118
Backing up opened files	119
Using wildcard characters	119

Chapter 5. Restoring your data.	121
Restoring an active or inactive backup	121
Restoring data using the GUI	122
Restoring data using the command line	122
Performing large restore operations	123
No query restore	124
Performing point-in-time restores.	126
Restoring an image	127
Performing an image restore using the GUI	128
Performing an image restore using the command line	129
Restoring an image to file	129
Performing an image restore using the GUI	129
Performing an image restore using the command line	130
Restoring data from a backupset	130
Restoring an entire or partial backupset	131
Restoring backupsets using the GUI.	132
Restoring backupsets using the command line client	133
Restoring data to a GPFS 3.1 file system with multiple storage pools	133
Restoring NAS file systems.	134
Restoring NAS file systems using the Web client	134
Restoring NAS file systems using the command line.	135
Restoring the WebSphere Application Server (WAS)	136
Restore: Additional considerations	138
Authorizing another user to restore or retrieve your files.	138
Restoring or retrieving files from another client node	139
Restore or retrieve files to another workstation	139
Restoring a disk in case of disk loss	140
Deleting file spaces	140

Chapter 6. Archiving and retrieving your data	143
Archiving files	143
Archiving data using the GUI.	143
Archiving data using the command line	144
Archiving data with client node proxy	145
Deleting archive data.	147
Archive: Advanced considerations	147
Retrieving archives	148
Retrieving data using the GUI.	149
Retrieving data using the command line	149
Understanding how your archives are managed	150

Chapter 7. Automating tasks	151
----------------------------------------------	------------

Specifying scheduling options	151
Displaying information about scheduled work	151
Displaying information about completed work	153
Scheduling options for commands	154
Enabling or disabling scheduled commands	154
Client return codes	155

Chapter 8. Understanding storage management policies 157

Using policy domains and policy sets	157
Using management classes and copy groups	158
Displaying information about management classes and copy groups	159
Copy group name	160
Copy type	160
Copy frequency	160
Versions data exists	160
Versions data deleted	160
Retain extra versions	161
Retain only version	161
Copy serialization	161
Copy mode	162
Copy destination	162
Retain versions	162
Selecting a management class for files	162
Assigning a management class to files	163
Overriding the management class for archived files	164
Selecting a management class for directories	164
Binding management classes to files	165
Rebinding backup versions of files	165
Using a retention grace period	166
Using event-based policy retention protection	166
Archiving files on a data retention server	167

Chapter 9. Using processing options 169

Overview of processing options	169
Communication options	169
TCP/IP options	170
Shared Memory options	170
Server and Node options	171
Server options	171
Node options	172
Backup and archive processing options	173
Restore and retrieve processing options	179
Scheduling options	182
Format options	184
Command processing options	184
Authorization options	184
Error processing options	185
Transaction processing options	186
Web client options	187
Diagnostics options	187
Using options with commands	187
Entering options with a command	188
Initial command line only options	193
Client options reference	194
Archmc	195
Archsymlinkasfile	196
Asnodename	197
Automount	199

Backupsetname	200
Changingretries	201
Class	202
Clusternode	203
Collocatebyfilespec	204
Commmethod	206
Commrestartduration	208
Commrestartinterval	209
Compressalways	210
Compression	211
Console	213
Dateformat	214
Defaultserver	217
Deletefiles	218
Description	219
Detail	221
Dirmc	222
Dironly	223
Diskbuffsize	224
Diskcachelocation	225
Domain	227
Domain.image	231
Domain.nas	232
Editor	234
Enablearchiveretentionprotection	235
Enablelanfree	237
Encryptiontype	239
Encryptkey	240
Errorlogmax	242
Errorlogname	243
Errorlogretention	244
Exclude options	246
Filelist	250
Filename	252
Filesonly	253
Followsymbolic	254
Fromdate	255
Fromnode	256
Fromowner	257
Fromtime	258
Groupname	259
Groups	260
Guitreeviewafterbackup	261
Httpport	262
Ifnewer	263
Imagetofile	264
Imagetype	265
Inactive	267
Inclxcl	268
Include options	269
Processing NAS file systems	273
Processing WebSphere Application Server (WAS) file systems	273
Incrbydate	274
Incremental	275
Lanfreecommmethod	276
Lanfreeshmport	278
Lanfreetcport	279
Lanfreetcserveraddress	280
Latest	281
Localbackupset	282

Location	283	Tcpserveraddress	367
Mailprog	284	Tcpwindow size	368
Makesparsefile	285	Timeformat	369
Manageservices	286	Toc	371
Maxcmdretries	288	Todate	373
Memoryefficientbackup	289	Totime	374
Mode	291	Txnbytelimit	375
Monitor	293	Type	376
Nasnodename	294	Users	377
Nfstimeout	295	V2archive	378
Nodename	297	Verbose	379
Nojournal	299	Verifyimage	380
Noprompt	300	Virtualfsname	381
Numberformat	301	Virtualmountpoint	382
Optfile	303	Virtualnodename	384
Password	304	Wasexphome	386
Passwordaccess	305	Washome	387
Passworddir	307	Wasndhome	388
Pick	308	Wasnode	389
Pitdate	309	Wastype	390
Pittime	310	Wasuser	391
Postschedulecmd/Postnschedulecmd	311	Webports	392
Postsnapshotcmd	313		
Preschedulecmd/Prenschedulecmd	314	Chapter 10. Using commands 393	
Preservelastaccessdate	316	Starting and ending a client command session	395
Preservepath	318	Process commands in batch mode	396
Presnapshotcmd	321	Process commands in interactive mode	396
Querschedperiod	322	Entering client commands	397
Quiet	323	Command name	397
Removeoperandlimit	324	Options	397
Replace	325	Parameters	398
Resourceutilization	327	File specification syntax	398
Retryperiod	329	Remembering previous commands	399
Revokeremoteaccess	330	Using wildcard characters	400
Schedcmddisabled	331	Entering commands	401
Schedlogmax	332	Client commands reference	401
Schedlogname	333	Archive	402
Schedlogretention	334	Backup Group	404
Schedmode	336	Backup Image	406
Schedrestretrdisabled	338	Backup NAS	410
Scrolllines	339	Backup WAS	413
Scrollprompt	340	Cancel Process	415
Servername	342	Cancel Restore	416
Sessioninitiation	344	Delete Access	417
Shmport	346	Delete Archive	418
Showmembers	347	Delete Backup	420
Skipacl	348	Delete Filespace	423
Skipaclupdatecheck	349	Delete Group	425
Snapshotcachesize	350	Expire	427
Snapshotroot	351	Help	429
Srvprepostscheddisabled	354	Incremental	430
Srvprepostsnapdisabled	355	Loop	435
Subdir	356	Macro	437
Tapeprompt	358	Monitor Process	438
Tcpadminport	360	Preview	439
Tcpbuffsize	361	Query Access	441
Tpcadaddress	362	Query Archive	442
Tcpclientaddress	363	Query Backup	444
Tcpclientport	364	Query Backupset	447
Tcpnodelay	365	Query Filespace	449
Tcpport	366	Query Group	451

Query Image	453
Query Includ	455
Query Mgmtclass	456
Query Node	457
Query Options	458
Query Restore	459
Query Schedule	460
Query Session	461
Query Systeminfo	462
Query WAS	464
Restart Restore	465
Restore	466
Restore Backupset	470
Restore Group	474
Restore Image	476
Restore NAS	479
Restore WAS	481
Retrieve	484
Schedule	486
Selective	488
Set Access	491
Set Event	493
Set Password	495
Set Waspassword	496

Appendix A. Using the Tivoli Storage Manager central scheduler 499

Overview of the Tivoli Storage Manager scheduler	499
Handling spaces in file names in schedule definitions	500
Preferential start times for certain nodes	500
Understanding scheduler processing options	501
Handling return codes from preschedulecmd and postschedulecmd scripts	502
Using the client acceptor to manage scheduler services versus the traditional scheduler services	503
Setting the client scheduler process to run as a background task and start automatically at boot time	503
Managing multiple schedule requirements on one machine	505
Restarting the scheduler process on a remote machine	507

Using the scheduler on clustered systems	507
Troubleshooting the Tivoli Storage Manager scheduler	508
Investigating abnormal scheduled events	509
Checking the server activity log	509
Inspecting the Client Schedule Log	510
Understanding the scheduler log	510

Appendix B. Configuring the backup-archive client in an HACMP takeover environment 511

Installing the backup-archive client	512
Configuring the backup-archive client to process local nodes	512
Configuring Tivoli Storage Manager backup-archive client to process cluster disk resources	512
Step 1: Register the client to a server	513
Step 2: Configure the client system options file	513
Step 3: Configure the client user options file	514
Defining the client as an HACMP application	514
Creating an HACMP resource group to add a client	514
Adding the client to an HACMP resource group	515
Using journal-based backup and HACMP	515

Appendix C. Journal daemon configuration 517

JournalSettings stanza	517
JournalExcludeList stanza	518
JournalFileSystemSettings stanza	519
Overriding stanzas	521

Notices 523

Trademarks	525
----------------------	-----

Glossary 527

Index 537

Tables

1.	Typographical conventions	xii	38.	Symbolic link management table for backup and restore	117
2.	IBM Tivoli Storage Manager client and server publications	xv	39.	Restore: Primary tasks	121
3.	Optimized option default values	xxvi	40.	Command line restore examples	123
4.	Client requirements	3	41.	NAS options and commands	136
5.	UNIX client installation reference	13	42.	Archive and retrieve: Primary tasks	143
6.	AIX client: Language codes for installation packages	15	43.	Command line archive examples	144
7.	Package Names and Directories	18	44.	Symbolic link management table for archive and retrieve	148
8.	HP-UX PA-RISC client: Language codes for installation packages	19	45.	Command line examples of retrieving archives	149
9.	HP-UX Itanium 2 client: Language codes for installation packages	22	46.	Client return codes and their meanings	155
10.	Package Names and Directories	25	47.	Default values in the standard management class	159
11.	Linux x86/x86_64 client: Language codes for installation packages	26	48.	TCP/IP options	170
12.	Package Names and Directories	28	49.	Shared Memory communication options	170
13.	Linux for zSeries client: Language codes for installation packages	29	50.	Server and Node Options	172
14.	Linux on POWER clients: Language codes for installation packages	32	51.	Backup and archive processing options	173
15.	Installation package names and descriptions	34	52.	Restore and retrieve processing options	179
16.	Tivoli Storage Manager server Installation Guides	40	53.	Scheduling options.	182
17.	Root user, Authorized User, and non-Authorized user tasks	41	54.	Format options	184
18.	Supported language locales	47	55.	Command processing options	184
19.	Platforms and default log directories	49	56.	Authorization options.	184
20.	GUID commands	58	57.	Error processing options	185
21.	Options for excluding file spaces and directories	60	58.	Transaction processing options.	186
22.	Options for controlling backup, archive, and image processing.	61	59.	Web client options	187
23.	Wildcard and other special characters	62	60.	Diagnostics options	187
24.	Using wildcard characters with include and exclude patterns	63	61.	Client command options.	188
25.	Options for controlling symbolic link processing	64	62.	Options that are valid on the initial command line only	194
26.	Options for controlling compression and encryption processing	64	63.	Encrypting or decrypting data	240
27.	Working with your files using the Tivoli Storage Manager GUI	76	64.	Entering commands	393
28.	Backup: Primary tasks	79	65.	Commands	393
29.	Planning your backups.	79	66.	Command recall and edit functions	400
30.	Supported file systems and ACL support	84	67.	Wildcard characters	400
31.	Maximum file size	86	68.	Archive command: Related options	402
32.	Command line backup examples	93	69.	Backup Group command: Related options	405
33.	Volume device type support for an image backup.	103	70.	Backup Image command: Related options	409
34.	Comparing incremental image backup methods	106	71.	Backup NAS command: Related options	411
35.	LVM1 and LVM2 image operation comparisons	107	72.	Backup WAS command: Related options	413
36.	NAS options and commands	110	73.	Delete Archive command: Related options	418
37.	Client command line informational messages	112	74.	Delete Backup command: Related options	420
			75.	Delete Filespace command: Related options	424
			76.	Delete Group command: Related options	425
			77.	Expire command: Related options.	427
			78.	Incremental command: Related options	433
			79.	Query Archive command: Related options	442
			80.	Query Backup command: Related options	444
			81.	Query Backupset command: Related options	447
			82.	Query Filespace command: Related options	449
			83.	Query Group command: Related options	451
			84.	Query Image command: Related options	453
			85.	Query Mgmtclass command: Related options	456
			86.	Query Node command: Related options	457
			87.	Query Options command: Related options	458
			88.	Query Systeminfo command: Related options	463

89. Query WAS command: Related options	464	95. Restore WAS command: Related options	482
90. Restore command: Related options	467	96. Retrieve command: Related options	484
91. Restore Backupset command: Related options	471	97. Schedule command: Related options	486
92. Restore Group command: Related options	474	98. Selective command: Related options	489
93. Restore Image command: Related options	477	99. Scheduled client event status codes	508
94. Restore NAS command: Related options	480		

About this book

IBM Tivoli Storage Manager is a client-server licensed product that provides storage management services in a multi-platform 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 either as a backup and archive server or migration server for distributed workstations and file servers. The server program also supplies hierarchical storage management (HSM) services. See “Related publications” on page xv 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 xv for available Tivoli Storage Manager Administrator’s Reference publications.
- An *application program interface (API)* that permits you to enhance an existing application with storage management services. When an application is registered with a server as a client node, the application can back up, restore, archive, and retrieve objects from storage. For more information about the Tivoli Storage Manager API, see *IBM Tivoli Storage Manager Using the Application Programming Interface*, SC32-0147.
- A *Web backup-archive client* that permits an authorized administrator, help desk person, or end user to perform backup, restore, archive, and retrieve services using a Web browser on a remote machine. See “Starting a Web client session” on page 72 for more information.

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

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

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 41.

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**. See “Root and authorized user tasks” on page 41 for more information about these tasks. 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*
```

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.

The Tivoli Storage Manager home page **Self help** section provides customers with a knowledge base of articles and information on issues with Tivoli Storage Manager that you might be experiencing.

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 \blacktriangleright — symbol indicates the beginning of a syntax diagram.
- The — \blacktriangleright symbol at the end of a line indicates the syntax diagram continues on the next line.
- The \blacktriangleright — symbol at the beginning of a line indicates a syntax diagram continues from the previous line.
- The — \blacktriangleleft 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.	\blacktriangleright —KEYWOrd— \blacktriangleleft																		
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.	<table> <tr><td>*</td><td>Asterisk</td></tr> <tr><td>{ }</td><td>Braces</td></tr> <tr><td>:</td><td>Colon</td></tr> <tr><td>,</td><td>Comma</td></tr> <tr><td>=</td><td>Equal Sign</td></tr> <tr><td>-</td><td>Hyphen</td></tr> <tr><td>()</td><td>Parentheses</td></tr> <tr><td>.</td><td>Period</td></tr> <tr><td></td><td>Space</td></tr> </table>	*	Asterisk	{ }	Braces	:	Colon	,	Comma	=	Equal Sign	-	Hyphen	()	Parentheses	.	Period		Space
*	Asterisk																		
{ }	Braces																		
:	Colon																		
,	Comma																		
=	Equal Sign																		
-	Hyphen																		
()	Parentheses																		
.	Period																		
	Space																		
Variables:																			
Italicized lowercase items (<i>var_name</i>) denote variables.	\blacktriangleright —KEYWOrd— <i>var_name</i> — \blacktriangleleft																		
In this example, you can specify a <i>var_name</i> when you enter the KEYWORD command.																			

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

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:

- 1 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.

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 diagram description	Example
<p>Syntax fragments:</p> <p>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.</p>	<p>►► The fragment name ◄◄</p> <p>The fragment name:</p>

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 Macintosh Backup-Archive Clients Installation and User's Guide</i>	SC32-0143
<i>IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide</i>	SC32-0148
<i>IBM Tivoli Storage Manager Using the Application Programming Interface</i>	SC32-0147
Server publications	
<i>IBM Tivoli Storage Manager for AIX Installation Guide</i>	SC32-0134
<i>IBM Tivoli Storage Manager for AIX Administrator's Reference</i>	SC32-0123
<i>IBM Tivoli Storage Manager for AIX Administrator's Guide</i>	SC32-0117
<i>IBM Tivoli Storage Manager for AIX Storage Agent User's Guide</i>	SC32-0129
<i>IBM Tivoli Storage Manager for HP-UX Installation Guide</i>	SC32-0135
<i>IBM Tivoli Storage Manager for HP-UX Administrator's Reference</i>	SC32-0124
<i>IBM Tivoli Storage Manager for HP-UX Storage Agent User's Guide</i>	SC32-0130
<i>IBM Tivoli Storage Manager for Linux Installation Guide</i>	SC32-0136
<i>IBM Tivoli Storage Manager for Linux Administrator's Reference</i>	SC32-0125
<i>IBM Tivoli Storage Manager for Linux Storage Agent User's Guide</i>	SC32-0131
<i>IBM Tivoli Storage Manager for z/OS Installation Guide</i>	SC32-0139
<i>IBM Tivoli Storage Manager for z/OS Administrator's Reference</i>	SC32-0128
<i>IBM Tivoli Storage Manager for Sun Solaris Installation Guide</i>	SC32-0137
<i>IBM Tivoli Storage Manager for Sun Solaris Administrator's Reference</i>	SC32-0126
<i>IBM Tivoli Storage Manager for Sun Solaris Administrator's Guide</i>	SC32-0120
<i>IBM Tivoli Storage Manager for Sun Solaris Storage Agent User's Guide</i>	SC32-0132
<i>IBM Tivoli Storage Manager for Windows Installation Guide</i>	SC32-0138
<i>IBM Tivoli Storage Manager for Windows Administrator's Guide</i>	SC32-0121
<i>IBM Tivoli Storage Manager for Windows Storage Agent User's Guide</i>	SC32-0133

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, please 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:

- Through the IBM Support Assistant

An IBM Support Assistant plugin is available for the Tivoli Storage Manager client to help you to do the following:

- Locate specific self-help information on the Web based on frequently asked questions and common problems
- Search for other Tivoli Storage Manager information on the Web
- Report a problem
- Automatically collect basic operating system and Tivoli Storage Manager related environmental and diagnostic information

You can find more information about the IBM Support Assistant at <http://www.ibm.com/software/support/isa>. You can also download the IBM Support Assistant from this page.

To download the IBM Support Assistant, you will need to login using your IBM ID (the same ID that you use for IBM web sites such as the MySupport Portal and developerWorks). If you do not already have an IBM ID, you can complete the free registration process to obtain one.

After downloading the IBM Support Assistant, install it on your computer.

Next, install the Tivoli Storage Manager client plugin. To do this, open the IBM Support Assistant, select the **Updater** tab and then the **New Products and Tools** tab. Locate **Tivoli Storage Manager Backup-Archive Client** and select the check box next to it. Click **Install** to install the plugin.

After the install finishes, restart the IBM Support Assistant. Then, click Product Information. The Tivoli Storage Manager client should appear in the Products list. To get assistance for the Tivoli Storage Manager client, explore the available items in the IBM Support Assistant.

- 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 the IBMLINK 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.

Go to the following Web site for the Tivoli Storage Manager home page:
<http://www.ibm.com/software/tivoli/products/storage-mgr/>

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

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

by selecting the Feedback link in the left navigation bar.

Reporting a problem

Please 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 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 77 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 5.4 - 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 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 client. See "Upgrade path for clients and servers" on page 1 for more information about the changes.

AIX journal support for AIX clients

AIX journal now supports JFS, JFS2, and VxFS and HSM-managed JFS2 file systems and the HACMP environment.

Linux LVM 2 volume manager image backup support

Image backup of volumes created using the Linux LVM 2 volume manager is now supported. This provides the same level of functionality as offered for LVM 1 logical volumes. Only specific Linux distributions (Kernel 2.6) are supported.

SLES 10 support on Linux clients

SuSE Linux Enterprise Server 10 is now supported on the Linux x86/x86_64, Linux zSeries, and Linux PPC clients.

USS V1R7 and V1R8 support

UNIX System Services V1R7 and V1R8 are now supported.

HACMP 5.4 support

HACMP 5.4 is now supported with the AIX Backup-Archive client and API, if the fix for IY89869 is applied first.

Asianux 2.0 support on Linux x86/x86_64 clients

Asianux 2.0 now supports the Linux x86/x86_64 clients.

Support for *srprepuestoscheddisabled*, *srprepuestosnapdisabled*, and *schedrestretrdisabled* options

Three new options: *srprepuestoscheddisabled*, *srprepuestosnapdisabled*, and *schedrestretrdisabled*, 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 "Srvprepuestoscheddisabled" on page 354, "Srvprepuestosnapdisabled" on page 355, and "Schedrestretrdisabled" on page 338 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 289 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 225 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 269 for more information about using the *include.fs* option.

Tivoli Storage Manager client support for Novell Storage Services file system extended attributes

The Tivoli Storage Manager Linux x86/x86_64 client can now back up and restore Novell Storage Services file system extended attributes. Novell Storage Services file system Trustee rights are packaged as part of the extended attributes, and they are also backed up and restored. The extended attributes information is automatically backed up with each object that has extended attribute data.

Be aware that the first incremental backup using this client will back up all Novell Storage Services files again, whether or not they have extended attributes.

Storage pools on GPFS file systems are supported

Tivoli Storage Manager now handles storage pools on GPFS file systems. See “Restoring data to a GPFS 3.1 file system with multiple storage pools” on page 133 for more information.

Local zone support for Solaris 10 x86 client

Local zone support has been added for the Solaris 10 x86 client. See “Solaris global zone and local zones backup” on page 91 for more information.

Dropped support for HP-UX 11i V1

Beginning with this release of Tivoli Storage Manager, we have dropped support for HP-UX 11i V1.

Dropped support for RHEL 3

Beginning with this release of Tivoli Storage Manager, we have dropped support for Red Hat Enterprise Linux 3.

Dropped support for SLES 8

Beginning with this release of Tivoli Storage Manager, we have dropped support for SuSE Linux Enterprise Server 8.

Dropped support for AIX Version 5.1

Beginning with this release of Tivoli Storage Manager, we have dropped support for AIX Version 5.1.

Dropped support for Solaris 8

Beginning with this release of Tivoli Storage Manager, we have dropped support for Solaris 8.

Dropped support USS V1R4, V1R5, and V1R6

Beginning with this release of Tivoli Storage Manager, we have dropped support for UNIX System Services V1R4, V1R5, and V1R6.

Technical changes for Version 5.3 - March 2006

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

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 348 and “*Skipaclupdatecheck*” on page 349 for more information.

Journal-based backup enhancements for AIX

Journal-based backup is supported for non-HSM AIX clients. See “*Configuring the journal daemon*” on page 56, “*Journal-based backup on AIX*” on page 89, and Chapter 10, “*Using commands,*” on page 393 for more information.

Zones support for Solaris 10 platforms

Zones support has been added for the following Solaris 10 platforms:

- Solaris 10 x86 32-bit and 64-bit limited support (Global Zone only, API 32-bit only, UFS only, TCP-IP only, no LAN-Free)
- Solaris 10 SPARC Local Zones
- Solaris 10 SPARC support for Space Management (HSM) (Global Zone only)
- Solaris HSM requires Veritas File system (VxFS) 4.1

The standard backup and archive operations can be performed from each Local Zone, just as these operations can be performed on any Solaris system without zoning.

HP-UX 11i V2 support

HP-UX 11i V2 is supported on the Backup-Archive clients.

Java GUI command line parameters

Tivoli Storage Manager Backup-Archive Java GUI (dsmj) now accepts command line parameters, such as the Java -X options. Because of this, you can also now modify the Java Heap Size.

Technical changes for Version 5.3 - October 2005

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

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 166, “*Enablearchiveretentionprotection*” on page 235 and “*Set Event*” on page 493 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.

See Chapter 1, “Installing Tivoli Storage Manager,” on page 1, Chapter 4, “Backing up your data,” on page 79, and Chapter 10, “Using commands,” on page 393 for more information.

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 Java GUI directory tree shows detailed information of included and excluded objects.

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

See Chapter 2, “Configuring Tivoli Storage Manager,” on page 41, Chapter 4, “Backing up your data,” on page 79, Chapter 6, “Archiving and retrieving your data,” on page 143, and Chapter 10, “Using commands,” on page 393 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 151 and Chapter 10, “Using commands,” on page 393 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 IBM GPFS. 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 197 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 95 for more information.

Backing up SAN File System data

You can now back up and restore SAN file system data belonging to different operating systems using the same client. You can use a single AIX or Linux machine to back up all SAN file system data. The object attributes will be correctly restored, even if backup and restore of Windows-owned files was performed by a UNIX machine. See “Full and partial incremental backup” on page 87 for more information.

NFSV4 support

Support has been added for NFSV4 file systems on AIX Version 5.3 and above. Support has also been added for NFSV4 global name space. NFSV4 ACLs are not supported. See “Backing up NFS file systems” on page 111 for more information.

Toleration support for Solaris 10

Toleration support has been added for Solaris 10 on the Backup-Archive clients and the API (no local zones).

HP-UX 11i V2 support

HP-UX 11i V2 is supported on the Backup-Archive clients.

Technical changes for Version 5.3 - December 2004

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

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

From the backup-archive client Java GUI and web client, 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 169 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 command line client, the web client acceptor and agent will not run without a writeable dsmerror.log. See Chapter 2, “Configuring Tivoli Storage Manager,” on page 41, Chapter 9, “Using processing options,” on page 169, and Chapter 10, “Using commands,” on page 393 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.

See Chapter 1, “Installing Tivoli Storage Manager,” on page 1, Chapter 4, “Backing up your data,” on page 79, and Chapter 10, “Using commands,” on page 393 for more information.

Shared memory protocol support

Administrators can now perform backups, restores, archives, and retrieves more quickly by improving communications between the Tivoli Storage Manager server for Linux and Solaris and the backup-archive client for Linux and Solaris. The server and the client must be configured for the shared memory protocol, and they must be on the same server machine. Shared memory is an area of memory that is accessible by several processes at the same time.

This support also enables the shared memory communication protocol in the Linux and Solaris storage agent for communication with the UNIX and Linux backup-archive client.

See Chapter 10, “Using commands,” on page 393 for more information.

Journal-based backup enhancements

Multiple journal-based backup sessions are now possible. See Chapter 4, “Backing up your data,” on page 79 and Chapter 10, “Using commands,” on page 393 for more information.

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 Java GUI directory tree shows detailed information of included and excluded objects.

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

See Chapter 2, “Configuring Tivoli Storage Manager,” on page 41, Chapter 4, “Backing up your data,” on page 79, Chapter 6, “Archiving and retrieving your data,” on page 143, and Chapter 10, “Using commands,” on page 393 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 151 and Chapter 10, “Using commands,” on page 393 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 IBM GPFS. 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 197 for more information.

Tivoli Storage Manager Administration Center

The hyperlink from the new Java based Tivoli Storage Manager Administration Center to a Tivoli Storage Manager client machine provides the administrator ID and encrypted password to the web client through a Java portlet. This enables the launching of the web client Java GUI without the administrator signing on again. The Tivoli Storage Manager Version 5.3 or higher web client is required for this support. See “Configuring the Web client” on page 50 and the *IBM Tivoli Storage Manager Administrator's Guide* for your operating system, for more information.

Java GUI and Web client enhancements

You can now perform the following functions from the Java GUI and the Web client:

- Find files in the backup, restore, archive or retrieve window.
- Back up, restore, archive, or retrieve your files by filtering file names or filtering the directory tree.
- Restore your data from backup sets without a server connection.
- Archive the file or directory where the symbolic link points.

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 95 for more information.

Tivoli Storage Manager backup-archive client for HP-UX Itanium 2

The Tivoli Storage Manager backup-archive client supported features on HP-UX Itanium 2 are the same as those supported by the existing HP-UX client. See Chapter 1, “Installing Tivoli Storage Manager,” on page 1 for environment and installation information.

Linux for zSeries® offline image backup

You can now perform offline image backups on Linux zSeries, including native zSeries volumes and partitions created by Logical Volume Manager. See “Volume device type support for an image backup” on page 103 for more information.

Optimized option default values

The *tcpnodelay* option specifies whether to send small transactions to the

server, without buffering them first. Setting the *tcpnodelay* option to *yes* (the default) can improve performance in higher-speed networks. This option has been extended to all UNIX and Linux clients, except UNIX System Services client for z/OS®. Table 3 lists the default values which have been optimized to improve performance. See “*Tcpnodelay*” on page 365 for more information.

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.	224
<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.	224
<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.	360
<i>tcpbuffsize</i>	The default value was changed from 31 kilobytes to 32 kilobytes.	361

Table 3. Optimized option default values (continued)

Option	Change	Page
<i>tcpnodelay</i>	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. 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.	365
<i>tcpwindowsize</i>	The default value was changed from 32 kilobytes to 63 kilobytes.	368
<i>txnbytelimit</i>	The default value was changed from 2048 kilobytes to 25600 kilobytes.	375

Technical changes for Version 5.2.2 - December 2003

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

Tivoli Storage Manager backup-archive client Linux on iSeries

The Tivoli Storage Manager backup-archive client supported features on Linux for iSeries are the same as those supported by the Linux for pSeries client. See Chapter 1, "Installing Tivoli Storage Manager," on page 1 for environment and installation information.

Tivoli Storage Manager backup-archive client for Linux on Intel™ Itanium

The Tivoli Storage Manager backup-archive client for Linux on Intel Itanium (Linux IA64 client) supports these file systems and their ACLs: xfs, ext2, and ext3. As with the Linux x86/x86_64 client, the Linux IA64 client supports the reiserfs file system, but not its ACLs.

The Tivoli Storage Manager Linux IA64 client supports the same functions as the Linux x86/x86_64 client, except the following:

- Backup-archive Native Graphical User Interface (Motif GUI).
- LAN-free data transfer.
- Cluster support.

See Chapter 1, "Installing Tivoli Storage Manager," on page 1 for environment and installation information.

Backup and restore support for IBM TotalStorage® SAN File Systems

Backup and restore for IBM TotalStorage SAN File Systems is supported on the AIX 5.1 client. See “File system and ACL support” on page 84 for more information.

Support for controlling symbolic link processing

Tivoli Storage Manager treats symbolic links as actual files and backs them up. However, the file referenced by the symbolic link is not backed up. In some cases symbolic links can be easily recreated and need not be backed up. In addition, backing up these symbolic links can increase backup processing time and occupy a substantial amount of space on the Tivoli Storage Manager server. You can use the *exclude.attribute.symlink* option to exclude a file or a group of files that are symbolic links from backup processing. If necessary, you can use the *include.attribute.symlink* option to include symbolic links within a broad group of excluded files for backup processing. For additional information, see :

- “Using include-exclude options” on page 60
- “Exclude options” on page 246
- “Include options” on page 269

Support for WebSphere® Application Server (WAS) security

If WAS security is enabled, user name and password validation for Data Protection for WebSphere Application Server is required. If you do not set the WAS password for the security, the backup will failover to an offline backup. It is recommended to set the WAS security password to perform consistent backups. Use the **set waspassword** command to set the user name and password for each installation of WAS on your machine. You only need to perform this task once, unless you change your WAS user name or password. You can only perform this task on the Tivoli Storage Manager command line. See “Set Waspassword” on page 496 for more information.

Removal of operand limits for backup and archive operations

Use the *removeoperandlimit* option to specify that Tivoli Storage Manager removes the 20-operand limit for UNIX and Linux-family platforms. 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 324

Multi-session backup session enhancements

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* 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 204 for more information.

Language support expanded to include Russian, Hungarian, Polish, and Czech

Tivoli Storage Manager client language packs are available in Russian, Hungarian, Polish, and Czech language locales. See “Setting language environment variables” on page 46 for more information.

Technical changes for Version 5.2 - April 2003

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

Support for backing up files from one or more file space origins

You can use the Tivoli Storage Manager command line client to create and back up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server. A *group backup* allows you to create a point-in-time backup of a group of files that is managed as a single logical entity. See “Group backup: Backing up files from one or more file spaces” on page 96 for more information.

Support for processing EMC Celerra NAS file system images

Through support of NDMP, Tivoli Storage Manager Windows, AIX, and Solaris servers can efficiently back up and restore NAS file system images to tape drives or libraries that are locally attached to Network Appliance and EMC Celerra NAS file servers. *NDMP support is available only on IBM Tivoli Storage Manager Extended Edition.* See “NDMP support requirements (Extended Edition only)” on page 12 for NDMP support requirements. See “Backing up NAS file systems” on page 108 for information on how to back up and restore NAS file system images using the Web client and command line client.

Support for backup and restore of the WebSphere Application Server

If you installed the Data Protection for WebSphere Application Server, you can use Tivoli Storage Manager to back up the Version 5.0 WebSphere Application Server Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server to the Tivoli Storage Manager server. You can restore this information from the Tivoli Storage Manager server and use it to recover a corrupted node application or an entire node (or nodes) in the event of an accident or disaster.

Data Protection for WebSphere Application Server is a separately priced and licensed product. Data Protection for WebSphere Application Server is only supported on a Tivoli Storage Manager Version 5.2 server and client. Data Protection for WebSphere Application Server is supported on the AIX 5.1 with Maintenance level 2; Solaris 8; Red Hat Advanced Server 2.1 for x86, SUSE LINUX 7.3 for x86, SUSE LINUX Enterprise Server 7 for x86, and United Linux 1.0 for x86 clients only.

See “Backing up the WebSphere Application Server (WAS)” on page 112 for more information.

Enhancements for command line image restore operations

You can use the *verifyimage* option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log. See “Verifyimage” on page 380 for more information.

If bad sectors are present on the target volume, you can use the *imagetofile* option with the **restore image** command to specify that you want to restore the source image to a file. Later, you can use a ‘dd’ utility (available on Unix) or its equivalent to copy data from this file to a logical volume. See “Imagetofile” on page 264 for more information.

Support for an external snapshot provider in the backup-archive clients

Use the *snapshotroot* option with the **incremental**, **selective**, or **archive** commands in conjunction with a third-party application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. The *snapshotroot* option does not provide any facilities to take a

volume snapshot, only to manage data created by a volume snapshot. See “Snapshotroot” on page 351 for more information.

Enhancements to the Web client interface

The following functions are now supported in the Tivoli Storage Manager Web client interface:

Access Another Node

Displays the backup versions and archive copies of another node. You can then restore the backup versions or retrieve the archives from the other user to your workstation. *You must have authorization to access the stored data of another node.*

Node Access List

Allows you to authorize other users to access your backup versions and archive copies.

View Policy Information

Displays storage management policy information for your node.

Enhanced firewall security

Security for back up and restore operations and Tivoli Storage Manager administrative functions is enhanced to allow the Tivoli Storage Manager server to control whether the server or client initiates sessions through a firewall. See “Sessioninitiation” on page 344 for more information.

Both the server and client can also specify 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 360 for more information.

Veritas file systems, ACLs, and Veritas Volume Manager support on AIX (32-bit and 64-bit) clients

Tivoli Storage Manager supports backup, restore, archive, and retrieve of Veritas file systems (VxFS) including ACLs on the AIX (32-bit and 64-bit) clients. Image backup and restore of Veritas Volume Manager logical volumes is also supported.

Automounter support for Linux x86/x86_64 and Linux zSeries clients

The Tivoli Storage Manager Linux x86/x86_64 and Linux zSeries clients now support the back up of automounted NFS and loopback file systems. Use the *automount* option with the *domain* option to specify all automounted file systems the Tivoli Storage Manager client tries to mount at the following points in time:

- When Tivoli Storage Manager client starts
- When the back up is started
- When the Tivoli Storage Manager client has reached an automounted file system during backup

See “Automount” on page 199 and “Domain” on page 227 for more information.

Support for displaying options and their settings via the command line

Use the **query options** command to display all or part of your options and their current settings. This command accepts an argument to specify a subset of options. The default is to display all options. See “Query Options” on page 458 for more information.

Support for gathering Tivoli Storage Manager system information

Use the **query systeminfo** command to gather information on one or more of the following items and output this information to a file name that you specify:

- DSMOPTFILE - The contents of dsm.opt file.

- ENV - Environment variables.
- ERRORLOG - The Tivoli Storage Manager error log file.
- FILE - Attributes for the file name that you specify.
- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- OPTIONS - Compiled options.
- OSINFO - Name and version of the client operating system (includes ULIMIT information for UNIX and Linux).
- POLICY - Policy set dump.
- DSMSYSFILE - The contents of the dsm.sys file.
- CLUSTER - AIX cluster information.

See “Query Systeminfo” on page 462 for more information.

Enhancements for the query filespace command

The **query filespace** command is enhanced to allow you to query a single file space on the Tivoli Storage Manager server. See “Query Filespace” on page 449 for more information.

Separately installable language packs available

Language packs are separately installable packages that contain only language-specific files (such as message catalog, resource file, help files, etc.). You can now install these additional language packs on top of your Tivoli Storage Manager client base install. To change your language preferences, specify the LANG environment variable or use the Preferences editor.

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 or server files that you do not currently need, preserve them in their current state, and retrieve them when necessary.

You can access Tivoli Storage Manager backup and archive features:

- Locally through the Java Graphical User Interface (Java GUI)
- Locally through the client command line interface
- Remotely or locally through the Web client interface

Migrating from earlier versions

The following sections explain what you need to do if you are migrating to Tivoli Storage Manager 5.4 from a previous version.

Upgrade path for clients and servers

As part of a migration plan from Tivoli Storage Manager Version 5.3 to Tivoli Storage Manager Version 5.4, 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, note the following:

- A Tivoli Storage Manager Version 5.3 client can perform backup, restore, archive, and retrieve functions to a Tivoli Storage Manager Version 5.4 server.
- A Tivoli Storage Manager Version 5.4 client can perform backup, restore, archive, retrieve, and query functions to a Tivoli Storage Manager Version 5.3 server.
- A Tivoli Storage Manager Version 5.3 HSM client can perform migrate and recall functions to a Tivoli Storage Manager Version 5.4 server.
- A Tivoli Storage Manager Version 5.4 HSM client can perform migrate and recall functions to a Tivoli Storage Manager Version 5.3 server.
- If you back up, archive, or migrate data from a Tivoli Storage Manager Version 5.2 or Version 5.3 client to any Tivoli Storage Manager server, you can restore, retrieve, or recall that data using a Tivoli Storage Manager Version 5.4 client.
- If you back up, archive, or migrate data from a Tivoli Storage Manager Version 5.4 client, except for the Windows, Linux, and NetWare clients, you can restore, retrieve, or recall that data using a Tivoli Storage Manager Version 5.3 client.
- All command line administrative clients can administer Tivoli Storage Manager Version 5.3 and Version 5.4 servers.
- The format of the journal database is changed from past client versions. The journal will be reset after starting the Tivoli Storage Manager 5.4 journal service for the first time, even if the *PreserveDbOnExit=1* setting is used. As a result, the next incremental backup will be a full incremental backup, after which journaled backups can resume.
- The format of the Tivoli Storage Manager password file changed in Tivoli Storage Manager Version 5.4. 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 Version 5.4 client. Once 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 307 for information about finding the location of the password file.
2. After installing, 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.

Note: The Tivoli Storage Manager Version 5.4 client can restore encrypted data that is backed up with a previous version of Tivoli Storage Manager. You do not need to install an older Tivoli Storage Manager client or to restore the older passwords file (as described above) to perform this task.

Additional migration information

Tivoli Storage Manager V5.3 contained the following three AIX packages:

- AIX without HSM
- AIX with JFS2 HSM
- AIX with GPFS HSM

All of these AIX packages are now contained in one package.

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

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

A command line administrative client is available on all client platforms. The primary intended interface to the server is the Web administrative interface and requires a Web browser. The Web administrative interface is packaged and installed with the server. For more information about the command line administrative client and the Web administrative interface, see “Related publications” on page xv for the appropriate Tivoli Storage Manager Administrator’s Reference.

See the `client_message.chg` file in the client package for a list of new and changed messages since the Tivoli Storage Manager Version 5.3 release.

Previous Tivoli Storage Manager client versions incorrectly allowed directories in the *domain* option, and that caused problems during domain backup. Now the client checks whether a domain name is valid. If the domain name is not a file system, the Tivoli Storage Manager client will issue the following message:

Client environment requirements

This section contains Tivoli Storage Manager client environment information, components, and hardware and software requirements for the UNIX and Linux platforms. Table 4 lists the location of the environment prerequisites for each supported platform.

Attention: 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>

Table 4. Client requirements

Operating system	Page
AIX client environment	3
HP-UX PA-RISC client environment	4
HP-UX Itanium 2 client environment	5
Linux x86/x86_64 client environment	6
Linux on POWER client environment	7
Linux for IBM eServer zSeries client environment	9
z/OS UNIX System Services client environment	10
Solaris client environment	10
NDMP support requirements (Extended Edition only)	12

AIX client environment

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

Client components

- Backup-archive client (command line)
- Administrative client (command-line)
- Backup-archive Java Graphical User Interface (Java GUI)
- Web backup-archive client
- Tivoli Storage Manager API (32-bit and 64-bit)
- X/Open API (32-bit)
- Hierarchical Storage Management client

Notes:

1. The IBM Tivoli Space Manager HSM Client and documentation for AIX, while contained on the client CD, is a separate product, and therefore subject to the fees and licensing required by Tivoli and IBM.
2. For more information about the Tivoli Storage Manager API and X/Open API, see *IBM Tivoli Storage Manager Using the Application Programming Interface*, SC32-0147.
3. For more information about the Administrative client, see *IBM Tivoli Storage Manager for AIX Administrator's Reference*, SC32-0123
4. For more information about the Hierarchical Storage Management client, see *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide*, SC32-0148

Hardware requirements

- A RISC System/6000, pSeries, iSeries, or compatible hardware, as supported by AIX
- For HACMP: One additional network adapter for each machine
- Memory: 128 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.
- The following are the minimum disk space requirements:
 - Backup-archive client with 32-bit API: 43 MB
 - 64-bit API: 9 MB
 - Web client: 32 MB
 - Journal-based backup: 3 MB
 - Non-English language packs: 82 MB
 - JFS2 HSM client: 82 MB
 - GPFS HSM client: 91 MB

Software requirements

- Backup-Archive client
 - AIX 5.2 (32-bit or 64-bit kernel mode)
 - bos.rte 5.2.0.30 (Maintenance Level 3) or higher
 - bos.rte.libpthreads 5.2.0.10
 - AIX 5.3 (32-bit or 64-bit kernel mode)
 - HACMP support:
 - AIX 5.2 (32-bit or 64-bit kernel mode), with AIX APAR IY37779
 - AIX 5.3 (32-bit or 64-bit kernel mode)
 - HACMP 5.3 (with IY80002) or 5.4 (with IY89869)
 - C Set ++ Runtime (libC) level 8 or above. **Note:** All of the needed packages are included in the Tivoli Storage Manager AIX package or they can be downloaded here:
 - <http://www.ibm.com/support/docview.wss?uid=swg24012187>
- Java JRE 1.5 or 1.4.x (where x>=1) for the Java GUI
- A Mozilla 1.4 or higher browser for the Web client and to access online help and documentation

Communication methods

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

HP-UX PA-RISC client environment

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

Client components

- Backup-archive client (command-line and GUI)
- Administrative client (command line)
- Backup-archive Java Graphical User Interface (Java GUI)
- Tivoli Storage Manager API (32-bit and 64-bit)

- X/Open API
- Web backup-archive client
- Hierarchical Storage Management client

Notes:

1. The IBM Tivoli Space Manager HSM Client and documentation for HP-UX PA-RISC, while contained on the client CD, is a separate product, and therefore subject to the fees and licensing required by Tivoli and IBM.
2. For more information about the Hierarchical Storage Management client, see *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide*, SC32-0148
3. For more information about the Tivoli Storage Manager API and X/Open API, see *IBM Tivoli Storage Manager Using the Application Programming Interface*, SC32-0147.
4. For more information about the Administrative client, see *IBM Tivoli Storage Manager for HP-UX Administrator's Reference*, SC32-0124

Hardware requirements

- An HP 9000 Series 700 or 800 workstation or server
- Disk space: 100 MB for installed files, database, recovery log, and utilities; 60 MB for swap space
- Memory: 256 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

- HP-UX 11i v2 (32-bit or 64-bit)
- Java JRE 1.4.x (where x>=1) for the Java GUI
- A Mozilla 1.4 or higher browser for the Web client and to access online help and documentation
- Hierarchical Storage Management client: Veritas Filesystem 4.1 (VxFS), Veritas Volume Manager (VxVM) 3.5 or higher, OnlineJFS 3.3 or 3.5

Communications methods

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

HP-UX Itanium 2 client environment

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

Client components

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

Notes:

1. For more information about the Tivoli Storage Manager API, see *IBM Tivoli Storage Manager Using the Application Programming Interface*, SC32-0147.
2. For more information about the Administrative client, see *IBM Tivoli Storage Manager for HP-UX Administrator's Reference*, SC32-0124

Hardware requirements

- An HP-UX Itanium 2 processor
- Disk space: 72 MB hard disk space (includes full install, 9 MB recovery log, 17 MB database)
- Memory: 256 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

- HP-UX 11i V2 for Itanium 2
- Java Runtime Environment (JRE) 1.4.x (where $x \geq 1$) for the Java GUI
- A Mozilla 1.4 or higher browser for the Web client and to access online help and documentation

Communications methods

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

Linux x86/x86_64 client environment

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

Client components

- Backup-archive client (command line)
- Administrative client (command line)
- Backup-archive Java Graphical User Interface (Java GUI)
- Tivoli Storage Manager API (32-bit and 64-bit)
- Web backup-archive client
- Hierarchical Storage Management client

Notes:

1. The Tivoli Space Manager HSM Client and documentation for Linux x86/x86_64, while contained on the client CD, is a separate product, and therefore subject to the fees and licensing required by Tivoli and IBM.
2. For more information about the Tivoli Storage Manager API, see *IBM Tivoli Storage Manager Using the Application Programming Interface*, SC32-0147.
3. For more information about the Administrative client, see *IBM Tivoli Storage Manager for Linux Administrator's Reference*, SC32-0125
4. For more information about the Hierarchical Storage Management client, see *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide*, SC32-0148

Hardware requirements

- x86-based PC (Pentium or newer) or AMD64/EM64T processor
- The following are the minimum disk space requirements:
 - 8 MB for the API 32-bit
 - 13 MB for the API 64-bit (this includes the API 32-bit)
 - 26 MB for the Backup-Archive client
 - 102 MB for the HSM client
- Memory: 128 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

The backup-archive client requires the following software to run:

- One of the following:
 - Red Hat Enterprise Linux 4.0
 - SUSE SLES 9 or 10
 - Asianux 2.0
 - Novell OES
- Java JRE 1.4.x (where $x \geq 1$)
- For GPFS ACL support on Linux x86/x86_64, the Tivoli Storage Manager client uses libgpfs.so library (which comes in the standard GPFS package). For Ext2/Ext3/XFS ACL support on Linux x86/x86_64, the Tivoli Storage Manager client uses libacl.so (which comes in the libacl package).
- A Mozilla 1.4 or higher browser for the Java GUI and to access online help and documentation

For Red Hat Enterprise Linux 4.0, you need to install the compat-libstdc++-296-2.96-132.7.2.i386.rpm package. You also need to install the compat-libstdc++-33-3.2.3-47.3.i386.rpm package, before you install the API and Backup-Archive packages, for the RHEL 4 AMD64/EM64T environment.

For SuSE Linux Enterprise Server 9 64-bit support, you need to install the glibc-locale-32bit-9 package.

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

The libdevmapper.so is no longer a prerequisite. The required level is only needed for image function for LVM2 support. This check is now performed at runtime. If the library is lower level than required or is not available, image function will not support LVM2 volumes.

Communication methods

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

Linux on POWER client environment

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

Client components

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

Notes:

1. For more information about the Tivoli Storage Manager API, see *IBM Tivoli Storage Manager Using the Application Programming Interface*, SC32-0147.
2. For more information about the Administrative client, see *IBM Tivoli Storage Manager for Linux Administrator's Reference*, SC32-0125

Hardware requirements

- Machine: Linux on POWER - pSeries, iSeries, JS20 Blades, OpenPower
- The following are minimum requirements for the Backup-Archive client:
 - 64 MB of memory (RAM)
 - 80 MB of available disk space

- The following are minimum requirements for the API:
 - 64 MB of memory (RAM)
 - 20 MB of available disk space
 - 20 MB of available disk space for the API 64-bit

The default installation path for the Backup-Archive client is
`/opt/tivoli/tsm/client/ba`.

The default installation path for the API is `/opt/tivoli/tsm/client/api`.

- The GUI interface requires at least 16-bit color to display properly on the local console

Software requirements

The backup-archive client requires the following software to run:

- SUSE SLES 9 or 10 or RHEL 4
- Java JRE 1.5 or 1.4.x (where $x \geq 1$) for the Java GUI
- A Mozilla 1.4 or higher browser for the Java GUI and to access online help and documentation.

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

The Tivoli Storage Manager client uses the `libacl.so` library, so it is searched for in the following locations:

- A colon-separated list of directories in the user's `LD_LIBRARY_PATH` environment variable
- The list of libraries cached in `/etc/ld.so.cache`, `/usr/lib`, followed by `/lib`

Ensure that there is a valid library or symbolic link present in each of the specified locations.

GPFS file data is backed up by this client; however, there is no ACL support.

Attention: To use the Shared Memory Commmethod or LAN-free functions, the client must meet the Tivoli Storage Manager server or Storage Agent software requirements.

Communication methods

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

Linux for IBM eServer zSeries client environment

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

Client components

- Backup-archive client (command line)
- Administrative client (command line)
- Tivoli Storage Manager API (32-bit and 64-bit)
- Web backup-archive client

Notes:

1. For more information about the Tivoli Storage Manager API, see *IBM Tivoli Storage Manager Using the Application Programming Interface*, SC32-0147.
2. For more information about the Administrative client, see *IBM Tivoli Storage Manager for Linux Administrator's Reference*, SC32-0125

Hardware requirements

- A 9672 G5 or G6, Multiprise 3000, or zSeries 800 or 900 (31-bit and 64-bit mode)
- The following are the minimum disk space requirements:
 - 6 MB for the API 32-bit
 - 11 MB for the API 64-bit (this includes the API 32-bit)
 - 15 MB for the Backup-Archive client
- Memory: 128 MB
- A minimum screen resolution of 800 by 600 pixels is required to display the Tivoli Storage Manager Backup-Archive Client Web GUI.

Software requirements

The backup-archive client requires the following software to run:

- One of the following:
 - SUSE SLES 9 or 10 for zSeries
For SuSE Linux Enterprise Server 9 64-bit support, you also need to install the `glibc-locale-32bit-9` package.
 - Red Hat Enterprise Linux 4.0 for zSeries
For Red Hat Enterprise Linux 4 support, you also need the following compatibility packages:
`compat-libstdc++-33-3.2.3-47.3.s390.rpm`
`libgcc-3.4.3-22.1.s390.rpm`
- Java JRE 1.4.x (where $x \geq 1$) for the Java GUI
- A Mozilla 1.4 or higher browser for the Web client and to access online help and documentation
For EXT2/EXT3/XFS ACL support, you need the `acl-2.0.19-17` package.

Communication methods

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

z/OS UNIX System Services client environment

This section contains client environment information, Tivoli Storage Manager client components, and hardware and software requirements for the z/OS UNIX System Services platform.

Client components

- Backup-archive client (command-line)
- Administrative client (command-line)
- Tivoli Storage Manager API (32-bit and 64-bit)
- Web backup-archive client

Notes:

1. For more information about the Tivoli Storage Manager API, see *IBM Tivoli Storage Manager Using the Application Programming Interface*, SC32-0147.
2. For more information about the Administrative client, see *IBM Tivoli Storage Manager for z/OS Administrator's Reference*, SC32-0128

Hardware requirements

- Any zSeries architecture CPU
- The following are the disk space requirements:
 - 50 MB of available disk space for the Backup-Archive client
 - 15 MB of available disk space for the API
- A minimum screen resolution of 800 by 600 pixels is required to display the Tivoli Storage Manager Backup-Archive Client Web GUI.

Software requirements

- z/OS V1R7 or V1R8
- Java JRE 1.4.x (where $x \geq 1$) for the Java GUI

Communication methods

To use this communication method:	Install this software:	To connect to these Tivoli Storage Manager servers:
TCP/IP	IBM TCP/IP	AIX, HP-UX, Linux, Solaris, Windows, z/OS

Solaris client environment

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

Client components

- Backup-archive client (command line)
- Administrative client (command-line)

- Backup-archive Java Graphical User Interface (Java GUI)
- Tivoli Storage Manager API (32-bit and 64-bit)
- X/Open API (SPARC only)
- Web backup-archive client
- Hierarchical Storage Management client (SPARC only)
- IBM Tivoli Enterprise Space management console (GUI) (SPARC only)

Notes:

1. The IBM Tivoli Space Manager HSM Client for Solaris SPARC, while contained on the client CD, is a separate product, and therefore subject to the fees and licensing required by Tivoli and IBM.
2. For more information about the Tivoli Storage Manager API and X/Open API, see *IBM Tivoli Storage Manager Using the Application Programming Interface*, SC32-0147.
3. For more information about the Administrative client, see *IBM Tivoli Storage Manager for Sun Solaris Administrator's Reference*, SC32-0126
4. For more information about the Hierarchical Storage Management client, see *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide*, SC32-0148

Hardware requirements

- A SPARCstation based on sun4u architecture or compatible workstation
- x86 32bit and 64bit (AMD64/EM64T) based system that is explicitly compatible and fully capable of running the Solaris 10 x86 operating system
- 100 MB of available disk space

Note: The installation procedure requires up to 200 MB of temporary disk space in the /var file system.

- Memory: 512 MB
- A minimum screen resolution of 800 by 600 pixels is required to display the Tivoli Storage Manager Backup-Archive Client Java GUI and Web GUI.

Software requirements

One of the following operating systems:

- Sun Solaris 9, 10 32bit or 64bit kernel mode (SPARC)
- Sun Solaris 10 32bit and 64bit kernel mode (x86)

The following Solaris patches are required:

- Solaris 9 SPARC:
 - 111711-12 SunOS 5.9 32-bit shared library patch for C++
 - 111712-12 Sun OS 5.9 64-bit shared library patch for C++
 - 112963-22 Linker Patch
- Solaris 10 SPARC:
 - 120034-01 mntfs patch
- Solaris 10 x86:
 - 120035-01 mntfs patch

Make sure to install these patches in the given order. Solaris patches are available from this Website:<http://sunsolve.sun.com>

The backup-archive client Java GUI requires:

- The Java Runtime Environment (JRE) 1.4.x that you can download from the following Website: <http://java.sun.com/j2se/1.4.2/download.html>

Note: The "java" executable needs to be found in the system PATH. To verify that you have the right JRE installed you can use the "java -version" command from a UNIX system window.

- A Mozilla 1.4 or higher browser for the Java GUI and to access online help and documentation

Communication methods

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

NDMP support requirements (Extended Edition only)

Through support of Network Data Management Protocol (NDMP), Tivoli Storage Manager can efficiently back up and restore NAS file systems to tape drives or libraries that are locally attached to Network Appliance and EMC Celerra NAS file servers. *NDMP support is available only on IBM Tivoli Storage Manager Extended Edition.*

NDMP support requires the following hardware and software:

- Tivoli Storage Manager Extended Edition Version 5.2 AIX or Sun Solaris server or higher.

Note: A Tivoli Storage Manager Version 5.3 server is required for directory level backup. See *IBM Tivoli Storage Manager for AIX Administrator's Guide* or *IBM Tivoli Storage Manager for Sun Solaris Administrator's Guide* for details on this function.

- Tivoli Storage Manager Extended Edition Version 5.2 Sun Solaris or AIX 32-bit client or higher.
- Network Appliance or EMC Celerra NAS file server. For supported models and operating systems, refer to:
<http://www.ibm.com/tivoli/storage>
- Tape drive and tape library. For supported combinations, refer to:
<http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>

See "Backing up NAS file systems" on page 108 for further information, including how to back up and restore NAS file system images using the web client and the command line client.

Client installation media on product CD-ROMs

The client images are contained on the following Tivoli Storage Manager product CD-ROMs:

IBM Tivoli Storage Manager AIX Clients

Contains the AIX client images. The AIX client images are in the /usr/sys/inst.images directory

IBM Tivoli Storage Manager HP Clients

Contains the HP-UX client images. The images reside in the

tsmcli/'platform'/ directory structure, where 'platform' is one of the following platform designations: **hp11, hp11ia64**.

IBM Tivoli Storage Manager Solaris Clients (SPARC and x86)

Contains the Solaris client images. The images reside in the tsmcli/'platform'/ directory structure, where 'platform' is one of the following platform designations: **solaris, solaris x_86**.

IBM Tivoli Storage Manager Linux Clients

Contains the Linux x86/x86_64, Linux for zSeries and LinuxPPC client images. The images reside in the tsmcli/'platform'/ directory structure, where 'platform' is one of the following platform designations: **Linux x86/x86_64, Linux for zSeries, LinuxPPC (for Linux on POWER clients)**.

See the following Web site for the program directories with information about available UNIX client installation media:

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

You can install the clients using any of the following methods:

- Install directly from the CD-ROM.
- Create client images to install.
- Transfer installable files from the UNIX and Linux CD-ROMs to a target workstation.
- `xcopy h:\setup* x:\ /s`

Installing the backup-archive client

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

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

After installation completes, see Chapter 2, “Configuring Tivoli Storage Manager,” on page 41 for required and optional tasks to perform before using Tivoli Storage Manager.

Table 5 lists the supported UNIX clients and the location of the installation instructions for each client.

Table 5. UNIX client installation reference

Installation instructions	Page
Installing the AIX client	14
Installing the HP-UX PA-RISC client	18
Installing the HP-UX Itanium 2 client	22
Installing the Linux x86/x86_64 client	25
Installing the Linux for zSeries client	28
Installing the Linux on POWER clients	31
Installing the Solaris clients	34

For information on Installing the z/OS UNIX System Services Client, see the program directories on the following Web site:

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

Installing the AIX client

There is no longer a separate 64-bit AIX Backup-Archive client package. The client can be used on either 32-bit or 64-bit machines. If you are migrating from versions prior to Version 5.3 of the 32-bit client, do a migrate install, or uninstall the older client and then install the new client. If you are migrating from older versions of the 64-bit client, you need to uninstall your older client before installing the new client.

Due to the compiler upgrade for the AIX client in Tivoli Storage Manager V5.4, the newer version of libC (C Set ++) runtime is required. All of the needed packages are included in the Tivoli Storage Manager AIX client package, and they will overwrite the older runtime on your system during installation.

Note: The packages can also be downloaded here:

<http://www.ibm.com/support/docview.wss?uid=swg24012187>

When using INSTALLP to install this client, do not change the default field values for the following two choices:

- "AUTOMATICALLY install requisite software?"
- "OVERWRITE same or newer versions?"

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

You can install only *one* of these clients:

- tivoli.tsm.client.ba (includes the HSM JFS2 client)
- tivoli.tsm.client.ba.gpfs (required when using the GPFS HSM client)

You can install these APIs:

- tivoli.tsm.client.api.32bit
- tivoli.tsm.client.api.64bit

Notes:

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

AIX client

tivoli.tsm.client.ba

Installs the following Tivoli Storage Manager files:

- Backup-archive base files
- Backup-archive common files
- Backup-archive Web client
- Image backup client
- NAS backup client

tivoli.tsm.client.ba.gpfs

Installs the following Tivoli Storage Manager files:

- Backup-archive GPFS files

- GPFS common files

tivoli.tsm.client.hsm.gpfs

Installs the Tivoli Storage Manager Hierarchical Storage Management GPFS client.

tivoli.tsm.client.hsm

Installs the Tivoli Storage Manager Hierarchical Storage Management JFS2 client.

tivoli.tsm.client.jbb

Installs the Tivoli Storage Manager journal-based backup component.

tivoli.tsm.client.api.32bit

Installs the 32-bit API.

tivoli.tsm.client.api.64bit

Installs the 64-bit API.

tivoli.tsm.client.msg.lang

Installs the NL messages for API, where *lang* is the appropriate language code from Table 6.

During installation:

- The Tivoli Storage Manager backup-archive, Web client, and administrative client (dsmadm) files are installed in the `/usr/tivoli/tsm/client/ba/bin` directory.
- The administrative client is installed in the `/usr/tivoli/tsm/client/admin/bin` directory.
- The Tivoli Storage Manager API files are installed in the `/usr/tivoli/tsm/client/api/bin` directory.
- The Tivoli Storage Manager API 64 files are installed in the `/usr/tivoli/tsm/client/api/bin64` directory.

Table 6. AIX client: Language codes for installation packages

Language	Language code
Simplified Chinese	zh_CN
Traditional Chinese	zh_TW, Zh_TW
Czech	cs_CZ
French	fr_FR
German	de_DE
Hungarian	hu_HU
Italian	it_IT
Japanese	ja_JP, Ja_JP, JA_JP
Korean	ko_KR
Polish	pl_PL
Brazilian Portuguese	pt_BR
Russian	ru_RU
Spanish	es_ES

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

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

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

A .toc file is created in that directory.

To install Tivoli Storage Manager from the CD-ROM:

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

Notes:

1. When installing the clients in a non-English locale, the corresponding local message filesets are required in addition to the clients from the **Install and Update From ALL Available Software** panel. For example, when installing the Tivoli Storage Manager client in es_ES locale, the following message fileset is also required:

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

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

The Tivoli Storage Manager files are installed in the `/usr/tivoli/tsm/client/ba/bin` directory. If you move the Tivoli Storage Manager files to another directory, you must perform the following steps:

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

After installation completes, see Chapter 2, “Configuring Tivoli Storage Manager,” on page 41 for required and optional tasks to perform before using Tivoli Storage Manager.

Uninstalling the AIX client

1. From the AIX command line, type:
 `smitty remove`
2. Hit the ENTER key.
3. In the SOFTWARE name field, press F4 to list the TSM file sets that you want to uninstall; hit the ENTER key.
4. Select the TSM file sets that you want to uninstall; press the ENTER key.
5. In the PREVIEW only? (remove operation will NOT occur) field, select No; hit the ENTER key.

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

Installing the HP-UX PA-RISC client

If a previous version of Tivoli Storage Manager is installed, remove it before installing a new version.

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

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

Table 7. Package Names and Directories

This software package	Is in this directory
TIVsmCapi	/opt/tivoli/tsm/client/api
TIVsmCapi64	/opt/tivoli/tsm/client/api
TIVsmCba	/opt/tivoli/tsm/client/ba
TIVsmChsm	/opt/tivoli/tsm/client/hsm
TIVsmC.msg<xx_Xx>	/opt/tivoli/tsm/client/lang/<xx_Xx>

Where <xx_Xx> defines the language installed

Installation steps: Follow these steps to install Tivoli Storage Manager on HP-UX PA-RISC clients:

1. Login as the root user and mount the Tivoli Storage Manager UNIX client CD-ROM to /cdrom.
2. Change to the directory where the packages are stored: `cd /cdrom/tsmcli/hp11`
If the Tivoli Storage Manager UNIX client CD-ROM is not mounted to /cdrom, or if the packages are stored in a different directory (e.g. downloaded by ftp), change to the correct directory.
3. To install the Tivoli Storage Manager API and the Tivoli Storage Manager Backup-Archive clients, enter the following command:
 - `/usr/sbin/swinstall -x mount_all_filesystems=false -v -s `pwd`/TIVsmC TIVsm`

Note: ``pwd`` can be replaced with the absolute name of the current directory.

If you want to install only the API, omit the last **TIVsm** from the above command, and mark just the fileset in the **swinstall** user interface that you want to install:

CLIENT for the backup-archive client
CLIENT_API for the API, or
CLIENT_API64 for the API64

Notes:

- a. The client needs the API for Raw Logical Volume backup. If you mark **CLIENT** for installation, the API is also installed.

- b. Make sure that you install these packages in the given order. This is required because the Tivoli Storage Manager API package is a prerequisite of the backup-archive client package.
 - c. The Web client is an integral part of the backup-archive client package and cannot be installed without it.
4. Refer to *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide* for information about installing the HSM client.
 5. The Tivoli Storage Manager default language is English. To install additional languages (optional), enter the commands listed for each additional language.

To install this language	Enter this command
NLS French messages	TIVsmC.msg.fr_FR
NLS German messages	TIVsmC.msg.de_DE
NLS Italian messages	TIVsmC.msg.it_IT
NLS Spanish messages	TIVsmC.msg.es_ES
NLS Portuguese messages	TIVsmC.msg.pt_BR
NLS Czech messages	TIVsmC.msg.cs_CZ
NLS Hungarian messages	TIVsmC.msg.hu_HU
NLS Polish messages	TIVsmC.msg.pl_PL
NLS Russian messages	TIVsmC.msg.ru_RU
NLS Japanese messages	TIVsmC.msg.ja_JP
NLS Korean messages	TIVsmC.msg.ko_KR
NLS Traditional Chinese messages (EUC)	TIVsmC.msg.zh_CN
NLS Traditional Chinese Big5 messages	TIVsmC.msg.zh_TW

Table 8. HP-UX PA-RISC client: Language codes for installation packages

Language	Language code
French	fr_FR
German	de_DE
Italian	it_IT
Spanish	es_ES
Brazilian Portuguese	pt_BR
Czech	cs_CZ
Hungarian	hu_HU
Polish	pl_PL
Russian	ru_RU
Japanese	ja_JP
Korean	ko_KR
Traditional Chinese	zh_TW

After installation completes, see Chapter 2, “Configuring Tivoli Storage Manager,” on page 41 for required and optional tasks to perform before using Tivoli Storage Manager.

| **Increasing the default limit of the data segment size:** The default limit of the
| data segment size of a process in HP-UX PA-RISC 11i V2 is 64 MB. When backing
| up large file systems, the Tivoli Storage Manager client might exceed this limit and
| run out of memory. To increase this limit you can modify the kernel as follows:

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

Uninstalling the HP-UX PA-RISC client

Follow these steps to uninstall all the packages related to Tivoli Storage Manager and HSM, including the command line, GUI, Web GUI and administrative client components. Make sure that you uninstall the packages in the given order.

1. HSM package: If the HSM client is installed, it needs to be removed first. The steps for removing the HSM client are listed in *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide*.

Or, see “Uninstalling Tivoli Storage Manager on Sun Solaris clients” on page 40 for detailed information.

To completely remove the HSM client, without reinstalling a new version, you must remove HSM support from all managed file system. If you do not completely remove HSM support from all managed file systems, the data of migrated files will not be accessible after removing the HSM package. To remove HSM support from all managed file systems, enter the following command:

```
dsmmigfs remove <fs>
```

Then enter the following command, as root user:

```
/usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm.CLIENT_HSM
```

To upgrade to the latest version of the HSM client, you do not need to remove HSM support from all managed file systems. In this case, enter the following command as root user:

```
/usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm.CLIENT_HSM
```

To install the newer version of the HSM client, follow the instructions in

Notes:

- a. To upgrade the HSM client, you have to uninstall and install the backup-archive client and the API, so that the product will be at the same version and level. Mixed installation is not supported.
 - b. Activity on the HSM-managed file systems should be avoided during software installation. Access to migrated files during the upgrade will not be possible (the process of trying to read a stub would hang).
 - c. There are special considerations for installing HSM. See *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide* for these special considerations.
2. Tivoli Storage Manager backup-archive package: Enter the following command to uninstall the Tivoli Storage Manager backup-archive client. This will uninstall all of the components of the backup-archive client (command-line, GUI, Web GUI and the administrative client).

```
• /usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm.CLIENT
```

Note: You cannot uninstall a single component of this package (e.g. Web GUI) without uninstalling the complete package.

3. Tivoli Storage Manager Language packages: If one or more language messages packages are installed, they must be removed before removing the API package. Enter the following command as root user:

```
• /usr/sbin/swremove -x mount_all_filesystems=false -v  
TIVsm.CLIENT_msg_xx_XX . Refer to Table 8 on page 19 for the language  
codes.
```

Installing the HP-UX Itanium 2 client

The following source packages are available on the installation media:

tsmcli/hp11/TIVsmC

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

TIVsm.CLIENT

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

TIVsm.CLIENT_API64

Contains the 64-bit API with the English message catalogues.

Note: Additional language support is available under the top level product name of **tsmcli/hp11/TIVsmC.msg.lang**. In this package, the additional language support is available for all components under the top level product name TIVsm. The component name is TIVsm.CLIENT_msg_lang. Replace lang with the appropriate language code from Table 9.

tsmcli/hp11/TIVsmCapi64

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

TIVsm.CLIENT_API64

Contains the 64-bit API with the English message catalogues.

Table 9. HP-UX Itanium 2 client: Language codes for installation packages

Language	Language code
Simplified Chinese	zh_CN
Traditional Chinese	zh_TW
Czech	cs_CZ
French	fr_FR
German	de_DE
Hungarian	hu_HU
Italian	it_IT
Japanese	ja_JP
Korean	ko_KR
Polish	pl_PL
Brazilian Portuguese	pt_BR
Russian	ru_RU
Spanish	es_ES

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

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

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

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

Replace *lang* with the appropriate language code from Table 9 on page 22.

To install from the CD-ROM, log in as the root user, mount the CD-ROM to /cdrom, and change directory to tsmcli/hp11. If you downloaded from ftp, go to the directory where the installable image is located. Enter the following command:

```
/usr/sbin/swinstall -x mount_all_filesystems=false -v -s `pwd`/TIVsmC TIVsm64
```

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

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

CLIENT for the backup-archive client, or
CLIENT_API64 for the API

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

To install additional languages, execute the following command:

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

Notes:

1. '*pwd*' can be used instead of the absolute name of the current directory.
2. Replace *lang* with the appropriate language code from Table 9 on page 22.

During installation:

- The Tivoli Storage Manager backup-archive, Web client, and administrative client (dsmadm) files are installed in the /opt/tivoli/tsm/client/ba/bin directory.
- The Tivoli Storage Manager API 64 files are installed in the /opt/tivoli/tsm/client/api/bin64 directory.

Increasing default limit of data segment size: The default limit of the data segment size of a process in HP-UX 11i v2 is 64 MB. When backing up large file systems, the Tivoli Storage Manager client might exceed this limit and run out of memory. To increase this limit you can modify the kernel as follows:

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

After installation completes, see Chapter 2, "Configuring Tivoli Storage Manager," on page 41 for required and optional tasks to perform before using Tivoli Storage Manager.

Uninstalling the HP-UX Itanium 2 client

Follow these steps to uninstall all of the packages related to Tivoli Storage Manager including command line, GUI, Web GUI and administrative Client components. Make sure that you uninstall the packages in the given order. Please note that you cannot uninstall a single component of this package (e.g. Web GUI) without uninstalling the complete package.

1. Enter the following command to uninstall the Tivoli Storage Manager Backup-Archive Client. This will uninstall all of the components of the Backup-Archive Client (command-line, Web GUI and the administrative Client).

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

2. If one or more Tivoli Storage Manager language messages packages have been installed, these must be removed before removing the API package. Execute the following command as root user:

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

Where xx_XX is the language you wish to remove.

Current supported languages are cs_CZ, de_DE, es_ES, it_IT, fr_FR, hu_HU, ja_JP, ko_KR, pl_PL, pt_BR, ru_RU, zh_CN and zh_TW.

3. If you also want to remove the CLIENT_API fileset, execute the following command:

```
/usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm64.CLIENT_API64
```

Installing the Linux x86/x86_64 client

If a previous version of Tivoli Storage Manager is installed, remove it before installing a new version.

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

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

Table 10. Package Names and Directories

This software package	Is in this directory
TIVsm-API.i386.rpm	/opt/tivoli/tsm/client/api/bin
TIVsm-API64.i386.rpm	/opt/tivoli/tsm/client/api/bin64
TIVsm-BA.i386.rpm	/opt/tivoli/tsm/client/ba
TIVsm- msg.<xx_XX>.i386.rpm	/opt/tivoli/tsm/client/lang/<xx_XX> Where <xx_XX> defines the language installed

Installation steps: Follow these steps to install Tivoli Storage Manager:

1. Login as the root user and mount the Tivoli Storage Manager UNIX client CD-ROM to /cdrom.
2. Change to the directory where the packages are stored: `cd /cdrom/tsmcli/linux86`

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

Notes:

- a. Make sure that you install these packages in the given order. This is required because the Tivoli Storage Manager API package is a prerequisite of the Backup-Archive client package.
 - b. The Web client is an integral part of the backup-archive client package and cannot be installed without it.
3. Enter the following commands to install the 32-bit API and the Tivoli Storage Manager Backup-Archive (B/A) client. This will install the command-line and the administrative client:
 - `rpm -i TIVsm-API.i386.rpm`
 - `rpm -i TIVsm-BA.i386.rpm`
 4. Enter the following commands to install the 64-bit API and the Tivoli Storage Manager Backup-Archive client. This will install the command-line and the administrative client:
 - `rpm -i TIVsm-API.i386.rpm` (required for 32-bit and 64-bit)
 - `rpm -i TIVsm-API64.i386.rpm` (required for 64-bit only)
 - `rpm -i TIVsm-BA.i386.rpm`
 5. The Tivoli Storage Manager default language is English. To install additional languages (optional), enter the commands listed for each additional language.

To install this language	Enter this command
NLS French messages	<code>rpm -i TIVsm-msg.fr_FR.i386.rpm</code>
NLS German messages	<code>rpm -i TIVsm-msg.de_DE.i386.rpm</code>
NLS Italian messages	<code>rpm -i TIVsm-msg.it_IT.i386.rpm</code>
NLS Spanish messages	<code>rpm -i TIVsm-msg.es_ES.i386.rpm</code>
NLS Portuguese messages	<code>rpm -i TIVsm-msg.pt_BR.i386.rpm</code>
NLS Czech messages	<code>rpm -i TIVsm-msg.cs_CZ.i386.rpm</code>
NLS Hungarian messages	<code>rpm -i TIVsm-msg.hu_HU.i386.rpm</code>
NLS Polish messages	<code>rpm -i TIVsm-msg.pl_PL.i386.rpm</code>
NLS Russian messages	<code>rpm -i TIVsm-msg.ru_RU.i386.rpm</code>
NLS Japanese messages	<code>rpm -i TIVsm-msg.ja_JP.i386.rpm</code>
NLS Korean messages	<code>rpm -i TIVsm-msg.ko_KR.i386.rpm</code>
NLS Traditional Chinese messages (EUC)	<code>rpm -i TIVsm-msg.zh_CN.i386.rpm</code>
NLS Traditional Chinese Big5 messages	<code>rpm -i TIVsm-msg.zh_TW.i386.rpm</code>

Table 11. Linux x86/x86_64 client: Language codes for installation packages

Language	Language code
French	fr_FR
German	de_DE
Italian	it_IT
Spanish	es_ES
Brazilian Portuguese	pt_BR
Czech	cs_CZ
Hungarian	hu_HU
Polish	pl_PL
Russian	ru_RU
Japanese	ja_JP
Korean	ko_KR
Traditional Chinese	zh_TW

After installation completes, see Chapter 2, “Configuring Tivoli Storage Manager,” on page 41 for required and optional tasks to perform before using Tivoli Storage Manager.

Uninstalling the Linux x86/x86_64 client

Follow these steps to uninstall all the packages related to Tivoli Storage Manager and HSM, including the command line, GUI, Web GUI and administrative client components. Make sure that you uninstall the packages in the given order.

1. HSM package: Check the *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide* for instructions on removing HSM and installing a new version.
2. Tivoli Storage Manager backup-archive package: Enter the following command to uninstall the Tivoli Storage Manager backup-archive client. This will uninstall all of the components of the backup-archive client (command-line, GUI, Web GUI and the administrative client).

- **rpm -e TIVsm-BA**

Note: You cannot uninstall a single component of this package (e.g. Web GUI) without uninstalling the complete package.

3. Tivoli Storage Manager Language packages: If one or more language messages packages are installed, they must be removed before removing the API package. Enter the following command as root user:

- **rpm -e TIVsm-msg.xx_XX**

where **xx_XX** is the language you want to remove. Refer to Table 11 on page 26 for a list of the current supported languages.

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

- **rpm -e TIVsm-API64**

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

- **rpm -e TIVsm-API**

Installing the Linux for zSeries client

If a previous version of Tivoli Storage Manager is installed, remove it before installing a new version.

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

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

Table 12. Package Names and Directories

This software package	Is in this directory
TIVsm-API.s390.rpm	/opt/tivoli/tsm/client/api/bin
TIVsm-API64.s390.rpm	/opt/tivoli/tsm/client/api/bin64
TIVsm-BA.s390.rpm	/opt/tivoli/tsm/client/ba
TIVsm- msg.<xx_XX>.s390.rpm	/opt/tivoli/tsm/client/lang/<xx_XX> Where <xx_XX> defines the language installed

Installation steps: Follow these steps to install Tivoli Storage Manager:

1. Login as the root user and mount the Tivoli Storage Manager UNIX client CD-ROM to /cdrom.
2. Change to the directory where the packages are stored: `cd /cdrom/tsmcli/linux390`

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

3. To install the 32 Bit API and Backup-Archive client, enter the following commands:

```
rpm -i TIVsm-API.s390.rpm  
rpm -i TIVsm-BA.s390.rpm
```

4. To install the 64 Bit API and Backup-Archive client, enter the following commands:

```
rpm -i TIVsm-API.s390.rpm  
rpm -i TIVsm-API64.s390.rpm  
rpm -i TIVsm-BA.s390.rpm
```

Notes:

1. This will install the command-line, the GUI and the administrative client.
2. Make sure that you install these packages in the given order. This is required because the Tivoli Storage Manager API package is a prerequisite of the backup-archive client package.
3. The Web client is an integral part of the backup-archive client package and cannot be installed without it.

To install this language	Enter this command
NLS French messages	<code>rpm -i TIVsm-msg.fr_FR.s390.rpm</code>
NLS German messages	<code>rpm -i TIVsm-msg.de_DE.s390.rpm</code>

To install this language	Enter this command
NLS Italian messages	<code>rpm -i TIVsm-msg.it_IT.s390.rpm</code>
NLS Spanish messages	<code>rpm -i TIVsm-msg.es_ES.s390.rpm</code>
NLS Portuguese messages	<code>rpm -i TIVsm-msg.pt_BR.s390.rpm</code>
NLS Czech messages	<code>rpm -i TIVsm-msg.cs_CZ.s390.rpm</code>
NLS Hungarian messages	<code>rpm -i TIVsm-msg.hu_HU.s390.rpm</code>
NLS Polish messages	<code>rpm -i TIVsm-msg.pl_PL.s390.rpm</code>
NLS Russian messages	<code>rpm -i TIVsm-msg.ru_RU.s390.rpm</code>
NLS Japanese messages	<code>rpm -i TIVsm-msg.ja_JP.s390.rpm</code>
NLS Korean messages	<code>rpm -i TIVsm-msg.ko_KR.s390.rpm</code>
NLS Traditional Chinese messages (EUC)	<code>rpm -i TIVsm-msg.zh_CN.s390.rpm</code>
NLS Traditional Chinese Big5 messages	<code>rpm -i TIVsm-msg.zh_TW.s390.rpm</code>

Table 13. Linux for zSeries client: Language codes for installation packages

Language	Language code
French	fr_FR
German	de_DE
Italian	it_IT
Spanish	es_ES
Brazilian Portuguese	pt_BR
Czech	cs_CZ
Hungarian	hu_HU
Polish	pl_PL
Russian	ru_RU
Japanese	ja_JP
Korean	ko_KR
Traditional Chinese	zh_TW

After installation completes, see Chapter 2, “Configuring Tivoli Storage Manager,” on page 41 for required and optional tasks to perform before using Tivoli Storage Manager.

Uninstalling the Linux for zSeries client

Follow these steps to uninstall all the packages related to Tivoli Storage Manager and HSM, including the command line, GUI, Web GUI and administrative client components. Make sure that you uninstall the packages in the given order.

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

- **rpm -e TIVsm-BA**

Note: You cannot uninstall a single component of this package (e.g. Web GUI) without uninstalling the complete package.

2. Tivoli Storage Manager Language packages: If one or more language messages packages are installed, they must be removed before removing the API package. Enter the following command as root user:

- **rpm -e TIVsm-msg.xx_XX**

where **xx_XX** is the language you want to remove. Refer to Table 13 on page 29 for a list of the current supported languages.

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

- **rpm -e TIVsm-API**

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

- **rpm -e TIVsm-API64**

Installing the Linux on POWER clients

The following installation options are available in uncompressed packages on the CD:

tsmcli/linuxppc/TIVsm-BA.ppc64.rpm

Installs the backup-archive client (command-line), administrative client (command-line), the Web client, and documentation.

Note: Additional language support is available with component names **TIVsm-BA.msg.lang.ppc64.rpm**. Replace *lang* with the appropriate language code from Table 14 on page 32.

tsmcli/linuxppc/TIVsm-API.ppc64.rpm

Installs the Application Programming Interface (API) containing the Tivoli Storage Manager API shared libraries and samples.

Note: Additional language support is available with component names **TIVsm-API.msg.lang.ppc64.rpm**. Replace *lang* with the appropriate language code from Table 14 on page 32.

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

```
rpm -e TIVsm-BA
and for the API:
rpm -e TIVsm-API
```

If an additional Language package is installed, you must uninstall it first:

```
rpm -e TIVsm-BA-.msg.lang
rpm -e TIVsm-API.msg.lang
```

Replace *lang* with the appropriate language code from Table 14 on page 32.

Note: The package version number is not needed for uninstall.

Use the following procedure to install the Tivoli Storage Manager clients:

1. Log in as the root user and mount the CD-ROM to /cdrom.
2. Enter the following directory path where the installation packages reside on the CD:

```
/cdrom/tsmcli/linuxppc
```

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

```
First:
rpm -i TIVsm-API.ppc64.rpm
Then:
rpm -i TIVsm-BA.ppc64.rpm
```

To circumvent the dependence check use the `--nodeps` option. Please be aware to check the dependencies manually:

```
First:
rpm -i --nodeps TIVsm-API.ppc64.rpm
Then:
rpm -i --nodeps TIVsm-BA.ppc64.rpm
```

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

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

```
rpm -i TIVsm-msg.lang.ppc64.rpm
```

Replace *lang* with the appropriate language code from Table 14.

Table 14. Linux on POWER clients: Language codes for installation packages

Language	Language code
Simplified Chinese	zh_CN
Traditional Chinese	zh_TW
Czech	cs_CZ
French	fr_FR
German	de_DE
Hungarian	hu_HU
Italian	it_IT
Japanese	ja_JP
Korean	ko_KR
Polish	pl_PL
Brazilian Portugese	pt_BR
Russian	ru_RU
Spanish	es_ES

During installation :

- The Tivoli Storage Manager backup-archive, Web client, Java GUI, and administrative client (dsmadm) files are installed in the `/opt/tivoli/tsm/client/ba/bin` directory.
- The Tivoli Storage Manager API files are installed in the `/opt/tivoli/tsm/client/api/bin` directory.
- The Tivoli Storage Manager documentation files are installed in the `/opt/tivoli/tsm/client/books/html/tsmunix` and `/opt/tivoli/tsm/client/books/pdf/tsmunix` directories.

After installation completes, see Chapter 2, “Configuring Tivoli Storage Manager,” on page 41 for required and optional tasks to perform before using Tivoli Storage Manager.

Uninstalling the Linux on POWER clients

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

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

2. For the API, enter this command:

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

3. If an additional language package is installed, uninstall it by entering these commands:

```
rpm -e TIVsm-BA.msg.xx_XX
```

```
rpm -e TIVsm-API.msg.xx_XX
```

Replace xx_XX with the language that you installed earlier.

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

Installing the Solaris clients

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

The following installation packages are available on the CD in the following locations:

- /cdrom/tsmcli/solaris (Solaris SPARC packages)
- /cdrom/tsmcli/solaris_x86 (Solaris x86 packages)

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

Table 15. Installation package names and descriptions

Package	Package Name	Package Description
Tivoli Storage Manager Application Programming Interface (API)	TIVsmCapi.pkg	Contains the Tivoli Storage Manager API shared libraries and samples. Note: On Solaris SPARC and x86-based systems the installation process will install both the 32 bit and 64 bit Tivoli Storage Manager API.

Table 15. Installation package names and descriptions (continued)

Package	Package Name	Package Description
Tivoli Storage Manager Backup-Archive (B/A)	TIVsmCba.pkg	<p>Contains the following components:</p> <ul style="list-style-type: none"> • Tivoli Storage Manager Backup-Archive Client (command-line and GUI) • Tivoli Storage Manager Administrative Client (command-line) • Tivoli Storage Manager Web Backup-Archive Client <p>Notes:</p> <ol style="list-style-type: none"> 1. The Tivoli Storage Manager Backup-Archive package includes the 32-bit Backup-Archive Client which is used in 32-bit and 64-bit kernel modes on all supported Solaris releases (see “Software requirements” on page 11). 2. The 32 bit Backup-Archive Client works with 32-bit and 64-bit Tivoli Storage Manager servers. 3. TCP/IP and Shared memory are supported as communication methods.
Tivoli Storage Manager for Space Management (HSM) (SPARC only)	TIVsmChsm.pkg	<p>Contains the following components:</p> <ul style="list-style-type: none"> • Tivoli Storage Manager for Space Management (HSM) • Space Management Agent • IBM Tivoli Enterprise Space Management Console (HSM Java GUI) <p>For more information about the HSM Java GUI, refer to the README_hsm_enu.htm file.</p>
Tivoli Storage Manager French language	TIVsmClFr.pkg	<p>Contains French localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.</p>
Tivoli Storage Manager German language	TIVsmClDe.pkg	<p>Contains German localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.</p>

Table 15. Installation package names and descriptions (continued)

Package	Package Name	Package Description
Tivoli Storage Manager Italian language	TIVsmCIIt.pkg	Contains Italian localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Portuguese language	TIVsmCI Pt.pkg	Contains Portuguese localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Spanish language	TIVsmCIEs.pkg	Contains Spanish localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Czech language	TIVsmCICs.pkg	Contains Czech localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Hungarian language	TIVsmCIHu.pkg	Contains Hungarian localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Polish language	TIVsmCIPl.pkg	Contains Polish localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Russian language	TIVsmCIRu.pkg	Contains Russian localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Japanese language	TIVsmCIJa.pkg	Contains Japanese localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Korean language	TIVsmCIKo.pkg	Contains Korean localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Simplified Chinese language	TIVsmCI Sc.pkg	Contains Simplified Chinese localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.

Table 15. Installation package names and descriptions (continued)

Package	Package Name	Package Description
Tivoli Storage Manager Traditional Chinese (EUC) language	TIVsmCITc.pkg	Contains Traditional Chinese (EUC) localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.
Tivoli Storage Manager Traditional Chinese (BIG5) language	TIVsmCIBc.pkg	Contains Traditional Chinese (BIG5) localized messages for Tivoli Storage Manager API, Tivoli Storage Manager Backup-Archive Client and HSM Client.

Installation steps:

- First installation versus replacement installation

If a previous version of Tivoli Storage Manager has been installed, remove it before installing a new version. Refer to “Uninstalling Tivoli Storage Manager on Sun Solaris clients” on page 40 for information about removing previous Tivoli Storage Manager packages.

Follow these steps to install Tivoli Storage Manager:

1. Login as the root user and mount the Tivoli Storage Manager UNIX client CD-ROM to /cdrom.
2. Change to the directory where the packages are stored: `cd /cdrom/tsmcli/solaris`

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

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

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

This will install the command-line, the GUI and the administrative client.

Notes:

- a. Make sure that you install these packages in the given order. This is required because the Tivoli Storage Manager B/A is a prerequisite of the HSM client package and the Tivoli Storage Manager API package is a prerequisite of the backup-archive client package.
 - b. The Web client is an integral part of the backup-archive client package and cannot be installed without it.
5. To install the Tivoli Storage Manager HSM client (SPARC only), enter the following command:

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

This will install the Tivoli Storage Manager for Space Management command-line (HSM), the Space Management Agent (hsmagent), and the IBM Tivoli Enterprise Space Management Console (GUI).

Notes:

- a. The UNIX "IBM Tivoli Enterprise Space Management Console" and the "Space Management Agent" are parts of the UNIX HSM package and cannot be installed without it.
- b. Tivoli Storage Manager API and backup-archive packages are prerequisites of the HSM Client package and must be installed first.
- c. On Solaris 10 systems, the HSM package can be installed only by the administrator of the Global Zone and cannot be installed in the Local Zones. Refer to *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide* for more information.

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

- If Tivoli Storage Manager is already installed in the global zone, creation of the links is not necessary. The links are already present and they are pointing to the correct executable and libraries.
- If Tivoli Storage Manager is not installed in the global zone, a warning message will be provided asking the global administrator to create the required links. A list of the links with the related command will be provided in the messages.

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

To install this language	Enter this command
NLS French messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCIFr.pkg TIVsmCIFr</code>
NLS German messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCIDe.pkg TIVsmCIDe</code>
NLS Italian messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCIIt.pkg TIVsmCIIt</code>
NLS Spanish messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCIEs.pkg TIVsmCIEs</code>
NLS Portuguese messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCIPt.pkg TIVsmCIPt</code>
NLS Czech messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCICs.pkg TIVsmCICs</code>
NLS Hungarian messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCIHu.pkg TIVsmCIHu</code>
NLS Polish messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCIPl.pkg TIVsmCIPl</code>
NLS Russian messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCIRu.pkg TIVsmCIRu</code>
NLS Japanese messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCIJa.pkg TIVsmCIJa</code>
NLS Korean messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCIKo.pkg TIVsmCIKo</code>
NLS Simplified Chinese messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCISc.pkg TIVsmCISc</code>
NLS Traditional Chinese messages (EUC)	<code>pkgadd -a ./tsmadmin -d ./TIVsmCITc.pkg TIVsmCITc</code>

To install this language	Enter this command
NLS Traditional Chinese Big5 messages	<code>pkgadd -a ./tsmadmin -d ./TIVsmCIBc.pkg TIVsmCIBc</code>

After installation completes, see Chapter 2, “Configuring Tivoli Storage Manager,” on page 41 for required and optional tasks to perform before using Tivoli Storage Manager.

Uninstalling Tivoli Storage Manager on Sun Solaris clients

Follow these steps to uninstall all the packages related to Tivoli Storage Manager and HSM, including the command line, GUI, Web GUI and administrative client components. Make sure that you uninstall the packages in the given order.

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

1. Hierarchical Storage Manager (HSM) package: See *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide* for information about uninstalling HSM.
2. Tivoli Storage Manager backup-archive package: Enter the following command to uninstall the Tivoli Storage Manager backup-archive client. This will uninstall all of the components of the backup-archive client (command-line, GUI, Web GUI and the administrative client).
 - **pkgrm -n -a ./tsmadmin TIVsmCba**

Note: You cannot uninstall a single component of this package (e.g. Web GUI) without uninstalling the complete package.

3. Tivoli Storage Manager Language packages: If one or more language messages packages are installed, they must be removed before removing the API package. Execute the following command as root user:
 - **pkgrm -n -a ./tsmadmin TIVsmCIfR TIVsmCIdE TIVsmCIIt TIVsmCIPt \ TIVsmCIEs TIVsmCICs TIVsmCIHu TIVsmCIPi TIVsmCIRu \ TIVsmCIJa TIVsmCIKo TIVsmCISc TIVsmCITc TIVsmCIBc**
4. Tivoli Storage Manager API package: Enter the following command to uninstall the Tivoli Storage Manager API. The API cannot be removed if the TSM backup-archive client is installed. The Tivoli Storage Manager backup-archive client must be removed first.
 - **pkgrm -n -a ./tsmadmin TIVsmCapi**

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

Table 16. Tivoli Storage Manager server Installation Guides

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

Chapter 2. Configuring Tivoli Storage Manager

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

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

Required configuration tasks include the following:

Task	Page
Creating and modifying the client system options file (required root user or authorized user task)	43
Registering your workstation with a server (required root user or authorized user task)	57

Optional configuration tasks include the following:

Task	Page
Setting environment variables	46
Creating a default client user options file (optional root user or authorized user task)	44
Creating a customized client user options file (optional user task)	45
Configuring the Web client	50
Configuring the client scheduler	51
Configuring the journal daemon	56

Root and authorized user tasks

The phrases **root user** and **Authorized User** identify tasks that only root users and Authorized Users can perform. An **Authorized User** is any user running with a real user ID of 0 (root) or who owns the Tivoli Storage Manager executable with the owner execution permission bit set to `s`.

Table 17 shows the tasks that can and cannot be performed by the root user, Authorized User, and the non-Authorized User.

Table 17. Root user, Authorized User, and non-Authorized user tasks

Task	Root User	Authorized User	Non-Authorized User
Install the backup-archive client.	Yes.	No.	No.
Registering new nodes with Tivoli Storage Manager server.	Yes.	Yes.	No, even when the registration is set to <i>open</i> on the server.

Table 17. Root user, Authorized User, and non-Authorized user tasks (continued)

Task	Root User	Authorized User	Non-Authorized User
Set or recreate the Tivoli Storage Manager password for client workstations.	Yes, if the password file has been deleted.	Yes, if the password file has been deleted and you have write permission to the file.	No, not if the password file has been deleted.
Create and modify dsm.sys.	Yes.	Yes, if you have write permission to the file.	No.
Create and modify the client user options file (dsm.opt).	Yes.	Yes, if you have write permission to the file.	Yes, if you own the file.
Create and modify an include-exclude list.	Yes.	Yes.	No.
Backup	Yes.	Yes, if you have read permission, regardless of ownership.	Yes, if you own the file.
Restore	Yes. When restoring to a new location or the same location, file permission and ownership are preserved.	Yes. However, the operating system 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. Note: Retrieve for an authorized user, using the dsmj GUI only works for individual files and not when selecting directories. Retrieve each file individually using the dsmj client or use the dsme client to restore an entire package or directory structure.	Yes, if you archived the file. Ownership of all retrieved objects will be changed to the non-Authorized user.
Client scheduler	Yes.	Yes.	No.
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.

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

During installation, the sample, `dsm.sys.smp`, is placed in the installation directory.

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

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:

Servername	server_a
COMMethod	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 305 for more information.

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

You can also specify a default server and a migration server (if you have the HSM client installed on your workstation) in your `dsm.sys` file. For more information, see “Defaultserver” on page 217.

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 169 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** → **Preferences**.
- Use your favorite text editor.

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 45 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.

For systems other than Solaris, copy the `dsm.opt.smp` file to `dsm.opt` in your installation directory and modify the required options according to your needs.

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 169 for more information about these options.

If you are a root user, you can create or modify a default client user options file for all users on your workstation. From the UNIX command line:

- Change to the directory containing the sample file.
 - For systems other than Solaris, 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 in the same directory as the `dsm.sys.smp` file, or in any directory for which you have write access. If you rename your client user options file or place it in a different directory, you can set the `DSM_CONFIG` environment variable or use the `optfile` option to point to your new client user options file.
 - **For the client user options file:** You can rename your client user options file and store it 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 **Preferences**. The Preferences editor updates the client configuration files, `dsm.opt` and `dsm.sys`, if any options have changed. If you update the `dsm.opt` file during a session, you must restart the session to pick up the changes.

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

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 your home directory as dsm.opt, or a new file name of your choice. You can store your client user options file in any directory to which you have write access.
3. Set the DSM_CONFIG environment variable to point to your new client user options file. For instructions to set this variable, see section, “Setting environment variables” on page 46.
4. Edit your dsm.opt file as appropriate for your system or use the Tivoli Storage Manager Preferences editor by selecting **Edit** → **Preferences** from the Tivoli Storage Manager GUI.

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

1. Launch the GUI by either double-clicking the **TSM Backup Client** icon on your desktop, or going to the Windows® task bar and clicking **Start** → **Programs** → **Tivoli Storage Manager** → **Backup Archive GUI**.
2. From the **Utilities** menu, select **Setup Wizard**. This will launch the Tivoli Storage Manager Client Configuration Wizard.
3. On the first panel of the Tivoli Storage Manager Client Configuration Wizard, make sure **Help me configure the TSM Backup Archive Client** is selected.
4. Follow the prompts to finish creating your options file. Click the **Help** button if you need assistance completing the wizard. For systems other than Solaris, the wizard places an options file named dsm.opt in the same directory where you installed the client files. For Solaris systems, the configuration files are located in the /usr/bin directory, and links to the configuration files /usr/bin/dsm.opt and /usr/bin/dsm.sys are placed in the installation directory. This placement prevents the configuration files from being removed when the client is completely uninstalled.

Note: You need to know a few technical details, such as the TCP/IP address of the backup server you will be connecting to and the port it listens on, in order to successfully create your client options file. If you do not know this information, contact your Storage Manager server administrator.

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

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

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** → **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 /home /mfg /planning /mrketing /mgmt
```

- Enter one or more blank spaces between parameters.
- The maximum number of characters for a file name is 256. The maximum combined length of the file name and path name is 1024 characters.

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 458 for more information.

Setting environment variables

Setting language environment variables

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

You can use the **LANG** environment variable to specify the language for the AIX®, HP-UX, Linux®, and Solaris clients. For all other UNIX clients, only English (United States) is available.

Tivoli Storage Manager supports the following language locales:

Table 18. Supported language locales

Languages	AIX	HP-UX	Solaris	All Linux clients
English (United States)	en_US	en_US	en_US, en_US.ISO8859-1	en_US
Simplified Chinese	zh_CN	zh_CN	zh, zh_CN.EUC	zh_CN,
Traditional Chinese	zh_TW, Zh_TW.BIG5	zh_TW.eucTW	Traditional Chinese: zh_TW, zh_TW.EUC Traditional Chinese BIG5: zh_TW.BIG5	zh_TW.big5
Czech	cs_CZ	cs_CZ	cs_CZ	cs_CZ
French (Standard)	fr_FR	fr_FR	fr, fr_FR	fr_FR
German (Standard)	de_DE	de_DE	de, de_DE	de_DE
Hungarian	hu_HU	hu_HU	hu, hu_HU	hu_HU
Italian (Standard)	it_IT	it_IT	it, it_IT	it_IT
Japanese	ja_JP, Ja_JP	ja_JP.eucJP	ja, ja_JP.eucJP	ja_JP.eucJP
Korean	ko_KR	ko_KR	ko, ko_KR.EUC	ko_KR
Polish	pl_PL	pl_PL	pl, pl_PL	pl_PL
Portuguese (Brazil)	pt_BR	pt_BR	pt, pt_BR	pt_BR
Russian	ru_RU	ru_RU	ru, ru_RU	ru_RU
Spanish	es_ES	es_ES	es, es_ES	es_ES

Note:

- Traditional Chinese BIG5 has locale Zh_TW on Linux, while Zh_TW is used on other platforms for eucTW.
- Not all UNIX shells are able to show multi-byte characters.

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

```
export LANG=fr_FR
```

You also need to export the LC_ALL environment variable.

Notes:

1. To display the Tivoli Storage Manager help browser menus in the language of your current locale, insure that the NLSPATH environment variable in the /etc/profile file contains the following path:
NLSPATH=/usr/dt/lib/nls/msg/%L/%N.cat:\$NLSPATH
export NLSPATH
2. For double-byte languages (Chinese/Korean/Japanese) on UNIX and Linux, we recommend that you change the language via the **Options** button on the CDE login screen, rather than exporting LANG environment variable via the command line.

If the **LANG** environment variable is set to **C**, POSIX (limiting the valid characters to those with ASCII codes less than 128), or other values with limitations for valid characters, the backup-archive client skips files which have file names containing invalid characters with ASCII codes higher than 127.

If you are using a single-byte character set (SBCS) like English as your language environment, all file names are valid and backed up by the backup-archive client. Multi-byte characters are interpreted as a set of single bytes all containing valid

characters. If you are using multi-byte character sets (MBCS) as your language environment, the backup-archive client backs up file names that consist of valid characters in the current environment.

For example, a file name consisting of Japanese characters might contain invalid multi-byte characters if the current language environment is a Chinese character set. Files containing invalid multi-byte characters are not backed up and are not shown by the graphical user interface. If such files are found during backup, the `dsmerror.log` file will list the skipped files.

When using the backup-archive client scheduling mode to back up a whole system, it is strongly recommended to set the `LC_CTYPE` environment variable to `en_US` (or some other SBCS language) to avoid skipped files.

Setting processing environment variables

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 options file for the backup-archive client, the administrative client, or both.

Note: You can also specify an alternate client options file for the command line client (not the administrative client) using the `optfile` option. See “Optfile” on page 303 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`, the resource files, and the `dsm.sys` file reside. You *cannot* specify the root (`/`) directory for `DSM_DIR`. If `DSM_DIR` is not set, the executables are expected in the default installation directory.

When you request an image backup, image restore, NAS backup, or NAS restore, Tivoli Storage Manager uses the `DSM_DIR` environment variable to locate the corresponding plug-in library. If `DSM_DIR` is not set, the client will look for the plug-in library in the following directories:

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

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

DSM_CONFIG

Specifies the fully-qualified path and file name of the client user options file for users who create their own personalized options file. The root directory is *not* a valid value for `DSM_CONFIG`. If the root directory is specified, the log will be written to the default directory. Specify a directory other than the root directory. If `DSM_CONFIG` is *not* set, or the `-optfile` option is not used, the client user options file is expected to satisfy both of these requirements:

1. The options file must be named `dsm.opt`.

2. 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.

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.

Table 19 shows the default log directory for the various platforms.

Table 19. Platforms and default log directories

Platform	Client default log directory	API default log directory
Linux	current working directory	current working directory
Solaris	current working directory	current working directory
HP	current working directory	current working directory
AIX	current working directory	current working directory
z/OS	/var/ibm/tivoli/tsm/client/ba/log	/var/ibm/tivoli/tsm/client/api/bin/log

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

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.

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

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

Setting Bourne and Korn shell variables

Enter the environment variables in the .profile file (Korn shell) or .bash_profile file (Bourne shell) in your \$HOME directory. For example:

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

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

Setting C shell variables

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

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

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

Setting API environment variables

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

DSMI_DIR

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

DSMI_CONFIG

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

DSMI_LOG

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

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

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

Configuring the Web client

You can use the command line to configure the Web client.

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

1. Ensure that you specify *passwordaccess generate* in the dsm.sys file. For more information on *passwordaccess*, see “Passwordaccess” on page 305.
2. 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.

3. Start the client acceptor daemon (CAD) by entering:

```
dsmcad
```

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

The options applicable only to the **dsmcad** program are *httpport*, *managedservices*, and *webports*. Other options such as *optfile* or *errorlogname* can also be used. You can use the *managedservices* option to specify whether the Tivoli Storage Manager Client Acceptor daemon (CAD) also manages the Tivoli Storage Manager scheduler. See Chapter 9, “Using processing options,” on page 169 for more information about these options.

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

4. Add *MANAGEDServices* webclient to *dsm.opt* and *dsm.sys*.
5. To access the Web client, enter the following URL from any supported browser:
`http://your_machine_name:1581`
where *your_machine_name* is the host name of the machine running the Web client.

Port 1581 is the default port number. You can set a different port number using the *httpport* option. See “Httpport” on page 262 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) Java JRE 1.4.x (where $x \geq 1$ or higher). See “Starting a Web client session” on page 72 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 182 for more information.

See Chapter 7, “Automating tasks,” on page 151 for more information. See Appendix A, “Using the Tivoli Storage Manager central scheduler,” on page 499 for supplemental information about the Tivoli Storage Manager central scheduler.

The Tivoli Storage Manager Client Acceptor daemon (CAD) can manage the scheduler. In this case, the CAD serves as an external timer for the scheduler.

When the scheduler is started, it queries the server for the next scheduled event. The event is either executed immediately or the scheduler exits. The CAD 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 CAD management. It is recommended that you use the Client Acceptor daemon to manage the client scheduler.

Attention: You cannot use the `dsmcad` for scheduling when you set the `sessioninitiation` option to `serveronly`. Refer to “Sessioninitiation” on page 344 for more information.

Use the `managedservices` option in your client system options file (`dsm.sys`) to specify whether the CAD manages the scheduler. See “Managedservices” on page 286 for more information.

Perform the following steps to configure the CAD to manage the client scheduler:

1. Install the Web client. See “Configuring the Web client” on page 50 for more information.
2. Install the Scheduler. See “Starting the client scheduler” for more information.
3. From the Tivoli Storage Manager GUI, select **Edit** → **Preferences**. Then select the **Web Client** category. Check the **Schedule** option in the **ManagedServices** options section. If you wish to run the Web client also, check the **Both** option.
4. Start the Client Acceptor. See “Configuring the Web client” on page 50 for more information.

Notes:

1. 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 182.
2. See Chapter 7, “Automating tasks,” on page 151 for information about the following tasks:
 - Starting the client scheduler
 - 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

To start the client scheduler on your client node and connect to the server schedule, enter the following command:

```
dsmc schedule
```

If the Tivoli Storage Manager client executable directory is not in your `PATH` environment variable, change to the installation directory and enter the following command:

```
./dsmc schedule
```

For z/OS, enter the following path and command:

```
/usr/lpp/Tivoli/tsm/client/ba/bin/dsmc schedule
```

When you start the client scheduler, it runs continuously until you close the window, end the process, or log off your system.

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

```
nohup dsmc schedule 2> /dev/null &
```

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

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

- **For non-z/OS UNIX**, add the following entry to the */etc/inittab* file:

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

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

For z/OS UNIX:

1. Create a shell script called */tivoli/tsm/client/ba/bin/rundsmc* which contains the following entries:

```
cd /usr/lpp/Tivoli/tsm/client/ba/bin
sleep 60
./dsmc schedule
```

This prevents the creation of two jobs with the same name and enables automatic shutdown. You might need to customize the time for your system.

2. Add the following entries in the */etc/rc* file to set environment variables to retrieve the *servername* and *nodename* options from *dsm.sys* and to start the client scheduler, as follows:

```
# Set environment variables to retrieve the servername and
# nodename options from dsm.sys.
export DSM_DIR=/tivoli/tsm/client/ba/bin
export DSM_CONFIG=/tivoli/tsm/client/ba/bin/dsm.opt
# Start the TSM Client scheduler and redirect outputs to
# schedule.out instead of the /etc/log file.
_BPX_JOBNAME='ADSMCLNT' /tivoli/tsm/client/ba/bin/rundsmc
1>/tivoli/tsm/client/ba/bin/schedule.out 2>&1 &
```

Note: Enter the *_BPX_JOBNAME* entry on a single line in the */etc/rc* file.

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

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

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

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

Use the Client Acceptor daemon to manage the client scheduler. See “Configuring the client scheduler” on page 51 for more information.

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 366 for more information. This will allow Tivoli Storage Manager scheduler communications in both *polling* and *prompted* mode, CAD-managed schedulers, and regular backup-archive client operations.

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

HTTP port

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

To use the administrative Web interface for a server across a firewall, the Tivoli Storage Manager administrator must open the HTTP port for the server using the *httpport* option in the server options file. The default HTTP port is 1580.

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 392 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 360 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. If *sessioninitiation* is set to *serveronly*, the value for the *tcpclientaddress* client option must be the same as the value for the *HLAddress* option of the **update node** or **register node** server command. The value for the *tcpclientport* client option must be the same as the value for the *LLAddress* option of the **update node** or **register node** server command.
2. Using the *sessioninitiation* option requires a Tivoli Storage Manager Version 5.2 or higher server.
3. If you set the *sessioninitiation* option to *serveronly*, with the exception of CAD-managed schedulers, the command line client, backup-archive client Java GUI, and Web client GUI 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*.
4. When configuring the Tivoli Storage Manager scheduler on a client machine 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 machine to the server machine. See 54 for more information about setting up this communication. After the encryption key is set, you can use server-initiated sessions to back up the files using Tivoli Storage Manager encryption.

If you set the *sessioninitiation* option to *client*, the client 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 can be used to prompt the client to connect to the server.

See “Sessioninitiation” on page 344 for more information about the *sessioninitiation* option.

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

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

Configuring the journal daemon

Journal-based backup is supported on AIX clients for JFS, JFS2 and VxFS file system types. If you install the *journal daemon* and it is running, then by default the **incremental** command will automatically perform a journal-based backup on selected file systems which are being monitored by the journal engine service. See “Journal-based backup on AIX” on page 89 for more information.

Journal-Based Backup is enabled by installing and configuring the Tivoli Storage Manager daemon. To install the Tivoli Storage Manager daemon, you need to customize the file named *tsmjbbd.ini.smp*, which is placed in the install directory, and save it as *tsmjbbd.ini*. The install process modified */etc/inittab* to start the journal daemon on a reboot. You can also start the daemon on the command line by entering the name of the executable - *tsmjbbd*.

Journal-based backup is enabled by installing and configuring the Tivoli Storage Manager journal daemon. The Tivoli Storage Manager journal daemon can be configured by editing the journal daemon configuration sample file, *tsmjbbd.ini.smp*, and saving it as *tsmjbbd.ini*. Both the configuration sample file and the saved file should be in the default install directory.

The journal configuration file, (*tsmjbbd.ini*), needs as a minimum a list of the file systems to monitor. These two lines are sufficient:

```
[JournaledFilesystemSettings]
JournaledFileSystems=/home
```

Enter a *JournalDir* setting and a *JournalExcludeList* for the database files (see the sample file). An *Errorlog* setting will also explicitly name the errorlog file. After the configuration file is created, start the journal daemon using the script file: */usr/tivoli/tsm/client/ba/bin/rc.tsmjbb*. The journal will write initialization information to the errorlog. When you are satisfied that the journal is working correctly, you should run the script file, */usr/tivoli/tsm/client/ba/bin/jbbinittab*. Running the script file will make entries in */etc/inittab*, so that the journal will begin running when you restart your system.

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.

If you plan to use a Web client, you must have an administrative user ID with system privilege, policy privilege, client access authority, or client owner authority. When a new node is registered, an administrative user ID is automatically created for the node. By default, this node has client owner authority.

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 value returned by the **hostname** command 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, defined by your Tivoli Storage Manager administrator, 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.

Associating your client node with a host system (optional)

The GUIDs help to uniquely identify a particular machine (for reporting purposes), regardless of the following:

- How many node names are used on the machine
- Which network adapter is used to connect to the Tivoli Storage Manager server
- Which nodes are connected to the Tivoli Storage Manager servers

For example, if you use nodes GORDON, DONNA, and DAGORDON to connect to a Tivoli Storage Manager server from your desktop machine, all three nodes will have the same GUID. Similarly, if nodes GORDON, DONNA, and DAGORDON connect to multiple Tivoli Storage Manager servers, each server will show the same GUID for these nodes.

Note: The GUID is available only on HP-UX, Linux x86/x86_64, and Solaris. You must be a root user to run `tivguid`.

When you install the Tivoli software, the `tivguid` program is run to generate a GUID which is stored in the `/etc/tivoli` directory. The GUID for a client node on the server can change if the host system machine is corrupted, if the file entry is lost, or if a user uses the same node name from different host systems. You can perform the following functions from the command line:

- Create a new GUID
- View the current GUID
- Write a specific value
- Create another GUID even if one exists.

Table 20 describes the GUID functions and the associated commands.

Table 20. GUID commands

Function	Enter on the command line:
Create and store a new GUID on the host if one does not exist. If a GUID already exists, the current value is displayed.	<code>tivguid -Create</code>
Display help for the <code>tivguid</code> commands.	<code>tivguid -Help</code>
Return the value of the current GUID.	<code>tivguid -Show</code>
Write the GUID that is specified in the <code>-GUID</code> option to the file. For example, <code>-Write GUID = 'string'</code> uses the value in 'string' rather than creating a new GUID. The string must be a valid Tivoli GUID (32 hexadecimal values).	<code>tivguid -Write -guid=38.70.92.a1.9a.93.11.d6.a2.f9.00.04.ac.dd.76.38</code>
This function is useful in the following cases: <ul style="list-style-type: none">• If the Tivoli GUID is corrupted you can use the administrative client to query the server for the value using the <code>q node nodename f=d</code> command, and set that value on the current machine.• If you want to set up multiple physical machines with the same guid (for example on cluster).	
Create a new GUID even if one exists.	<code>tivguid -Write -New</code>

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

Authorized User

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 157.

You can create an include-exclude list to exclude a specific file or groups of files from backup services, and to assign specific management classes to files. Tivoli Storage Manager backs up any file that is not explicitly excluded. You should exclude Tivoli Storage Manager client directories from backup services. You can use the **query inclexcl** command to display a list of include and exclude statements in the order they are examined when determining whether an object is to be included.

Specify the include-exclude list in your `dsm.sys` file. If you define more than one server in your `dsm.sys` file, each server must have its own include-exclude list. This list can also contain include-exclude statements obtained from the include-exclude files you specify with the *inclexcl* option.

When the client processes include-exclude statements, the include-exclude statements within the include-exclude file are placed at the position occupied by the *inclexcl* option in `dsm.sys`, in the same order, and processed accordingly.

See “Inclexcl” on page 268 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 **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.

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 60. Tivoli Storage Manager evaluates all *exclude.fs* and *exclude.dir* statements *first* (regardless of their position within the include-exclude list), and removes the excluded file spaces, directories, and files from the list of objects available for processing. All other include-exclude statements are processed from the bottom of the list up. Therefore, it is important to enter all your

include-exclude statements in the proper order. For example, in the following include-exclude list the `includefile.txt` file *is not* backed up:

```
include /home/usr/includefile.txt
exclude /home/usr/.../*
```

However, in the following include-exclude list the `includefile.txt` file *is* backed up:

```
exclude /home/usr/.../*
include /home/usr/includefile.txt
```

4. Save the file and close it.
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 file spaces and directories

Use *exclude.fs* and *exclude.dir* statements to exclude file spaces and 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 file spaces, 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 21. Options for excluding file spaces and 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 246 for more information.	246
<i>exclude.dir</i>	Excludes a directory, its files, and all its subdirectories and their files from backup processing. For example, the statement <code>exclude.dir /test/dan/data1</code> excludes the <code>/test/dan/data1</code> 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.	246

Notes:

1. If you define an exclude statement without using a drive letter, such as `exclude.dir dirname`, this will exclude from processing any directory named `dirname` on any drive.

Controlling backup, archive, and image processing

After Tivoli Storage Manager evaluates all *exclude.fs*, *exclude.dir*, *exclude.attribute.symlink*, and *include.attribute.symlink* statements, the following options are evaluated against the remaining list of objects available for processing.

Table 22. Options for controlling backup, archive, and image 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 and space management services (if the HSM client is installed). The <i>exclude.backup</i> option only excludes files from normal backup, but not from HSM.	246
<i>include</i> <i>include.backup</i> <i>include.file</i>	Use these options to include files or assign management classes for backup processing.	269
<i>include.fs</i>	Controls how Tivoli Storage Manager processes your file space for incremental backups.	269
Archive processing		
<i>exclude.archive</i>	Excludes a file or group of files from archive services.	246
<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.	269
Image processing		
<i>exclude.image</i>	Excludes mounted file systems and raw logical volumes that match the specified pattern from full image backup operations. Incremental image backup operations are unaffected by <i>exclude.image</i> . This option is valid for AIX, HP-UX, Solaris, and all Linux clients..	246
<i>exclude.fs.nas</i>	Excludes file systems on the NAS filer from an image backup when used with the backup nas command. If you do not specify a NAS node name, the file system identified applies to all NAS filers. The backup nas command ignores all other exclude statements including <i>exclude.fs</i> and <i>exclude.dir</i> statements. This option is for AIX and Solaris clients <i>only</i> .	246
<i>include.image</i>	Includes a file space or logical volume, assigns a management class, or allows you to assign one of several image backup processing options to a specific logical volume when used with the backup image command. The backup image command ignores all other include options. This option is valid for AIX, HP-UX, Solaris, Linux x86/x86_64, and Linux on POWER only.	269
<i>include.fs.nas</i>	Use the <i>include.fs.nas</i> option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system image backup, using the <i>toc</i> option with the <i>include.fs.nas</i> option in your <i>dsm.sys</i> file. See "Toc" on page 371 for more information. This option is valid for AIX and Solaris clients <i>only</i> .	269

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 23. 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 400.

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 23. Wildcard and other special characters

Character	Function
?	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: <ul style="list-style-type: none">• The pattern <code>ab?</code>, matches <code>abc</code>, but does not match <code>ab</code>, <code>abab</code>, or <code>abzzz</code>.• The pattern <code>ab?rs</code>, matches <code>abfrs</code>, but does not match <code>abrs</code>, or <code>abllrs</code>.• The pattern <code>ab?ef?rs</code>, matches <code>abdefjrs</code>, but does not match <code>abefrs</code>, <code>abdefrs</code>, or <code>abefjrs</code>.• The pattern <code>ab??rs</code>, matches <code>abcdrs</code>, <code>abzzrs</code>, but does not match <code>abrs</code>, <code>abjrs</code>, or <code>abkkrs</code>.
*	The match-all character. For example: <ul style="list-style-type: none">• The pattern <code>ab*</code>, matches <code>ab</code>, <code>abb</code>, <code>abxxx</code>, but does not match <code>a</code>, <code>b</code>, <code>aa</code>, <code>bb</code>.• The pattern <code>ab*rs</code>, matches <code>abrs</code>, <code>abtrs</code>, <code>abrsrs</code>, but does not match <code>ars</code>, or <code>aabrs</code>, <code>abrss</code>.• The pattern <code>ab*ef*rs</code>, matches <code>abefrs</code>, <code>abefghrs</code>, but does not match <code>abefr</code>, <code>abers</code>.• The pattern <code>abcd.*</code>, matches <code>abcd.c</code>, <code>abcd.txt</code>, but does not match <code>abcd</code>, <code>abcdc</code>, or <code>abcdtxt</code>.
/...	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: <code>xxx[abc]</code> matches <code>xxxa</code> , <code>xxxb</code> , or <code>xxxc</code> .
-	The character-class range includes characters from the first character to the last character specified. For example: <code>xxx[a-z]</code> matches <code>xxxa</code> , <code>xxxb</code> , <code>xxxc</code> , ... <code>xxxz</code> .
\	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 <code>'</code> in a character class, enter <code>[...\...]</code> . The escape character removes the usual meaning of <code>'</code> 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 24 on page 63 contains examples of ways you might use wildcard characters with *include* and *exclude* patterns. For more information about using the *exclude.backup* option, see “Exclude options” on page 246.

Note: In the `dsm.sys` file, the *include* and *exclude* options do not work with symbolic links to directories. For example, do not use `/u` in your include or exclude statements because `/u` is a symbolic link to the `/home` directory. Instead of entering:

```
include /u/tmp/save.fil
```


enter:

```
include /home/tmp/save.fil
```

However, the *exclude* option does work with symbolic links to directories when you enter a backup command with the absolute path that contains the symbolic link.

Table 24. Using wildcard characters with include and exclude patterns

Task	Pattern
Exclude all files that end with .doc, except those found in the Documents directory of aleko's home directory.	EXCLUDE /.../*.doc INCLUDE "/home/aleko/Documents/*.doc"
Exclude all files during backup with an extension of bak, except those found on the /usr file system in the dev directory.	exclude /.../*.bak include /usr/dev/*.bak
Exclude all files and directories under any Documents directory that might exist, except for the Current file of user aleko.	EXCLUDE /.../Documents/.../* INCLUDE "/home/aleko/Documents/Current"
Exclude all files and directories in any tmp directory that might exist, except for the file /home/tmp/save.fil. 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 /fs1, /fs2, /fs3 and /fs4 file systems.	EXCLUDE /fs[1-4]/.../*.cpp
Exclude the .cpp files found in the /fs2/source directory.	EXCLUDE /fs2/source/*.cpp
Exclude any .o file in any directory on the /usr1, /usr2, and /usr3 file systems.	exclude /usr[1-3]/.../*.o
Exclude the .o files found in the root directory in the usr2 file system only.	exclude /usr2/*.o
Exclude any file that resides under the tmp directory found in any file system.	exclude /.../tmp/.../*
Exclude the entire directory structure /var/spool from all processing.	exclude.dir /var/spool
Exclude a single file system from backup processing.	exclude.fs /fs1 exclude.fs home:
Exclude all file systems mounted anywhere in the /test/myfs/fs01 and /test/myfs/fs02 directory tree from backup processing.	exclude.fs /test/myfs/fs01/.../* exclude.fs /test/myfs/fs02/*
Exclude the /home/mydir/test1 directory and any files and subdirectories under it.	exclude.dir /home/mydir/test1
Exclude all directories under the /home/mydir directory with names beginning with test.	exclude.dir /home/mydir/test*
Exclude all directories directly under the /mydir directory with names beginning with test, on any file system.	exclude.dir /.../mydir/test*
Exclude the raw logical volume from image backup.	exclude.image /dev/hd0
Exclude all symbolic links 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 processing

After Tivoli Storage Manager evaluates all *exclude.fs* and *exclude.dir* statements and removes the excluded file spaces and directories, Tivoli Storage Manager evaluates any include-exclude statements for controlling symbolic link processing (*exclude.attribute.symlink* and *include.attribute.symlink*) against the remaining list of objects available for processing. Table 25 defines options for controlling symbolic link processing.

Table 25. Options for controlling symbolic link processing

Option	Description	Page
<i>exclude.attribute.symlink</i>	Excludes a file or a group of files that are symbolic links from backup processing only.	246
<i>include.attribute.symlink</i>	Includes a file or a group of files that are symbolic links within broad group of excluded files for backup processing only.	269

Controlling compression and encryption processing

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

Table 26. 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.	246
<i>include.compression</i>	Includes files for compression processing if <i>compression=yes</i> is specified. This option applies to backups and archives.	269
Encryption processing		
<i>exclude.encrypt</i>	Excludes files from encryption processing.	246
<i>include.encrypt</i>	Includes files for encryption processing.	269
	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.	

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 on a .o file, and choose "View File Details" from the popup menu. The dialog shows that these files are included. Click on 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 439 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.fs* and *exclude.dir* statements *first*, and removes the excluded file spaces, directories, and files from the list of objects available for processing. See "Excluding file spaces and directories" on page 60 and "Exclude options" on page 246 for more information about the *exclude.fs* and *exclude.dir* options.

After evaluating all *exclude.fs* and *exclude.dir* statements, Tivoli Storage Manager evaluates the include-exclude statements for controlling symbolic link 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 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 169 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 incl excl** command. See "Query Incl excl" on page 455 for more information.

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

1. Files are checked; directories are *only* checked if the *exclude.dir* option is specified.
2. **File names are compared to the patterns in the include-exclude list from the bottom up.** When a match is found, the processing stops and checks whether the option is *include* or *exclude*. If the option is *include*, the file is backed up. If the option is *exclude*, the file *is not* backed up.

Note: 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 you defined the following statements for the *include* and *exclude* options:

```
exclude *.o
include /home/foo/.../*.o
exclude /home/foo/junk/*.o
```

The file being processed is: `/home/foo/dev/test.o`. Processing follows these steps:

1. Rule 3 (the last statement defined) is checked *first* because of bottom-up processing. The pattern `/home/foo/junk/*.o` does not match the file name that is being processed.
2. Processing moves to Rule 2 and checks. This time, pattern `/home/foo/.../*.o` matches the file name that is being processed. Processing stops, the option is checked, and it is *include*.
3. File `/home/foo/dev/test.o` is backed up.

Example 2

Assume that you defined the following statements for the *include* and *exclude* options:

```
exclude *.obj
include /home/foo/.../*.o
exclude /home/foo/junk/*.o
```

The file being processed is: `/home/widg/copyit.txt`. Processing follows these steps:

1. Rule 3 is checked and finds no match.
2. Rule 2 is checked and finds no match.
3. Rule 1 is checked and finds no match.
4. Because a match is not found, file `/home/widg/copyit.txt` is implicitly included and backed up.

Example 3

Assume that you defined the following statements for the *include* and *exclude* options:

```
exclude /.../*.o
include /home/foo/.../*.o
exclude /home/foo/junk/*.o
```

The current file being processed is: /home/lib/objs/printf.o. Processing follows these steps:

1. Rule 3 is checked and finds no match.
2. Rule 2 is checked and finds no match.
3. Rule 1 is checked and a match is found.
4. Processing stops, the option is checked, and it is *exclude*.
5. File /home/lib/objs/printf.o is not backed up.

Example 4

Assume that you defined the following statements for the *include* and *exclude* options:

```
exclude.attribute.symlink /.../*
exclude /.../*.o
include /home/foo/.../*.o
exclude /home/foo/junk/*.o
```

The current file being processed is: /home/lib/objs/printf.o. Processing follows these steps:

1. The *exclude.attribute.symlink* statement is checked first. If the printf.o file is a symbolic link it 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 /home/lib/objs/printf.o is not backed up.

Chapter 3. Getting started

This chapter includes instructions for the following tasks:

Task	Page
Starting a GUI session	70
Starting a command line session	71
Starting a Web client session	72
Starting the client scheduler automatically	75
Changing your password	75
Sorting file lists using the Tivoli Storage Manager Java GUI	76
Displaying online help	76
Ending a session	77

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 machine, 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 machine, you must have an administrative user ID with client access authority over the node name for the remote client machine. 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

machine, specify the *revokeremoteaccess* option in your client user options file (dsm.opt). See “Revokeremoteaccess” on page 330 for more information.

Client access authority only allows Tivoli Storage Manager administrators to back up and restore files on remote machines. They do not have physical access to the data. That is, they cannot restore the data belonging to the remote machine to their own machines. To restore data belonging to a remote machine to your own machine, 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

The Tivoli Storage Manager GUI must be run from the X Window System. If you see the Tivoli Storage Manager icon on your desktop, Tivoli Storage Manager is already running. Double-click the icon to open the Tivoli Storage Manager window. If the Tivoli Storage Manager icon does not appear on your desktop, you should start Tivoli Storage Manager using the **dsmj** (starts the Java GUI) command. Tivoli Storage Manager can run as either a *foreground* or *background* process.

To run the Java GUI in the foreground, enter:

```
dsmj
```

To run the Java GUI in the background, enter:

```
dsmj &
```

Tivoli Storage Manager locates the *client user options* file (dsm.opt) and starts with the options specified in that file. See Chapter 2, “Configuring Tivoli Storage Manager,” on page 41 for more information about the client user 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 75.

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 machine), 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 `/usr/bin` directory contains a symbolic link to the Tivoli Storage Manager executable, and all DSM environment variables are set, you can enter the **dsmc** command from any directory; otherwise, enter the fully qualified path.

- On the command line enter **dsmc** followed by the command, to execute a single command (*batch mode*). If the `/usr/bin` (Solaris) or `opt/bin` directory contains a symbolic link to the Tivoli Storage Manager installation directory, you can enter the **dsmc** command from any directory. Otherwise you can enter the fully qualified name.
- 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 `/home/proj1/file1.txt`, enter the command:

```
dsmc archive /home/proj1/file1.txt
```

To issue the **incremental** command, enter the following at the command 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 305 for more information.

When you type your password and press Enter, 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 **Enter**, 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 `archive /home/proj1/file1.txt`, enter the command:

```
archive /home/proj1/file1.txt
```

Note: In interactive mode, you cannot enter a file specification that contains national language characters. If a command contains national characters, process the command in batch mode by preceding the command with the executable program name, **dsmc**.

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 Enter, the password is not displayed on your screen.

See “Options handling in interactive mode” on page 397 for a discussion of how options are handled in interactive mode.

See Chapter 10, “Using commands,” on page 393 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
```

Tivoli Storage Manager Backup-Archive Java GUI (**dsmj**) now accepts command line parameters, such as the Java `-X` options. Because of this, you can also now modify the Java Heap Size. For example:

```
dsmj -Xmx512M
```

Starting a Web client session

After installing the Web client on your workstation (see “Configuring the Web client” on page 50) you can use the Web client to perform backup, archive, restore, and retrieve operations from any browser that is running Java Runtime Environment (JRE) 1.4.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:

- Microsoft Internet Explorer 5.0 or higher with JRE 1.4.x, where *x* is greater than or equal to 1.

- Mozilla 1.4 or higher with JRE 1.4.x , where x is greater than or equal to 1 (AIX and Linux clients only)

To run the Web Client from Mozilla browsers, **Enable JavaScript™** must be checked. This setting is enabled by default, but to verify it:

1. Open Mozilla's **Edit** menu and select **Preferences**.
2. In the Preferences dialog under Category, select **Advanced**, then **Scripts & Plugins**.
3. Ensure there is a check mark next to **Enable JavaScript for Navigator**.
1. Open the **Tools** menu and select **Internet Options**
2. From the Internet Options dialog, select the **Security** tab.
3. Click the Web content zone in which you will be using the Storage Manager Web client and then click the **Custom Level** button.
4. In the Security Settings dialog, ensure that **Enable** is selected under the **Scripting of Java applets** setting.

If your browser does not have the correct JRE level, the Web client will notify you and if possible, will try to automatically install the correct JRE for you.

- The Web Client will let you know if your browser does not have the correct JRE level. For Microsoft Internet Explorer on Windows platforms, the Web client will automatically download and install JRE 1.4.x, where x is greater than or equal to 1
- On all platforms except for AIX and Internet Explorer browsers on Windows platforms, the Web client will tolerate the JRE at level 1.3.0 and will not prompt to install a later version to avoid installing more than one JRE on a machine. If you experience problems with the Web Client using JRE 1.3.0, you should install JRE 1.4.x, where x is greater than or equal to 1.

If you have JRE 1.3.0 or lower installed, consider the following:

- Uninstall any existing JREs before installing the new JRE.
- Do not install the new JRE in the same directory as the existing JRE. This can cause JRE-related problems or crashes.

You can download and install JRE 1.4.x, where x is greater than or equal to 1 (unless otherwise noted) manually from the appropriate URL:

- For Windows, Solaris, Linux:
<http://java.sun.com/javase/index.jsp>
- AIX:
<http://www.ibm.com/developerworks/java/jdk/aix/index.html>
- HP-UX:
http://www.hp.com/products1/unix/java/?jumpid=reg_R1002_USEN

Note: Proxy server users: The JRE 1.4.x , where x is greater than or equal to 1 might return a security exception or a class not found exception if the Storage Manager Web Client attempts to open a TCP/IP socket to a socks server to communicate with the Storage Manager Remote Client Agent. To avoid this, you can use one of the following methods to bypass a proxy server, allowing the Web client to establish a direct connection to the Agent machine:

- **Change your Java plug-in settings:**

For Windows:

1. Open the Windows **Start** menu and select **Settings** → **Control Panel**.

2. In the Control Panel, double-click **Java Plugin**.
3. In the Java Plug-In Control Panel, select the **Proxies** tab and uncheck the **Use browser settings** check box.

For UNIX and Linux:

1. Change directory to the installation directory of your JRE, and then change directory to the bin directory.
 2. Run the JavaPluginControlPanel executable and click the **Proxies** tab.
 3. Uncheck **Use browser settings**.
- **Change your browser settings to enable a direct connection to the Internet:**
 - For Mozilla: Open the **Edit** menu and select **Preferences**. Under **Category**, expand the **Advanced** section, select **Proxies**, and click **Direct connection to the Internet**.
 - For Internet Explorer: Open the **Tools** menu and select **Internet Options....** Select the **Connections** tab, and click the **LAN Settings** button. Uncheck the **Use a proxy server** check box.

Additional information about running Swing applets can be found in Sun's Java Tutorial:

<http://java.sun.com/docs/books/tutorial/uiswing/index.html>

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. However, Tivoli Storage Manager backup and restore activities running in conjunction with a NAS box will continue after disconnect.

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 69 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 330 for more information about the *revokeremoteaccess* option. See "Tivoli Storage Manager client authentication" on page 69 for more information about access authorities.

Starting the client scheduler automatically

Root User

You can add a command to the inittab file to start the client scheduler when the system boots up without any users logged in. 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 151 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 52 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. From the main window, open the **Utilities** menu and select **Change password**.
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:

```
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 304 for additional password information.

The following are additional password information sources:

- “Starting the client scheduler automatically”
- “Starting: Additional considerations” on page 72
- “Password” on page 304
- “Set Password” on page 495

Sorting file lists using the Tivoli Storage Manager Java GUI

Table 27. 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 126).
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 267 and “Pick” on page 308 for more information.

Displaying online help

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

- **From the Tivoli Storage Manager Java GUI:**
 - Open the **Help** menu.
 - Click the **Help** button in the current window.
- **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 429 for more information about the **Help** command.

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 **Exit**.
 - Open the **System** menu and select **Close**.
 - **For the Web client:** Open a different URL or close the browser.
- **From the DSMC command line:** In batch mode, each **dsmc** command you enter is a complete session. Tivoli Storage Manager ends the session when it finishes processing the command. To end an interactive session, enter **quit** at the **tsm>** prompt.

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://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://rz.uni-karlsruhe.de) (mirror - Germany)
[ftp.wu-wien.ac.at](ftp://wu-wien.ac.at) (mirror - Austria)
[ftp.cac.psu.edu](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 xvi for product support information.

Chapter 4. Backing up your data

Use Tivoli Storage Manager to maintain a set of backup versions of your file system or raw logical volumes on your workstation. You can recover these older file versions in the event your current files are lost or damaged. This chapter covers different ways to back up files, how to restore these files, and the difference between backing up and archiving files.

All client backup procedures in this chapter also apply to the Web client, except the following:

- Estimate
- Preferences editor

See “Starting a Web client session” on page 72 for information on starting the Web client.

The following table identifies tasks described in this chapter:

Table 28. Backup: Primary tasks

Task	Page
Planning your backups	79
Pre-backup considerations	81
Performing an incremental, selective, or incremental-by-date backup	87
Deleting backup data	95
Group backup: Backing up files from one or more file spaces	96
Performing an image backup	102
Backing up NAS file systems	108
Backing up the WebSphere Application Server (WAS)	112

Planning your backups

If you are a first-time user, or if you only back up files occasionally, you might want to use Table 29 as a checklist of preliminary steps to consider before performing a backup.

Table 29. Planning your backups

—	Decide whether you want to back up files or archive them. See “When to back up and when to archive files” on page 80 for more information.
—	See “Pre-backup considerations” on page 81 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 83 for more information.

Table 29. Planning your backups (continued)

—	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 87• “Group backup: Backing up files from one or more file spaces” on page 96• “Performing an image backup” on page 102• “Backing up NAS file systems” on page 108• “Backing up the WebSphere Application Server (WAS)” on page 112
—	If you are performing an image backup, ensure that you have accounted for unique considerations. See “Performing an image backup” on page 102 for more information.
—	For further backup considerations, see “Backup: Additional considerations” on page 114.

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 machine

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 130 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 143 for more information about archiving and retrieving files.

Pre-backup considerations

LAN-free data movement

AIX, HP-UX, Linux and Solaris (global zones and non-global zones) clients support LAN-free data movement, which shifts the movement of client data from the communications network to a storage area network (SAN). This decreases the load on the Tivoli Storage Manager server.

The SAN provides a path that allows you to back up, restore, archive, and retrieve data to and from a SAN-attached storage device. Client data moves over the SAN to the storage device via the Tivoli Storage Manager Storage Agent. The Tivoli Storage Manager Storage Agent must be installed on the same system as the client.

LAN-free prerequisites

To enable LAN-free support, you must install and configure the Tivoli Storage Manager Managed System for SAN Storage Agent on the client workstation. For more information, refer to the following:

- *IBM Tivoli Storage Manager for AIX Storage Agent User's Guide, SC32-0129*
- *IBM Tivoli Storage Manager for HP-UX Storage Agent User's Guide, SC32-0130*
- *IBM Tivoli Storage Manager for Linux Storage Agent User's Guide, SC32-0131*
- *IBM Tivoli Storage Manager for Sun Solaris Storage Agent User's Guide, SC32-0132*

The Tivoli Storage Manager Managed System for SAN Storage Agent and documentation, while contained on the client CD, is a separate product, and therefore subject to the fees and licensing required by Tivoli and IBM.

LAN-free options

After installing and configuring the Tivoli Storage Manager Managed System for SAN feature on the client workstation, you can use the following options to enable LAN-free data movement:

enablelanfree

Specifies whether to enable an available LAN-free path to a SAN-attached storage device. See “Enablelanfree” on page 237 for more information.

lanfreecommmethod

Specifies a communication protocol between the client and the Storage Agent. See “Lanfreecommmethod” on page 276 for more information.

lanfreeshmport

Specifies the unique number that is used by the client and the storage agent to identify shared memory area used for communications. See “Lanfreeshmport” on page 278 for more information.

lanfreetcport

Specifies the TCP/IP port number where the Storage Agent is listening. See “Lanfreetcport” on page 279 for more information.

lanfreetcserveraddress

Specifies the TCP/IP address for the Tivoli Storage Manager storage agent.

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*.
- UNIX clients can use the *virtualmountpoint* client option to define multiple virtual mount points within a single file system, each of which can be backed up independently by the Tivoli Storage Manager client.

- If the client option *resourceutilization* is set to a value greater than 4, and there are multiple file systems being backed up, then reducing *resourceutilization* to 4 or lower 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 151 for more information on management classes and how to assign them to files.

For more information on creating an include-exclude options file, see "Creating an include-exclude list (optional root user or authorized user task)" on page 59.

Encrypting data during backup or archive operations

Tivoli Storage Manager Version 5.3 and higher supports 128-bit AES (Advanced Encryption Standard) 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 239 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 269.

Note: AES is not supported on z/OS.

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.

For more information about this option, see “Encryptkey” on page 240.

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.

File system and ACL support

Special file systems contain dynamic information generated by the operating system; they contain no data or files. The UNIX and Linux clients ignores special file systems and their contents. Special file systems include the following:

- the /proc file system on most of the UNIX platforms
- the /dev/fd file system on Solaris
- the /dev/pts on Linux

The Tivoli Storage Manager client is designed and tested to work on specific file system types which are commonly used. For the list of supported file system types, see table Table 30.

Table 30. Supported file systems and ACL support

Platform	File System	ACL Support
AIX	GPFS	yes
	JFS	yes
	JFS2	yes
	VxFS	yes

Table 30. Supported file systems and ACL support (continued)

Platform	File System	ACL Support
HP-UX	HFS	no
	VxFS (JFS Veritas)	yes (V3.3.Layout 4)
Linux x86/x86_64	XFS	yes
	EXT2	yes
	EXT3	yes
	ReiserFS	no
	GPFS	yes
	JFS	no
	VxFS	no
Linux on POWER	XFS	yes
	EXT2	yes
	EXT3	yes
	ReiserFS	no
	JFS	no
Linux zSeries	EXT2	yes
	EXT3	yes
	ReiserFS	no
	JFS	no
Solaris	UFS	yes
	VxFS	yes
	QFS	no

Notes:

- In order to process all other file systems, the *virtualmountpoint* option needs to be used to enable limited support. This includes backup, restore, archive, and retrieve of file data and directory tree structure. No other file system specific attributes, such as the ACL, will be handled. The file system type for such file systems will be set to "UNKNOWN". For example, if the /media/nss/DATA1 file system is not supported by Tivoli Storage Manager, add the following statement to dsm.sys to back up or archive the data in this file system:

```
VIRTUALMOUNTPOINT /media/nss/DATA1
```

This will also include the /usr/home file system in ALL-LOCAL domain.
Please note that this limited support is only available if the file system itself supports some basic POSIX system calls, such as read() or write(), on your system.
- Cross-platform backup and restore is not supported. For example, data backed up by an AIX client will not be available for restore by a Windows client and vice versa.
- Cross file system type restore or retrieve of the ACL information is allowed provided that both the original file system and the destination file system support compatible ACLs. For example, on Solaris, the ACL information backed up from a VxFS file system will be restored to an UFS file system because these file systems support compatible ACLs. The ACL information will not be restored during cross file system type restores or retrieves if the original file system and the destination file system are incompatible with regard to ACL support or if the destination file system does not support ACLs.
- The standalone package LSCqfs 3.5.0 is the only supported version of QFS. In addition, the following restrictions also apply to the QFS file system:
 - Image backup is not supported on QFS file systems.

- The Solaris backup-archive client does not support the combination of QFS and SAM needed to archive files onto tertiary background storage, such as tapes. Instead, it recalls files from tape to disk automatically if it finds migrated files during a backup.
 - A QFS file system contains two hidden system files and a system directory that cannot be backed up. This is acceptable because a backup of these files is not needed. They contain internal data to manage the file system. This data will be automatically excluded from a backup and recreated automatically by the file system itself if a restore of files in that file system is invoked.
5. Incremental, selective, filelist back up, archive, restore, and retrieve processing of the Veritas file system and its ACLs on AIX are supported. Restore of a Veritas volume on a Logical Volume Manager volume and vice-versa is allowed provided both have the same file system type.

Attention: If you are running GPFS for AIX or GPFS for Linux x86/x86_64 in a multi-node cluster, and all nodes share a mounted GPFS file system, Tivoli Storage Manager processes this file system as a local file system. Tivoli Storage Manager backs up the file system on each node during an incremental backup. To avoid this, you can do *one* of the following:

- Explicitly configure the *domain* statement in the client user options file (dsm.opt) to list the file systems you want that node to back up.
- Set the *exclude.fs* option in dsm.sys to exclude the GPFS file system from backup services.

Maximum file size for operations

Table 31 specifies the maximum file sizes for the native file systems on Tivoli Storage Manager UNIX and Linux client platforms.

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 31. Maximum file size

Platform	Max file size (in bytes)
AIX 5.2 and 5.3 (GPFS)	109,951,162,777,600 (100 TB)
AIX 5.2 and 5.3 (JFS2)	1,099,511,627,264 (1TB - 512)
HP-UX	1,099,511,627,775 (1TB-1)
All Linux clients	9,223,372,036,854,775,807 (8EB-1)
UNIX	4,294,967,295 (4GB)
Solaris	1,099,511,627,775 (1TB-1)

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.

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

Your administrator might have set up schedules to automatically back up files on your workstation. See Chapter 7, “Automating tasks,” on page 151 for information on checking and running the schedules available to you. The following sections discuss how to back up files without using a schedule.

There are two types of incremental backup: *full incremental* and *partial incremental*.

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 directories on the file systems 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 157 for more information about storage management policies.

To perform a full or partial incremental backup using the client Java GUI, select the *Incremental (complete)* option from the pull-down menu at the top of the

screen on the backup window, or use the **incremental** command in the command line interface. Specify file systems, directory trees, or individual files to include in the backup.

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

- Back up new files
- Back up files whose contents changed since the last backup. The client considers a file changed if any of the following attributes changed since the last backup:
 - File size
 - Date or time of last modification
 - Access Control List

If *only* the following items change, they are updated without causing the entire file to be backed up to the server:

- File owner
- File permissions
- Last access time
- Inode

Inode updates are ignored for files stored on an IBM TotalStorage SAN File System.

- Group ID

- 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

You can use the *preservelastaccessdate* option during a backup or archive operation to specify whether to reset the last access date to its original value following a backup or archive operation. By default, the Tivoli Storage Manager client *will not* reset the last access date of any backed up or archived files to their original value prior to the backup or archive operation. See “Preservelastaccessdate” on page 316 for more information.

Directories 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 246.

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 157 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 158 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 158 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 158 for more information.

Journal-based backup on AIX

Journal-based backup is supported on JFS, JFS2, and VxFS and HSM-managed JFS2 file systems and the HACMP environment (not supported on GPFS file systems). A backup for a particular file system will be journal-based when the Tivoli Storage Manager journal daemon has been installed and configured to journal the particular file system, and a valid journal has been established for the file system. See “Configuring the journal daemon” on page 56 for more information on configuring file systems for journal-based backup.

If you install the journal daemon and it is running, then by default the **incremental** command will perform a journal-based backup on selected file systems which are being monitored by the journal daemon. In order to successfully perform a journal-based backup, several conditions must be met. These include:

1. The journal daemon must be set up to monitor the file system that contains the files and directories being backed up.
2. A full incremental backup has to have been run successfully at least once on the file system being backed up.
3. The file space image of the file system at the server cannot have been modified by an administrative command since the last full incremental
4. The storage management policy for the files being backed up cannot have been updated since the last full incremental.

The journal daemon records changes to an object or its attributes in a journal database. During a journal-based backup, the client obtains a list of files that are eligible for backup from the journal database. Journal-based backup can increase backup performance because the client does not scan the local file system or contact the server to determine which files to process. Journal-based backup also reduces network traffic between the client and server.

The client filters the list based on the current include-exclude list and processes, expires, and updates the resulting files according to policy constraints, such as serialization. However, the client ignores the server frequency attribute during a journal-based backup. The reason for this is because a journal-based backup eliminates the backup version query to the server. The client, therefore, does not know how many days have transpired since the last backup of the file. Thus it is not possible for the journal-based backup to use the frequency attribute.

To support journal-based backup you must install and configure the Tivoli Storage Manager daemon. See “Configuring the journal daemon” on page 56 for information about installing and configuring the daemon.

Notes:

1. Journal-based backup is supported on a Version 4.2 or higher Tivoli Storage Manager server.
2. A journal-based backup might not fall back to the traditional incremental backup if the policy domain of your node is changed on the server. This depends on when the policy set within the domain was last updated and the date of the last incremental backup. In this case, you must force a full traditional incremental backup to rebind the files to the new domain. Use the *nojournal* option with the **incremental** command to specify that you want to perform a traditional full incremental backup, instead of the default journal-based backup.

Incremental-by-date backup

For a file system to be eligible for incremental-by-date backups, you must have performed at least one full incremental backup of that file system. Running an incremental backup of only a directory branch or individual file will not make the file system eligible for incremental-by-date backups.

The client backs up only those files whose modification date and time is later than the date and time of the last incremental backup of the file system on which the file resides. Files added by the client after the last incremental backup, but with a modification date earlier than the last incremental backup, are not backed up.

Files that were renamed after the last incremental backup, but otherwise remain unchanged, will not be backed up. Renaming a file does not change the modification date and time of the file. However, renaming a file does change the modification date of the directory in which it is located. In this case, the directory is backed up, but not the files it contains.

If you run an incremental-by-date backup of the whole file system, the server updates the date and time of the last incremental backup. If you perform an incremental-by-date backup on only part of a file system, the server does not update the date of the last full incremental backup. In this case, the next incremental-by-date backup will back up these files again.

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

Full incremental, partial incremental, journal-based, 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.

The memory requirements for an initial journaling environment are the same as the memory requirements for a full file space incremental, because journal-based backups must complete the full file space incremental in order to set the journal database as valid, and to establish the baseline for journaling.

The memory requirements for subsequent journal-based backups are much less. Journal backup sessions run in parallel and are governed by the *resourceutilization* client option in the same manner as normal backup sessions. The size of the

journal database file reverts to a minimal size (less than 1k) when the last entry has been deleted from the journal. Since entries are deleted from the journal as they are processed by the client, the disk size occupied by the journal should be minimal after a complete journal backup. A full incremental backup with journaling active takes less time to process than an incremental-by-date backup. An incremental-by-date backup might not place exactly the same backup files into server storage because the incremental-by-date backup:

- Does not expire backup versions of files that you delete from the workstation.
- Does not rebind backup versions to a new management class if you change the management class.
- Does not back up files with attributes that change, unless the modification dates and times also change.
- Ignores the copy group frequency attribute of management classes. (Journal-based backups also ignore this attribute.)

Selective backup

Use a selective backup when you want to back up specific files or directories regardless of whether a current copy of those files exists on the server. Incremental backups are generally part of an automated system to back up entire file systems. In contrast, selective backups allow you to manually select a set of files to back up regardless of whether they have changed since your last incremental backup.

To perform a selective backup using the client Java GUI, see “Backing up data using the Java GUI” on page 92 for more information. Use the **selective** command to perform a selective backup from the client command line. See “Selective” on page 488 for more information.

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.

Solaris global zone and local zones backup

For Solaris zones support, the incremental and selective backup of file systems should be performed where the file systems are locally mounted. Each local zone should be treated as a separate system with its own Tivoli Storage Manager nodename and the backup should be performed within each of the local zones.

To perform an incremental or selective backup of the local zones from the global zone, the administrator needs to consider which files from the zone directories should be included or excluded in the backup. For example, system files and kernel files would not automatically be excluded.

Saving access permissions

When you back up your files, Tivoli Storage Manager also saves standard UNIX access permissions assigned to the files. Depending on your operating system, it also saves extended permissions. For example, for files on an AIX workstation, Tivoli Storage Manager saves access control lists.

Note: It is possible for an Authorized User to back up files for another user, but this should not cause ownership conflicts. The backup server will properly record that the file belongs to the original owner. The Authorized User does not need to grant the other user access to the backup versions.

Setting a virtual mount point

If you are an Authorized User and you want to back up files beginning with a specific directory within a file system, you can define that directory as a virtual mount point (see “Virtualmountpoint” on page 382).

Defining a virtual mount point within a file system provides a direct path to the files you want to back up, thus saving processing time. It is more efficient than defining the file system with the *domain* option and then using an *exclude* option to exclude the files you do not want to back up. It also allows you to store backups and archives for specific directories in separate storage file spaces.

Backing up data using the Java GUI

You can back up specific files, entire directories, or entire file systems from the directory tree. You can locate the files you want to back up by searching or filtering. Filtering displays only the files matching the filter criteria for your backup. Use the Tivoli Storage Manager client Java GUI to back up your data as follows:

1. Click **Backup** in the Tivoli Storage Manager window. The Backup window appears.
2. Expand the directory tree if necessary. Click on the selection boxes next to the object or objects you want to back up. To search or filter files, click the **Find** icon on the tool bar.

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 Java 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 157 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 227 for more information.
- You can use the Preferences editor to exclude file systems in your default domain from backup processing.

Notes:

1. The Java GUI does not have an **Estimate** button.
2. During installation, Tivoli Storage Manager creates the .adsmrc file to record statistics from the backup-archive client estimate function. The .adsmrc file resides in the directory named in the \$HOME environment variable.

Backing up data using the command line

You can use the **incremental** or **selective** commands to perform backups. Table 32 shows examples of using these commands to perform different tasks. See “Incremental” on page 430 and “Selective” on page 488 for more information about these commands.

Table 32. 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 430 for more information about the incremental command. See “Full and partial incremental backup” on page 87 for detailed information about incremental backups.
Back up the /fs1 and /fs2 file systems in addition to the /home, /usr, and /datasave file systems defined in your client domain.	dsmc incremental -domain="/fs1 /fs2"	See “Domain” on page 227 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 /home file system.	dsmc incremental -domain="all-local -/home"	You cannot use the (-) operator in front of the domain keyword all-local. See “Domain” on page 227 for more information.
Back up <i>only</i> the /fs1 and /fs2 file systems.	dsmc incremental /fs1 /fs2	None
Back up all files in the /home directory and all its subdirectories.	dsmc incremental /home/ -subdir=yes	See “Subdir” on page 356 for more information about the <i>subdir</i> option.

Table 32. Command line backup examples (continued)

Task	Command	Considerations
Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, run an incremental backup of all files and directories under the local snapshot and manage them on the Tivoli Storage Manager server under the file space name /usr.	dsmc incremental /usr -snapshotroot=/snapshot/day1	Tivoli Storage Manager considers the <i>snapshotroot</i> value as a file space name. See "Snapshotroot" on page 351 for more information.
<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 274 for more information about the <i>incrbydate</i> option.
<i>Selective backups</i>		
Back up all files in the /home/proj directory.	dsmc selective /home/proj/	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 488 for more information about the selective command.
Back up all files in the /home/proj directory and all its subdirectories.	dsmc selective /home/proj/ -subdir=yes	If you specify <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. 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 356 for more information about the <i>subdir</i> option.
Back up the /home/dir1/h1.doc and /home/dir1/test.doc files.	dsmc selective /home/dir1/h1.doc /home/dir1/test.doc	If you specify the <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 324 for more information about this option.
Back up a list of files in the /home/filelist.txt file.	selective -filelist=/home/filelist.txt	Use the <i>filelist</i> option to process a list of files. See "Filelist" on page 250 for more information.

Table 32. Command line backup examples (continued)

Task	Command	Considerations
Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, run a selective backup of the /usr/dir1/sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name /usr.	dsmc selective /usr/dir1/sub1/ -subdir=yes -snapshotroot=/snapshot/ day1	Tivoli Storage Manager considers the <i>snapshotroot</i> value as a file space name. See “Snapshotroot” on page 351 for more information.

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 420 for more information.

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 /tmp file space, you are

deleting every backup for every file in that file system and every file you archived from that file system. Carefully consider whether you want to delete a file space.

You can also delete a file space using the **delete filesystem** command. See “Delete Filespace” on page 423 for more information. Use the *class* option with the **delete filesystem** command to delete NAS file spaces. See “Class” on page 202 for more information.

You can delete individual backup versions by using the delete backup command. See “Delete Backup” on page 420 for more information.

You can delete file spaces using the Tivoli Storage Manager GUI or command line clients. To delete NAS file spaces, use the Web client or command line client.

To delete a file space using the GUI, perform the following steps:

1. Select **Utilities** → **Delete Filespaces** from the main window.
2. Click the selection boxes next to the file spaces you want to delete.
3. Click the **Delete** button. Tivoli Storage Manager prompts you for confirmation before deleting the file space.

Group backup: Backing up files from one or more file spaces

You can use the **backup group** command to create and back up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server. See “Backup Group” on page 404 for more information.

A *group backup* allows you to create a consistent point-in-time backup of a group of files that is managed as a single logical entity:

- All objects in the group are assigned to the same management class. See “Include options” on page 269 for more information about using the *include* option to bind a group to a management class.
- Existing *exclude* statements for any files in the group are ignored.
- All objects in the group are exported together.
- All objects in the group are expired together as specified in the management class. No objects in a group are expired until all other objects in the group are expired, even when another group they belong to gets expired.

A group backup can be added to a backup set. See “Restoring data from a backupset” on page 130 for more information about backup sets.

You can perform a full or differential backup using the *mode* option. See “Backup Group” on page 404 and “Mode” on page 291 for more information.

For example, to perform a full backup of all the files named in the `/home/dir1/filelist1` file to the virtual file space `/virtfs` containing the group leader `/home/group1` file, enter:

```
dsmc backup group -filelist=/home/dir1/filelist1 -groupname=group1 -virtualfsname=
/virtfs -mode=full
```

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 machine responsible for performing the backup can change over time, such as with a cluster. The *asnodename* option also allows data to be restored from a different system than the one which performed the backup.

An agent node is a client node which has been granted authority to perform client operations on behalf of a target node.

A target node is a client node which grants authority to one or more agent nodes to perform client operations on its behalf.

Use the *asnodename* option with the appropriate command to back up, archive, restore, and retrieve data under the target node name on the Tivoli Storage Manager server. This support is only available with Tivoli Storage Manager Version 5.3 and higher server and client. To enable this option, follow these steps:

- Install the backup-archive client on all nodes in a 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 target nodes as traditional nodes, especially if you encrypt your files before backing them up to the server.

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** → **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.

Follow these steps to verify that your client node is now accessing the server as the target node:

1. Open the tree window and check that the target node name specified by the **As Node Name** field appears, or
2. Verify the target node name in the **Accessing As Node** field in the **Connection Information** window.

To return to single node operation, delete the **As Node Name** from the **Accessing As Node** field in the **General** → **Preferences** tab.

Suggestions:

- All agent nodes in the multiple node environment should be of the same platform type.
- Do not use target nodes as traditional nodes. Use them only for multiple node processing.

Restrictions enforced within a proxied session:

- You cannot access another node (either from GUI drop down or use of the Fromnode option).
- You cannot use the *clusternode* option.
- You cannot perform NAS backup or restore.

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

1. Specify *encryptkey=save* in the options file.
2. Back up at least one file with *asnode=ProxyNodeName* to create a local encryption key on each agent node in the multiple node environment.

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

1. Specify *encryptkey=prompt* in the options file.
 2. Ensure that users of the agent nodes in the multiple node environment are using the same encryption key.
- If you change the encryption key, you must repeat the previous steps.
 - Use the same encryption key for all files backed up in the shared node environment.

See “Asnodename” on page 197 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 operation 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
```

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=C:
  startdate=05/21/2005 starttime=01:00
define schedule standard proxy2 description="NODE-B proxy schedule"
  action=incremental options="-asnode=NODE-Z" objects=D:
  startdate=05/21/2005 starttime=01:00
define schedule standard proxy3 description="NODE-C proxy schedule"
  action=incremental options="-asnode=NODE-Z" objects=E:
  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*.

Start the schedules using the following client command:

```
dsmc sched
```

You can also use the client acceptor daemon (dsmcad), with managedservices set to schedule in the systems options file.

Notes:

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

Scheduling for backing up a GPFS filesystem

The following examples show the administrative client server commands using the scheduler to back up a GPFS filesystem, using three nodes in the GPFS cluster which participate in the backup.

1. The administrator defines four nodes on the Tivoli Storage Manager server: node_1, node_2, node_3 and node_gpfs. node_1, node_2 and node_3 are only used for authentication; all filespaces are stored with node_gpfs.

```
REGISTER NODE node_1 mysecretpa5s
REGISTER NODE node_2 mysecretpa5s
REGISTER NODE node_3 mysecretpa5s
REGISTER NODE node_gpfs mysecretpa5s
```

2. The administrator defines a proxynode relationship between the nodes:

```
GRANT PROXYNODE TARGET=node_gpfs AGENT=node_1, node_2, node_3
```

3. The administrator defines the node name and asnode name for each of the machines in their respective dsm.sys files:

```
nodename node_1
asnodename node_gpfs
```

4. The administrator defines a schedule for only node_1 to do the work:

```
DEFINE SCHEDULE STANDARD GPFS_SCHEDULE ACTION=MACRO
OBJECTS="gpfs_script"
DEFINE ASSOCIATION STANDARD GPFS node_gpfs
```

5. To execute the schedule on node node_gpfs, enter the client command:

```
DSMC SCHED
```

Another way to back up GPFS is to use Tivoli Storage Manager to look for the incremental changes. The GPFS filesystem can be divided into three branches and each branch can be statically assigned to each node using the *virtualmountpoint* option. In the following example, you have a file system called /gpfs with three branches: /gpfs/branch_1, /gpfs/branch_2, and /gpfs/branch_3.

1. The administrator defines four nodes on the Tivoli Storage Manager server: node_1, node_2, node_3 and node_gpfs. node_1, node_2 and node_3 are only used for authentication; all filespaces are stored with node_gpfs.

```
REGISTER NODE node_1 mysecretpa5s
REGISTER NODE node_2 mysecretpa5s
REGISTER NODE node_3 mysecretpa5s
REGISTER NODE node_gpfs mysecretpa5s
```

2. The administrator defines a proxynode relationship between the nodes:

```
GRANT PROXYNODE TARGET=node_gpfs AGENT=node_1,node_2,node_3
```

- The administrator defines the node name, virtualmountpoint and domain for each of the three machines in their respective dsm.sys files:

```
nodename          node_1
virtualmountpoint /gpfs/branch_1
domain            /gpfs/branch_1
```

Note: The user does not want to define *asnodename* in the options file. In this case the *asnodename* must be on the schedule so that each one of the nodes can have its own schedule associated with its real node name.

- The administrator defines a schedule for all nodes: node_1, node_2 and node_3

```
DEFINE SCHEDULE STANDARD GPFS_SCHEDULE OPTIONS="-asnode=node_gpfs"
DEFINE ASSOCIATION STANDARD GPFS node_1,node_2,node_3
```

- To start the scheduler on the three nodes, enter the client command:

```
DSMC SCHED
```

Scheduling for backing up a HACMP™ cluster

In the following examples, HACMP is configured for two AIX hosts, host_a and host_b. Along with their own local data, the hosts are sharing SSA storage which has two filespaces: /ssa1 and /ssa2.

The CLUSTERNODE example shows how the *clusternode* option is used in a current HACMP environment.

- The administrator defines 3 nodes on the Tivoli Storage Manager server: host_a, host_b, cluster_group

```
REGISTER NODE host_a mysecretpa5s
REGISTER NODE host_b mysecretpa5s
REGISTER NODE cluster_group mysecretpa5s
```

- The administrator defines a dsm.opt file on host_a and host_b (note that the opt files are different on each host):

```
NODENAME host_a (option can be left as default)
DOMAIN   /home /usr ... etc.
```

- The administrator defines a dsm.opt file located somewhere on one of the cluster disk groups, e.g., /ssa1/tsm/dsm.opt,

```
NODENAME      cluster_group
DOMAIN        /ssa1 /ssa2
CLUSTERNODE   YES
```

- The administrator defines a schedule on the Tivoli Storage Manager server:

```
DEFINE SCHEDULE STANDARD CLUSTER_BACKUP
```

- The administrator associations for each of the 3 nodes:

```
DEFINE ASSOC STANDARD CLUSTER_BACKUP host_a,host_b,cluster_group
```

At any one time, there are three instances of the Tivoli Storage Manager Backup-Archive client schedule running (with the scheduler for cluster_group being part of the cluster resources that will failover whenever the cluster group disk resources failover. Thus, it would be running on either host_a or host_b but not both simultaneously).

- All three node names contain data on the Tivoli Storage Manager server.

The ASNODE example shows how this can be configured without the use of CLUSTERNODE; this is a generic solution which could be applied to UNIX cluster solutions to which we do not have support, e.g., Veritas Cluster Server for Solaris.

- The administrator defines 3 nodes on the Tivoli Storage Manager server host_a, host_b, cluster_group:

- ```
REGISTER NODE host_a mysecretpa5s
REGISTER NODE host_b mysecretpa5s
REGISTER NODE cluster_group mysecretpa5s
```
- The administrator defines a proxy node relationship between host\_a and host\_b to hacmp\_cluster

```
GRANT PROXYNODE TARGET=cluster_group AGENT=host_a,host_b
```
  - The administrator defines a dsm.opt file on host\_a and host\_b to handle the local file systems:

```

NODENAME host_a (option can be left as default)
DOMAIN /home /usr ... etc.

NODENAME host_b (option can be left as default)
DOMAIN /home /usr ... etc.
```
  - The administrator defines a dsm.opt file on the cluster resource to handle the backup of the clustered resources, e.g. /ssa1/tsm/dsmcluster.opt (note the nodename will be the default nodename which will either be host\_a or host\_b depending on which machine contains the cluster group at any given time):

```

DOMAIN /ssa1 /ssa2
ASNODE cluster_group
```
  - The administrator defines a schedule on the Tivoli Storage Manager server:

```
DEFINE SCHEDULE STANDARD CLUSTER_BACKUP
```
  - The administrator defines associations for each one of the 3 nodes.

```
DEFINE ASSOC STANDARD CLUSTER_BACKUP host_a,host_b,cluster_group
```
  - At any one time, there are three instances of the Tivoli Storage Manager Backup-Archive client schedule running with the scheduler for node hacmp\_cluster running on either host\_a or host\_b but not both (it is included in the cluster resources that would failover). This scheduler would point to the dsmcluster.opt that is defined on each host. The three instances would be started as:

```

[host_a] dsmc sched
[host_b] dsmc sched
[cluster_group] dsmc sched -optfile=/ssa/tsm/dsmcluster.opt
```
  - All three node names will contain data on the Tivoli Storage Manager server.

For more information about the server scheduler commands, see *IBM Tivoli Storage Manager Administrator Command Reference*.

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

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

```
def sched domain_name sched_name options='-asnode=target'
```
- Each agent node must have its schedule associated with node

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

---

## Associating a local snapshot with a server file space

Use the *snapshotroot* option with the **incremental** and **selective** commands in conjunction with a third-party application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. The *snapshotroot* option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot. See “Snapshotroot” on page 351 for more information.



---

## Performing an image backup

### Root User

From your local workstation, you can back up a logical volume as a single object (image backup) on your system. The traditional *static* image backup prevents write access to the volume by other system applications during the operation.

An image backup provides the following benefits:

- Backs up file systems containing a large number of files faster than a full file system incremental back up.
- Improves the speed with which Tivoli Storage Manager restores file systems containing many small files.
- Conserves resources on the server during backups since only one entry is required for the image.
- Provides a point-in-time picture of your logical volume, which might be useful if your enterprise needs to recall that information.
- Restores a corrupt file system or raw logical volume. Data is restored to the same state it was when the last logical volume backup was performed.

The traditional *static* image backup prevents write access to the volume by other system applications during the operation. Use the *imagetype=dynamic* option to back up the volume *as is* without remounting it read-only. Corruption of the backup can occur if applications continue to write to the volume while the backup is in progress, which can result in inconsistent data and data loss after the restore. This option overrides the copy serialization value in the management class to perform an image backup. When backing up an image using *imagetype=dynamic*, after a restore always run the *fsck* utility.

In order to restore an image backup of a volume, the Tivoli Storage Manager client must be able to obtain an exclusive lock on the volume being restored.

**Attention:** Tivoli Storage Manager dynamic image backups should not be used for file systems, because the file system's internal design could provide inconsistent data even when there is no write activity, such as the AIX JFS2 file system. Dynamic image backup could result in a fuzzy image, and is not guaranteed to be valid or complete when restored. If the Tivoli Storage Manager client fails to mount the file system after restoring image, you can run *fsck*. However, running *fsck* can affect the integrity of large amounts of data.

**Recommendation:** To prevent data loss, avoid using dynamic *imagetype*, and ensure that there is no write activity on the volume while the backup is in progress.

For Linux clients *only*: By default, Tivoli Storage Manager performs a *snapshot* image backup of file systems residing on a logical volume created by the Linux Logical Volume Manager, during which the volume is available to other system applications.

You can use the *imagetype* option with the **backup image** command or the *include.image* option to specify whether to perform a static, dynamic, or snapshot image backup. See "Imagetype" on page 265 for more information.

## Before you perform an image backup

Before you perform an image backup, consider the following:



- Ensure that no other application is using the volume when you run a static image backup. To ensure a consistent image during backup processing, the client will unmount and remount the volume as read only, so that no other applications can write to it. If the volume is in use when the client attempts to unmount, the backup will fail.

If the client cannot unmount and remount the volume as read only because it is in use, and snapshot image backup is not available, you can use the *imagetype* option to force the client to perform an image backup without unmounting and remounting the volume in read-only mode. Set the *imagetype* option to *dynamic* in an *include.image* statement or from the command line. The backup can be corrupted if applications write to the volume while the backup is in progress. This can be corrected by running **fsck** after a restore to fix any corrupted blocks. See “Include options” on page 269 for more information.

*Do not* include system files in an image backup because file systems being actively used cannot be unmounted.

For AIX and Linux only: If you perform an image backup of a mounted file system which is mounted to another mount point and specified in the file system table, then after completing the image backup, all mount options for this file system, except read or write state, will be lost.

**Important:** If a mounted file system has nested mount points, unmount them before attempting a backup. Otherwise, Tivoli Storage Manager will be unable to unmount the volume. The file system is rendered *busy* if it contains any mounts.

- Use the *include.image* option to assign a management class to the volume image. If you do not assign a management class, the default management class is used for the image. See “Include options” on page 269 for more information. See Chapter 8, “Understanding storage management policies,” on page 157 for more information about management classes.

**Note:** Copy serialization is controlled by the client only for image backup using the *imagetype* option. See “Imagetype” on page 265 for more information.

- You can exclude a volume from image backup using the *exclude.image* option. See “Exclude options” on page 246 for more information.
- You must use the mount point for the file system volume on which you want to perform an image backup. Tivoli Storage Manager will not back up a file system volume without the use of a mount point. Back up file systems using the mounted name. For example, if `/dev/lv01` is formatted as a file system mounted on `/home`, enter this command to perform an image backup of this volume:

```
dsmc backup image /home
```

Back up raw volumes using the device name. For example, if `/dev/lv02` is a raw volume, enter this command to perform an image backup of this volume:

```
dsmc backup image /dev/lv02
```

### Volume device type support for an image backup

The following table lists devices supported by the **backup image** command. A raw device might be a disk slice, a partition, or a logical volume.

Table 33. Volume device type support for an image backup

| Logical volume manager                | Raw device types | Sample device name          | Backup image command support |
|---------------------------------------|------------------|-----------------------------|------------------------------|
| AIX Logical Volume Mgr                | Logical Volumes  | <code>/dev/lv00</code>      | AIX                          |
| Sun Solstice DiskSuite Volume Manager | Metadevices      | <code>/dev/md/dsk/dl</code> | Solaris                      |

Table 33. Volume device type support for an image backup (continued)

| Logical volume manager                       | Raw device types | Sample device name                                                  | Backup image command support                    |
|----------------------------------------------|------------------|---------------------------------------------------------------------|-------------------------------------------------|
| Solstice DiskSuite or Solaris Volume Manager | Metadevices      | /dev/md/dsk/dl                                                      | Solaris                                         |
| Veritas Volume Mgr                           | Logical Volumes  | /dev/vx/dsk/rootg/vol01<br>/dev/vg00/lvol01<br>/dev/vx/rootdg/vol01 | Solaris SPARC<br>HP-UX,<br>HP-UX Itanium 2, AIX |
| Raw Disk                                     | Partitions       | /dev/hda1, /dev/sda3                                                | Linux x86/x86_64,<br>Linux on POWER             |
| Raw Disk                                     | Partitions       | /dev/dasdx                                                          | Linux zSeries                                   |
| Linux Logical Volume Mgr                     | Logical Volumes  | /dev/myvolgroup/<br>myvolume                                        | All Linux                                       |
| Raw Disk                                     | Disk Slices      | /dev/dsk/c0tld0s0                                                   | Solaris SPARC/x86                               |

The client must support the raw device type on the specific platform in order to perform an image backup of a raw device. If you want to perform an image backup for a file system mounted on a raw device, the raw device must be supported. Remember to specify raw devices by their block device name.

**Notes:**

1. For the Linux clients, image backup is only supported on partitions with id 0x83 or logical volumes created with the Linux Logical Volume Manager. Backing up other partitions, such as extended partitions that contain mounted file systems or database data, might produce inconsistent backup data if the data changes during the image backup operation.
2. On HP-UX, raw logical volume backup does not support devices other than logical volumes, such as /dev/dsk/c0t0d1. Logical volume devices usually take the form /dev/vgXY/lvolAB.
3. You should not back up disk slices containing cylinder 0 on Solaris because the volume table of contents (VTOC) will be overwritten after a restore.
4. For AIX, HP-UX, and Solaris: You can perform image operations on volumes created using Veritas Volume Manager. Tivoli Storage Manager will initially support static (default) and dynamic image type for backup.
5. For Solaris 10 clients, image backup is supported only for file systems that are assigned from the global zone to the non-global zone by exporting the device (add device + set match). Image backup is *not* supported for other file systems in the non-global zones because the non-global zone does not have the authority to mount or unmount the file system. An image backup on a file system other than one assigned from the global zone to the non-global zone will result in a message such as *ANS1068E Device is not local*
6. Also, for Solaris 10 clients, the overlap device of the root disk (c0t0d0s2) is not supported for raw device backup. Avoid using this feature on disks or slices that are used as swapping devices.

Meta devices created by the Veritas Volume Manager must be listed, including the disk group in /etc/vfstab, to be recognized by the Tivoli Storage Manager Backup-Archive client for an image backup of file systems. Raw devices should not be listed in /etc/vfstab. For example, the following is the correct meta device name to be used in the /etc/vfstab file:

```
/dev/vx/dsk/<disk group>/<meta device name>
```

| Specifying `/dev/vx/dsk/` would not be recognized correctly, and you would  
| receive an error (ANS1134E).

| Disk slices containing cylinder 0 should not be backed up or restored. In this  
| case the VTOC will be overwritten. If you need to back up the first disk slice,  
| exclude cylinder 0 by starting the disk slice from cylinder 1 (use the format  
| utility). The Tivoli Storage Manager Backup-Archive client does not check  
| whether cylinder 0 is contained in the device that is overwritten during a  
| restore.

## Utilizing image backup to perform file system incremental backup

There are two methods of utilizing image backups to perform efficient incremental backups of your file system. These backup methods allow you to perform point-in-time restore of your file systems and improve backup and restore performance. You can perform the backup only on formatted volumes; not on raw logical volumes. You can use one of the following methods to perform image backups of volumes with mounted file systems.

### Method 1: Using image backup with file system incremental

1. Perform a full incremental backup of the file system (See “Backing up data using the Java GUI” on page 92 for instructions). This establishes a baseline for future incremental backups.
2. Perform an image backup of the same file system to make image restores possible. See “Performing an image backup using the GUI” on page 106 for instructions.
3. Perform incremental backups of the file system periodically to ensure that the server records additions and deletions accurately.
4. Perform an image backup periodically to ensure faster restore.
5. Restore your data by performing an incremental restore (See “Performing an image restore using the GUI” on page 128 for instructions). Ensure that you select the **Image plus incremental directories and files** and **Delete inactive files from local** options in the Restore Options window before beginning the restore. During the restore, the client does the following:
  - Restores the most recent image on the server.
  - Deletes all of the files restored in the previous step which are inactive on the server. These are files which existed at the time of the image backup, but were subsequently deleted and recorded by a later incremental backup.
  - Restores new and changed files from the incremental backups.

| **Note:** If an incremental backup is performed several times after backing up image,  
| make sure that the Tivoli Storage Manager server’s backup copy group has  
| enough versions for existing and deleted files on the server so that the  
| subsequent restore image with *incremental* and *deletefiles* options will be  
| able to delete files correctly.

### Method 2: Using image backup with incremental-by-date image backup

1. Perform an image backup of the file system. See “Performing an image backup using the GUI” on page 106 for instructions.
2. Perform an incremental-by-date image backup of the file system (See “Performing an image backup using the GUI” on page 106 for instructions). This sends only those files that were added or changed since the last image backup to the server.

3. Periodically, perform full image backups (See “Performing an image backup using the GUI” for instructions).
4. Restore your volume by performing an incremental restore (See “Performing an image restore using the GUI” on page 128 for instructions). Ensure that you select the **Image plus incremental directories and files** option in the Restore Options window before beginning the restore. This will first restore the most recent image and will then restore all the incremental backups performed since that date.

**Note:** You should perform full image backups periodically in the following cases:

- When a file system changes substantially (more than 40%), as indicated in step 4 of method 1 and step 3 of method 2. On restore, this would provide a file system image close to what existed at the time of last incremental-by-date image backup and also improves restore time.
- As appropriate for your environment.

This will improve restore time because fewer changes are applied from incrementals.

The following restrictions apply when using method 2:

- The file system can have no previous full incremental backups.
- Incremental-by-date image backup does not inactivate files on the server; therefore, when you restore an image with the *incremental* option, files deleted after the original image backup will be present after the restore.
- If this is the first image backup for the file system, a full image backup is performed.
- If file systems are running at or near capacity, an out-of-space condition could result during the restore.

## Comparing methods 1 and 2

To help you decide which method is appropriate for your environment, Table 34 is a comparison of methods 1 and 2.

Table 34. Comparing incremental image backup methods

| Method 1: Using image backup with file system incremental                                                                                                           | Method 2: Using image backup with incremental-by-date image backup                                                                                                                                                                                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Files are expired on the server when they are deleted from the file system. On restore, you have the option to delete files which are expired on server from image. | Files are not expired on server. After the image incremental restore completes, all files deleted on the file system after the image backup will be present after the restore. If file systems are running at or near capacity, an out-of-space condition could result. |
| Incremental backup time is the same as regular incremental backups.                                                                                                 | Incremental image backup is faster because the client does not query the server for each file that is copied.                                                                                                                                                           |
| Restore is much faster compared to a full incremental file system restore.                                                                                          | Restore is much faster compared to a full incremental file system restore.                                                                                                                                                                                              |
| Directories deleted from the file system after the last image backup are not expired.                                                                               | Directories and files deleted from the file system after the last full image backup are not expired.                                                                                                                                                                    |

## Performing an image backup using the GUI

When you perform an image backup using the client Java GUI **Image Backup** option, Tivoli Storage Manager honors the image type setting of the *include.image*

or *imagetype* options in your `dsm.sys`. If you set the image type to *snapshot*, the client performs a snapshot image backup of file systems residing on a logical volume created by the Linux Logical Volume Manager, during which the volume is available to other system applications. If you set the image type to *static*, the client will unmount and remount the volume as read-only, so that no other applications can access it. If you do not specify either of these options, the client performs a snapshot image backup.

For the AIX, HP-UX, and Solaris clients, selecting the **Image Backup** option performs a static or dynamic image backup depending on the image type setting of the *include.image* or *imagetype* options in `dsm.sys`. If the image type is set to *static*, the client will unmount and remount the volume as read-only, so that no other applications can access it. If the image type is set to *dynamic*, the client performs the image backup without making the file system read-only during the backup. If you do not specify either of these options, the client performs a static image backup.

To create an image backup of your file system or raw logical volume, perform the following steps:

1. Click on the **Backup** button in the Tivoli Storage Manager main window. The Backup window appears.
2. Expand the directory tree and select the objects you want to back up. To back up a raw logical volume, locate and expand the **RAW** directory tree object.
  - To perform a static image backup, select **Image Backup** from the drop-down list.
  - To perform a snapshot image backup, select **Snapshot Image Backup** from the drop-down list. (Linux clients)
  - To perform an incremental-by-date image backup, select **Incremental image (date only)** from the drop-down list.
3. Click **Backup**. The Backup **Task List** window displays the backup processing status. The Backup Report window displays a detailed status report.

### Considerations

- To modify specific backup options, click the **Options** button. The options you select are effective during the current session *only*.

**Linux only:** The Tivoli Storage Manager Version 5.4 client will not recognize any LVM1 volumes for image operations. However, it will allow prior image backups of LVM1 volumes to be restored on LVM2 volumes. Table 35 shows the combinations involving the old and new client levels handling LVM1 and LVM2 volumes for different image operations.

Table 35. LVM1 and LVM2 image operation comparisons

| Tivoli Storage Manager client version | LVM1 Backup and Restore            | LVM2 Backup and Restore           | Mixed Volumes                                               |                                     |
|---------------------------------------|------------------------------------|-----------------------------------|-------------------------------------------------------------|-------------------------------------|
|                                       |                                    |                                   | Backup: LVM1, Restore: LVM2                                 | Backup: LVM2, Restore: LVM1         |
| V5.3 and prior                        | YES                                | Only static image for file system | NO                                                          | NO - raw volumes are not supported  |
| V5.4 and beyond                       | NO<br>Error msg ANS1090E displayed | YES                               | YES<br>LVM1 vol must have been backed up using prior client | NO<br>Restore to LVM1 vol will fail |

## Performing an image backup using the command line

Use the **backup image** and **restore image** commands to perform image backup and restore operations on a single volume. See “Backup Image” on page 406 and “Restore Image” on page 476 for more information.

You can use the *imagetype* option with the **backup image** command or the *include.image* option in your `dsm.sys` file or on the command line to specify whether to perform a static, dynamic, or snapshot image backup. See “Imagetype” on page 265 for more information.

Use the *mode* option with the **backup image** command to perform an incremental-by-date image backup that backs up only new and changed files after the last full image backup. However, this only backs up files with a changed date, not files with changed permissions. See “Mode” on page 291 for more information.

---

## Backing up NAS file systems

Through support of Network Data Management Protocol (NDMP), Tivoli Storage Manager Windows, AIX, and Solaris servers can efficiently back up and restore network attached storage (NAS) file system images to automated tape drives or libraries that are locally attached to Network Appliance and EMC Celerra NAS file servers. *NDMP support is available only on IBM Tivoli Storage Manager Extended Edition.* See “NDMP support requirements (Extended Edition only)” on page 12 for NDMP support requirements.

For information on how to configure NDMP support on the Tivoli Storage Manager server, see the following publications:

- *IBM Tivoli Storage Manager for AIX Administrator's Guide, SC32-0117*
- *IBM Tivoli Storage Manager for Sun Solaris Administrator's Guide, SC32-0120*
- *IBM Tivoli Storage Manager for Windows Administrator's Guide, SC32-0121*

After configuring NDMP support, the server connects to the NAS device and uses NDMP to initiate, control, and monitor each backup and restore operation. The NAS device performs outboard data transfer to and from the NAS file system to a locally attached library.

The benefits of performing backups using NDMP include the following:

- LAN-free data transfer.
- High performance and scalable backups and restores.
- Backup to local tape devices without network traffic.

The following support is provided:

- Full file system image backup of all files within a NAS file system.
- Differential file system image backup of all files that have changed since the last full image backup.
- Parallel backup and restore operations when processing multiple NAS file systems.
- Choice of interfaces to initiate, monitor, or cancel backup and restore operations:
  - Web client
  - Backup-archive command line client
  - Administrative command line client (backup and restore operations can be scheduled using the administrative command scheduler)
  - Administrative Web client

The following functions are *not* supported:



- Archive and retrieve
- Client scheduling. Use server commands to schedule a NAS backup.
- Detection of damaged files.
- Data-transfer operations for NAS data stored by Tivoli Storage Manager:
  - Migration
  - Reclamation
  - Storage pool backup and restore
  - Move data
  - Export
  - Backup set generation

## Backing up NAS file systems using the Web client GUI

For information on how to install and configure the Web client, see “Configuring the Web client” on page 50.

For both the Web client GUI and the command line client, you must specify *passwordaccess=generate* (which is a current Web client restriction for the client node) and the *authentication=on* must be specified at the server. You are always prompted for a user ID and password. To display NAS nodes and perform NAS functions, you must enter an authorized administrative user ID and password. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web.

You can use the *toc* option with the *include.fs.nas* option in *dsm.sys* to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup. See “Toc” on page 371 for more information. If you save TOC information, you can use Tivoli Storage Manager Web client to examine the entire file system tree and select files and directories to restore. Creation of a TOC requires that you define the TOCDESTINATION attribute in the backup copy group for the management class to which this backup image is bound. Note that TOC creation requires additional processing, network resources, storage pool space, and possibly a mount point during the backup operation.

To back up NAS file systems using the Web client GUI:

1. Click **Backup** from the main window. The Backup window is displayed.
2. Expand the directory tree if necessary.

### Notes:

- a. The root node called **Nodes** is not selectable. This node only appears if a NAS plug-in is present on the client machine.
  - b. NAS nodes display on the same level as the client workstation node. Only nodes for which the administrator has authority appear.
  - c. You can expand NAS nodes to reveal file spaces, but no further expansion is available (no file names).
3. Click the selection boxes next to the nodes or file systems you want to back up.
  4. Click the type of backup you want to perform in the backup type pull-down menu. The NAS backup type list is active only when you first select NAS backup objects. **Full backup** backs up the entire file system. **Differential** backs up the changes since the most recent full backup.
  5. Click **Backup**. The NAS Backup **Task List** window displays the backup processing status and progress bar. The number next to the progress bar indicates the number of bytes backed up so far. After the backup completes, the NAS Backup Report window displays processing details, including the *actual* size of the backup including the total bytes backed up.

**Note:** If it is necessary to close the Web browser session, current NAS operations will continue after disconnect. You can use the **Dismiss** button on the NAS Backup **Task List** window to quit monitoring processing without ending the current operation.

- (Optional) To monitor processing of an operation from the GUI main window, open the **Actions** menu and select **TSM Activities**. During a backup, the status bar indicates processing status. A percentage estimate is not displayed for differential backups.

Considerations:

- Workstation and remote (NAS) backups are mutually exclusive in a Backup window. After selecting an item for backup, the next item you select must be of the same type (either NAS or non NAS).
- Details will not appear in the right-frame of the Backup window for NAS nodes or file systems. To view information about objects in a NAS node, highlight the object and select **View** → **File Details** from the menu.
- To delete NAS file spaces, select **Utilities** → **Delete Filespaces**.
- Backup options do not apply to NAS file spaces and are ignored during a NAS backup operation.

To restore NAS file system images using the Web client GUI, see “Restoring NAS file systems” on page 134.

## Backing up NAS file systems using the command line

Table 36 lists the commands and options you can use to back up NAS file system images from the command line.

Table 36. NAS options and commands

| Option or command     | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                         | Page |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>domain.nas</i>     | Use the <i>domain.nas</i> option to specify the volumes to include in your default domain for NAS backups.                                                                                                                                                                                                                                                                                                                                         | 232  |
| <i>exclude.fs.nas</i> | Use the <i>exclude.fs.nas</i> option to exclude file systems on the NAS file server from an image backup when used with the <b>backup nas</b> command. This option is for AIX and Solaris clients <i>only</i> .                                                                                                                                                                                                                                    | 246  |
| <i>include.fs.nas</i> | Use the <i>include.fs.nas</i> option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system image backup, using the <i>toc</i> option with the <i>include.fs.nas</i> option in your dsm.sys. See “Toc” on page 371 for more information. This option is valid for AIX and Solaris clients <i>only</i> . | 269  |
| <b>query node</b>     | Use the <b>query node</b> command to display all the nodes for which a particular administrative user ID has authority to perform operations. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using.                                                                                                                                             | 457  |
| <b>backup nas</b>     | Use the <b>backup nas</b> command to create an image backup of one or more file systems that belong to a Network Attached Storage (NAS) file server.                                                                                                                                                                                                                                                                                               | 410  |



Table 36. NAS options and commands (continued)

| Option or command        | Definition                                                                                                                                                                                                                          | Page |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>toc</i>               | Use the <i>toc</i> option with the <b>backup nas</b> command or the <i>include.fs.nas</i> option to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup.                   | 371  |
| <b>monitor process</b>   | Use the <b>monitor process</b> command to display current backup and restore processes for all NAS nodes for which an administrative user has authority. The administrative user can then select one process to monitor.            | 438  |
| <b>cancel process</b>    | Use the <b>cancel process</b> command to display current backup and restore processes for all NAS nodes for which an administrative user has authority. From the display, the administrative user can select one process to cancel. | 415  |
| <b>query backup</b>      | Use the <b>query backup</b> command with the <i>class</i> option to display information about file system images backed up for a NAS file server.                                                                                   | 444  |
| <b>query filesystem</b>  | Use the <b>query filesystem</b> command with the <i>class</i> option to display a list of file spaces belonging to a NAS node.                                                                                                      | 449  |
| <b>delete filesystem</b> | Use the <b>delete filesystem</b> command with the <i>class</i> option to display a list of file spaces belonging to a NAS node so that you can choose one to delete.                                                                | 423  |

- NAS nodes represent a new node type. The NAS node name uniquely identifies a NAS file server and its data to Tivoli Storage Manager. You can prefix the NAS node name to the file specification to specify the file server to which the include statement applies. If you do not specify a NAS node name, the file system you specify applies to all NAS file servers.
- Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.

**Note:** When you initiate a NAS backup operation using the command line client, GUI client, or Web client the server starts a process to initiate, control, and monitor the operation. It might take several moments before you notice progress at the command line client interface because the server must perform mount and other necessary tasks before data movement occurs.

## Backing up NFS file systems

Tivoli Storage Manager supports NFSV4 file systems on AIX 5.3 and above. NFSV4 file systems are backed up as NFS file systems. The NFSV4 ACLs are not supported.

The Global Name Space Feature (referral) of AIX 5300-03 is supported

NFS V4 clients can back up NFS file systems mounted using the "referral" server. All file systems in the global name space will be backed up under a single file space. The following are some examples:

```
server 'publications' has /doc file system
server 'projects' has /projects file system
server 'data' has /data file system
```

Server 'account1' is the main NFS server that exports all these files systems using referral, and it is the server that all of the clients know about. /etc/exports on 'account1' looks like this:

```

/doc -vers=4,refer=/doc@publications
/projects -vers=4,refer=/projects@projects
/data -vers=4,refer=/data@data

```

Client system 'payroll' mounts directories from 'account1' server and can access all three file systems:

```

payroll:/#mount -o vers=4 account1:/ /mnt
payroll:/#ls /mnt
doc/ projects/ data/

```

Client 'payroll' can back up /mnt as one NFS file system to back up all other file systems

Use of the *virtualmountpoint* option is recommended when backing up NFSV4 file systems using global name space for management granularity and improved performance. For example, the following entries could be added for the above example in the dsm.sys stanza to back up each mounted directory as separate filesystems:

```

VIRTUALLYMOUNTPOINT /doc
VIRTUALLYMOUNTPOINT /projects
VIRTUALLYMOUNTPOINT /data

```

---

## Backing up the WebSphere Application Server (WAS)

If you installed the Data Protection for WebSphere Application Server, you can use Tivoli Storage Manager to back up the Version 5.0 WebSphere Application Server Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server (contains setup, application files, and configuration information) to the Tivoli Storage Manager server. You can use this information to recover a corrupted node application or an entire node (or nodes) in the event of an accident or disaster.

For information on backing up the WebSphere Application Server, see *IBM Tivoli Storage Manager for Application Servers: Data Protection for WebSphere Application Server Installation and User's Guide*.

---

## 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. Click the **Help** button in the Backup Report window for context help.

On the backup-archive command line the name of each file is displayed after it is sent to the server. The progress indicator shows overall progress. Informational messages might display as follows:

Table 37. Client command line informational messages

| Informational message | Meaning                                   |
|-----------------------|-------------------------------------------|
| Directory-->          | Indicates the directory that you back up. |

Table 37. Client command line informational messages (continued)

| Informational message              | Meaning                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Normal File-->                     | Any file that is not a directory, symbolic link or special file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Special File-->                    | Special files define devices for the system or temporary files created by processes. There are three basic types of special files: FIFO (first-in, first-out), block, and character. FIFO files are also called pipes. Pipes are created by one process to temporarily allow communication with another process. These files cease to exist when the first process finishes. Block and character files define devices. Tivoli Storage Manager processes only device and named pipe special files. Socket special files are not processed. |
| Symbolic Link-->                   | Indicates that Tivoli Storage Manager backs up a symbolic link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Updating-->                        | Indicates that only the file meta data is sent, not the data itself.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Expiring-->                        | Indicates an object (file or directory) on the server that no longer exists on the client is expired and made inactive on the server.                                                                                                                                                                                                                                                                                                                                                                                                     |
| Total number of objects inspected: | As indicated. When using journal-based backup, the number of objects inspected might be less than the number of objects backed up.                                                                                                                                                                                                                                                                                                                                                                                                        |
| 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 165 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 87 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 37. 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><b>Note:</b> 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 file space or other set of files. When multiple sessions are open during backup, the data transfer time represents the sum of the times reported by all sessions. In this case, aggregate data transfer time is incorrectly reported as higher. However, when running with a single session, the aggregate data transfer rate should always be reported as lower than the network data transfer rate.</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: | <p>As indicated.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| LanFree bytes transferred:         | <p>The total number of data bytes transferred during a lan-free operation. If the <i>enablelanfree</i> option is set to <i>no</i>, this line will not appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## Backup: Additional considerations

This section includes topics related to incremental and selective backups. 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 or virtual mount point from which the files originated.

For example, if you have a file system named `/home`, and you back up a file named `doc1` in the `/home/monnett` directory, Tivoli Storage Manager stores the file in a file space named `/home`. If an Authorized User later defines `/home/monnett` as a virtual mount point, any files you back up from the `/home/monnett` directory, such as `doc2`, are stored in a file space named `/home/monnett`. If you enter this command:

```
dsmc query backup "/home/monnett/*"
```

Tivoli Storage Manager looks for files in the `/home/monnett` file space. It always looks for a file in the file space with the longest name that matches the file specification you include in a command. It locates the file named `doc2` that was

backed up after the virtual mount point was defined. However, it does not locate the file named doc1 because that file was backed up before the virtual mount point was defined and the backup was stored in the /home file space.

To list or restore the doc1 file using a command, you must explicitly specify the file space name by enclosing it in braces. For example:

```
dsmc query backup "{/home}/monnett/*"
dsmc restore {/home}/monnett/doc1
```

If the authorized user subsequently removes the /home/monnett virtual mount point, and you then back up additional files in the /home/monnett directory, the backups are once again stored in the /home file space. For example, if you now back up a file named doc3 in the /home/monnett directory, it is stored in the /home file space. It is not stored in the existing /home/monnett file space.

However, because the /home/monnett file space already exists, when you try to query or restore the doc3 file, Tivoli Storage Manager looks for the file in the /home/monnett file space unless you specify the correct file space name. For example:

```
dsmc query backup "{/home}/monnett/*"
dsmc restore {/home}/monnett/doc2
```

**Note:** You must explicitly specify the file space name only when there can be more than one resolution to the file specification.

For example, if the following file spaces exist in storage:

```
/home
/home/monnett
/home/monnett/project1
/home/monnett/project1/planning
```

then enter:

```
dsmc query backup "/home/monnett/project1/planning/*"
```

Tivoli Storage Manager looks for files only in the /home/monnett/project1/planning file space, even if one or more of the other file spaces contains a path with the same name. But, when you enter one of the following:

```
dsmc query backup "{/home}/monnett/project1/planning/*"
dsmc query backup "{/home/monnett}/project1/planning/*"
dsmc query backup "{/home/monnett/project1}/planning/*"
```

Tivoli Storage Manager looks for files only in the /home file space, the /home/monnett file space, or the /home/monnett/project1 file space, depending on which form you use.

## How special file systems are handled

Special file systems contain dynamic information generated by the operating system; they contain no data or files. The Tivoli Storage Manager client ignores special file systems and their contents. Special file systems include the following:

- the /proc file system on most of the UNIX platforms
- the /dev/fd file system on Solaris
- the /dev/pts on Linux

## 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** from the Java 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 157 for more information on management classes, how an Authorized User assigns management classes to files, and storage management policies.

## Understanding how symbolic links are handled

A UNIX symbolic link is a file that contains a pointer to another file or directory. Tivoli Storage Manager handles symbolic links differently than it does regular files and directories. In some operations, such as a backup, only the path information that the symbolic link contains is backed up. In other operations, such as archive, the file to which the symbolic link points is archived, but under the name of the symbolic link. For more information on how symbolic links are handled during an archive operation, see “Archsymlinkasfile” on page 196.

### *Incremental backup*

— When you run an incremental backup, Tivoli Storage Manager backs up only the path information to a file or directory to which a symbolic link points. The contents of the file or the contents of files in the directory are not backed up.

### *Selective backup*

— When you run a selective backup on a symbolic link that points to a file, Tivoli Storage Manager backs up only the path information to that file. The contents of the file are not backed up.

*Restore* — When you restore a symbolic link that originally pointed to a file, the symbolic link is restored, whether or not the file it points to still exists. If you restore a symbolic link that originally pointed to a directory:

- Without the files in the directory (for example, the /home/gillis/symdir directory), and the symbolic link does not exist on your file system, nothing is returned.
- Along with the files in the directory (for example, /home/gillis/symdir/), and the symbolic link does not exist on your file system, Tivoli Storage Manager builds the directory on your workstation and puts the files in that directory. If the *subdir* option is set to yes, Tivoli Storage Manager recursively restores all subdirectories of the directory.
- And the symbolic link already exists on your workstation, the result depends on how the *followsymbolic* option is set; if it is set to:

*Yes* — The symbolic link is restored and overwrites the directory on your workstation. If the *followsymbolic* option is set to yes, a symbolic link can be used as a virtual mount point.

*No* — Tivoli Storage Manager displays an error message. No is the default.

**Note:** On UNIX systems, when a symbolic link is created its modification time is set to current system time and can not be

changed. When restoring a symbolic link, its modification date and time is set to the date and time of the restore, not to the date and time of the symbolic link when it was backed up. As a result, Tivoli Storage Manager will back up the symbolic link again during the next incremental backup because its modification time changed since the last backup.

The following table shows symbolic link backup and restore functions along with the action taken:

*Table 38. Symbolic link management table for backup and restore*

| Function                                                                                                                     | Action taken                                                                                                                                                                                          |
|------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Selective backup of a file.                                                                                                  | Backs up the symbolic link only, the file is not backed up.                                                                                                                                           |
| Selective backup of a directory.                                                                                             | Backs up the directory only, the symbolic link is not backed up.                                                                                                                                      |
| Incremental backup with <i>subdir=no</i> .                                                                                   | Backs up the symbolic links only, files and directories pointed to are not backed up.                                                                                                                 |
| Incremental backup with <i>subdir=yes</i> .                                                                                  | Backs up the symbolic links and directories and files they point to.                                                                                                                                  |
| Restore a symbolic link that points to a file.                                                                               | The symbolic link is restored, regardless of whether the file the symbolic link points to still exists.                                                                                               |
| Restore a symbolic link that points to a directory.                                                                          | The symbolic link is restored, regardless of whether the directory the symbolic link points to still exists.                                                                                          |
| Restore a symbolic link that points to a directory with <i>subdir=yes</i> , the directory still exists.                      | The symbolic link and files in the directory and subdirectories are restored.                                                                                                                         |
| Restore a symbolic link that points to a directory with <i>subdir=yes</i> , the directory and symbolic link no longer exist. | A directory is created in the directory in which the symbolic link resides and all files and subdirectories are restored to that directory; the symbolic link name is used as the new directory name. |

## 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 files will also be hard-linked even if neither of the files exists at the time of restore, if both of the files are restored together in a single command. The one exception to this procedure occurs if you back up two files that are hard-linked and then break the connection between them on your workstation. If you restore the two files from the server, Tivoli Storage Manager will respect the current file system and not re-establish the hard link.



**Attention:** 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 sparse files are handled

Sparse files do not have disk space allocated for every block in the whole address space, leading to holes within the file. Holes are detected by their content, which is always zeros, and these zeros take up space. The default is to restore the sparse file without the holes, which would leave more free disk space. The Tivoli Storage Manager client detects sparse files during a backup operation and marks them as sparse on the Tivoli Storage Manager server. See “Makesparsefile” on page 285 for more information.

Tivoli Storage Manager backs up a sparse file as a regular file if client compression is off. See “Compression” on page 211 for more information.

## 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.

**Note:** On Solaris and HP-UX, the *nfstimeout* option can fail if the NFS mount is hard. If a hang occurs, deactivate the *nfstimeout* option and mount the NFS file system soft mounted, as follows:

```
mount -o soft,timeo=5,retry=5 machine:/filesystem /mountpoint
```

The parameters are defined as follows:

**soft** Generates a soft mount of the NFS file system. If an error occurs, the `stat()` function returns with an error. If the option `hard` is used, `stat()` does not return until the file system is available.

**timeo=n**

Sets the time out for a soft mount error to *n* seconds

**retry=n**

Set the internal retries and the mount retries to *n*, the default is 10000.

## Backing up opened files

Tivoli Storage Manager looks for files that have changed between the start and the completion of the file's backup. Some files on your system might be in use, or open, when you try to back them up. Because an open file can change, a backup action might not reflect the correct contents of the file at a given time.

Consider if a file is important. Can you build the file again? If the file is not important, you might not want to back up the file. Or, if the file is important, a root user on your workstation can ensure the file is closed before backup.

If your backups run on a schedule, a root user can use the *preschedulecmd* option to enter a command to close the file. For example, if the open file is a database, use the database's **quiesce** command to shut down the database. A root user can use the *postschedulecmd* option to restart the application that uses the file after the backup completes. If you are not using a schedule for the backup, ensure that you close the application that uses the file before you start the backup.

Tivoli Storage Manager can back up the file even if it is open and gets changed during the backup. This is only useful if the file will be usable even if it changes during backup. To back up these files, assign the files a management class with the serialization *dynamic* or *shared dynamic*. See "Selecting a management class for files" on page 162 and "Displaying information about management classes and copy groups" on page 159 for information.

## 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 "/home/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.

**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>*</i> (Asterisk) | Zero or more characters that match all files:                                             |
|---------------------|-------------------------------------------------------------------------------------------|
| <i>*.cpp</i>        | With a <i>cpp</i> extension                                                               |
| <i>hm*.*</i>        | Starting with <i>hm</i> , regardless of extension, but must have the <i>'.'</i> character |
| <i>hm*</i>          | Starting with <i>hm</i> , whether an extension exists or not                              |

|                                                  |                                                                                 |
|--------------------------------------------------|---------------------------------------------------------------------------------|
| <b>*</b> ( <i>Asterisk</i> )                     | <b>Zero or more characters that match all files:</b>                            |
| <b>*h*.*</b>                                     | With an h somewhere in the file name, regardless of extension, but must have .  |
| <b>?</b> ( <i>Question mark</i> )                | <b>One character that matches all files with:</b>                               |
| <b>?.*</b>                                       | The extension cpp with one, and only one, character in the file name            |
| <b>hm?.cpp</b>                                   | Three-character names beginning with hm and that have the cpp extension         |
| <b>* ?</b> ( <i>Asterisk and question mark</i> ) | <b>Asterisk and question mark combinations matching:</b>                        |
| <b>??hm.*</b>                                    | 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

Use Tivoli Storage Manager to restore backup versions of specific files, a group of files with similar names, or entire directories. Select the files you want to restore using file specification (file path, name, and extension), a directory list, or a subdirectory path to a directory and its subdirectories. UNIX socket files are skipped during restore, including socket files that were backed up with earlier versions of Tivoli Storage Manager.

All client restore procedures in this chapter also apply to the Web client, except the following:

- Estimate
- Preferences editor

See “Starting a Web client session” on page 72 for information on starting the Web client.

Table 39 identifies tasks described in this chapter:

*Table 39. Restore: Primary tasks*

| <b>Task</b>                                                | <b>Page</b> |
|------------------------------------------------------------|-------------|
| Restoring data using the GUI                               | 122         |
| Restoring data using the command line                      | 122         |
| Performing point-in-time restores                          | 126         |
| Restoring an image                                         | 127         |
| Restoring an image                                         | 127         |
| Restoring data from a backupset                            | 130         |
| Restoring NAS file systems                                 | 134         |
| Restoring the WebSphere Application Server (WAS)           | 136         |
| Authorizing another user to restore or retrieve your files | 138         |
| Restoring or retrieving files from another client node     | 139         |
| Restore or retrieve files to another workstation           | 139         |
| Restoring a disk in case of disk loss                      | 140         |
| Deleting file spaces                                       | 140         |

---

### 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 on the **View** menu → **Display active files only** item. If you try to restore more than one version at a time, only the active version is restored.

On the Tivoli Storage Manager command line, use the *inactive* option to display both active and inactive objects. See “Inactive” on page 267 for more information.

---

## Restoring data using the GUI

To restore backup versions of individual files or subdirectories:

1. Click **Restore** from the main window. The Restore window appears.
2. Expand the directory tree. Select the selection boxes next to the files or directories you want to restore. To search or filter files, click the **Find** icon on the tool bar.
  - To search:*
    - a. Enter your search criteria in the Find Files (Restore) window.
    - b. Click the **Search** button. The Matching Files (Restore) window appears.
    - c. Click the selection boxes next to the files you want to restore and close the Matching Files (Restore) window.
  - To filter:*
    - a. Enter your filter criteria in the Find Files (Restore) window.
    - b. Click the **Filter** button. The Restore window displays the filtered files.
    - c. Click the selection boxes next to the filtered files or directories you want to restore.
3. To modify specific restore options, click the **Options** button. Any options you change are effective during the current session *only*.
4. Click **Restore**. The Restore Destination window appears. Enter the information in the Restore Destination window.
5. Click **Restore**. The Restore **Task List** window displays the restore processing status.

---

## Restoring data using the command line

Use the **restore** command to restore files. See “Restore” on page 466 for more information about the **restore** command. Table 40 on page 123 shows examples of using the **restore** command to restore objects from Tivoli Storage Manager server storage. See “Restore” on page 466 for additional examples.

Table 40. 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 281 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 308 and “Inactive” on page 267 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 356 for more information about the <i>subdir</i> option.                                                                                                                  |
| Restore all files in the /Users/gordon/Documents directory to their state as of 1:00 PM on August 17, 2003.           | <code>dsmc restore -pitd=8/17/2003 -pitt=13:00:00 /Users/gordon/Documents/</code>                     | See “Pitdate” on page 309 and “Pittime” on page 310 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. | <code>dsmc restore "/Users/mike/Documents/*.bak" /Users/mike/projectn/</code>                         | 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.                                          | <code>dsmc restore -filelist=/Users/user2/Documents/restorelist.txt /Users/NewRestoreLocation/</code> | See “Filelist” on page 250 for more information about restoring a list of files.                                                                                                                                                                                                                                                                                                                                            |

## 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 /home file system, enter:

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

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

For example, you could enter these commands:

```
dsmc restore /home/monnett/ -subdir=yes -replace=all -tapeprompt=no
dsmc restore /home/gillis/ -subdir=yes -replace=all -tapeprompt=no
dsmc restore /home/stewart/ -subdir=yes -replace=all -tapeprompt=no
```

You can also use the *quiet* option with the **restore** commands to save processing time. However, you will not receive informational messages for individual files.

**Note:** If you already have the appropriate values set for the *subdir*, *replace*, *tapeprompt*, and *quiet* options in your client user options file, you do not need to include those options in the commands.

When you enter multiple commands to restore your files, you must specify a unique part of the file space in each **restore** command. Be sure you do not use any overlapping file specifications in the commands.

To display a list of the directories in a file space, use the **query backup** command. For example:

```
dsmc query backup -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 how much memory you have and network utilization.

The speed at which you can restore the files also depends on how many tape drives are available on the server, and whether your administrator is using collocation to keep file spaces assigned to as few volumes as possible.

For example, if /Users/user1 and /Users/user2 are on the same tape, the restore for /Users/user2 must wait until the restore for /Users/user1 is complete. However, if /Users/user3 is on a different tape, and there are at least two tape drives available, the restore for /Users/user3 can begin at the same time as the restore for /Users/user1.

| Set the system ulimit values to unlimited (-1) if restoring very large (2 GB) files  
| with HSM or the backup-archive client. The Tivoli Storage Manager client can  
| restore these large files with enough system resources. If the ulimits are set to  
| lower values, there might be restore failures.

## No query restore

The no query restore process is outlined below.

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*, or *todate*, the client uses a *no query restore* method for restoring files and directories from the server. This method is called *no query restore* because instead of querying the server for each object to be restored, a single restore request is sent to the server. In this case, the server returns the files and directories to the client without



further action by the client. The client merely accepts the data coming from the server and restores it to the destination named on the **restore** command.

Using the command line client, an example of an unrestricted wildcard command would be: An example of an unrestricted wildcard file specification would be:

```
/home/mydocs/2004/*
```

An example of a restricted wildcard file specification would be:

```
/home/mydocs/2004/sales.*
```

## Standard restore process

The standard restore process (also known as classic restore) 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

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 327 for more information.

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 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. 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 466. 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 wish to complete
3. Click the **Restart** button at the bottom of the panel.

---

## 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. You can also perform a point-in-time restore of image backups. For more information see “Backup Image” on page 406.

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 you 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 157 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 restoreable 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 an image

Before you perform an image restore, consider the following:

- Restoring the image of a volume will restore the data to the same state that it was in when you performed your last image backup. Be absolutely sure that you need to restore an image, because it will replace your entire current file system or raw volume with the image on the server.
- Ensure that the volume to which you are restoring the image is at least as large as the image that is being restored.
- The file system or volume you are restoring to must be the same type as the original.
- Ensure that the target volume of the restore is not in use. The client will lock the volume before starting the restore. The client will unlock the volume after the restore completes. If the volume is in use when the client attempts to lock the file system, the restore will fail.
- You cannot restore an image to where the Tivoli Storage Manager client program is installed.
- If you have run progressive incremental backups *and* image backups of your file system, you can perform an incremental image restore of the file system. The process restores individual files after the complete image is restored. The individual files restored are those backed up after the original image. Optionally, if files were deleted after the original backup, the incremental restore can delete those files from the base image. Deletion of files will be performed correctly if the Tivoli Storage Manager server's backup copy group has enough versions for existing and deleted files. Incremental backups and restores can be performed only on mounted file systems, not on raw logical volumes.

- If for some reason a restored image is corrupted, you can use the **fsck** tool to attempt to repair the image.

You can use the *verifyimage* option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log. See “Verifyimage” on page 380 for more information.

If bad sectors are present on the target volume, you can use the *imagetofile* option with the **restore image** command to specify that you want to restore the source image to a file. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume. See “Imagetofile” on page 264 for more information.

## Performing an image restore using the GUI

Use the following procedure to restore an image of your file system or raw logical volume:

1. Click **Restore** from the main window. The Restore window appears.
2. Expand the directory tree.
3. Locate the object in the tree named **Image** and expand it. Click the selection box next to the image you want to restore. You can obtain detailed information about the object by highlighting the object and selecting **View** → **File Details...** from the main window or click the **View File details** button.
4. **(Optional)** To perform an incremental image restore, click the **Options** button to open the Restore Options window and select the **Image plus incremental directories and files** option. If you want to delete inactive files from your local file system, select the **Delete inactive files from local** check box. Click the **OK** button.
5. Click **Restore**. The Restore Destination window appears. The image can be restored to the volume with the mount point from which it was originally backed up. Alternatively, a different volume can be chosen for the restore location.
6. Click the **Restore** button to begin the restore. The **Task List** window appears showing the progress of the restore. The Restore Report window displays a detailed status report.

Considerations:

- You can select **View** → **File Details** from the main window or click the **View File details** button to display the following statistics about file system images backed up by the client:
  - Image Size - This is the volume size which was backed up.
  - Stored Size - This is the actual image size stored on the server. The stored image on the Tivoli Storage Manager server is the same size as the volume capacity.
  - File system type
  - Backup date and time
  - Management class assigned to image backup
  - Whether the image backup is an active or inactive copy
- To modify specific restore options, click the **Options** button. Any options you change are effective during the current session *only*.

- In the Restore Options window, you can choose to restore the image only or the image and incremental directories files. If you choose **Image Only**, you will restore the image from your last image backup only. This is the default.

If you ran incremental-by-date image backup on a volume or image backups on a volume with incrementals, you can choose the **Image plus incremental directories and files** option. If you choose **Image plus incremental directories and files**, you can also select **Delete inactive files from local** to delete the inactive files that are restored to your local file system. If incremental-by-date image backup was the only type of incremental backup you performed on the file system, deletion of files will not occur.

**Attention:** Be absolutely sure that you need to perform an incremental restore because it will replace your entire file system with the image from the server and then restore the files that you backed up using the incremental image backup operation.

## Performing an image restore using the command line

Use the **restore image** command to restore an image using the Tivoli Storage Manager command line client. See “Restore Image” on page 476 “Verifyimage” on page 380, and “Imagetofile” on page 264 for more information.

---

### Restoring an image to file

All LVM volumes contain Logical Volume Control Block (LVCB) on the first sector (512 bytes) of the volume which contains volume specific meta-data and should be preserved by applications using the volume. When performing image backup, Tivoli Storage Manager backs up LVCB, but on restore, it only restores the data (skipping LVCB), in order to preserve the original LVCB of the destination volume. When performing image restore to file, entire volume contents, including LVCB, are restored to file. When you copy the file containing image back onto a LVM volume, you will need to skip LVCB from both the file and destination volume. The following dd command can be used for this purpose.

```
dd if=<filename> of=/dev/<vol> bs=512 skip=1 seek=1
```

The above command sets the block size to 512 bytes, which will make copying very slow. It is better to use bs=1m or similar. Here is a faster (but more complicated) way to copy image data:

1. Save the original first sector to a file:

```
dd if=/dev/<vol> of=firstblk.tmp bs=512 count=1
```

2. Copy the restored image:

```
dd if=<filename> of=/dev/<vol> bs=1m
```

3. Restore the original first sector:

```
dd if=firstblk.tmp of=/dev/<vol> bs=512 count=1
```

## Performing an image restore using the GUI

Use the following procedure to restore an image of your file system or raw logical volume:

1. Click **Restore** from the main window. The Restore window appears.
2. Expand the directory tree.
3. Locate the object in the tree named **Image** and expand it. Click the selection box next to the image you want to restore. You can obtain detailed information about the object by highlighting the object and selecting **View** → **File Details...** from the main window or click the **View File details** button.

4. **(Optional)** To perform an incremental image restore, click the **Options** button to open the Restore Options window and select the **Image plus incremental directories and files** option. If you want to delete inactive files from your local file system, select the **Delete inactive files from local** check box. Click the **OK** button.
5. Click **Restore**. The Restore Destination window appears. The image can be restored to the volume with the mount point from which it was originally backed up. Alternatively, a different volume can be chosen for the restore location.
6. Click the **Restore** button to begin the restore. The **Task List** window appears showing the progress of the restore. The Restore Report window displays a detailed status report.

Considerations:

- You can select **View** → **File Details** from the main window or click the **View File details** button to display the following statistics about file system images backed up by the client:
  - Image Size - This is the volume size which was backed up.
  - Stored Size - This is the actual image size stored on the server. The stored image on the Tivoli Storage Manager server is the same size as the volume capacity.
  - File system type
  - Backup date and time
  - Management class assigned to image backup
  - Whether the image backup is an active or inactive copy
- To modify specific restore options, click the **Options** button. Any options you change are effective during the current session *only*.
- In the Restore Options window, you can choose to restore the image only or the image and incremental directories files. If you choose **Image Only**, you will restore the image from your last image backup only. This is the default.

If you ran incremental-by-date image backup on a volume or image backups on a volume with incrementals, you can choose the **Image plus incremental directories and files** option. If you choose **Image plus incremental directories and files**, you can also select **Delete inactive files from local** to delete the inactive files that are restored to your local file system. If incremental-by-date image backup was the only type of incremental backup you performed on the file system, deletion of files will not occur.

**Attention:** Be absolutely sure that you need to perform an incremental restore because it will replace your entire file system with the image from the server and then restore the files that you backed up using the incremental image backup operation.

## Performing an image restore using the command line

Use the **restore image** command to restore an image using the Tivoli Storage Manager command line client. See “Restore Image” on page 476 for more information.

---

## 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.

Portable media can be used on devices such as a tape, CD-ROM, DVD, and Iomega JAZ or ZIP drives. Current device support information is available at the following Web sites:

- [http://www.ibm.com/software/sysmgmt/products/support/IBM\\_TSM\\_Supported\\_Devices\\_for\\_AIXHPSUNWIN.html](http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_AIXHPSUNWIN.html)
- [http://www.ibm.com/software/sysmgmt/products/support/IBM\\_TSM\\_Supported\\_Devices\\_for\\_iSeries.html](http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_iSeries.html)
- [http://www.ibm.com/software/sysmgmt/products/support/IBM\\_TSM\\_Supported\\_Devices\\_for\\_Linux.html](http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html)

You can restore backupsets from the following locations:

- Directly from the server (You must be a root user to restore an entire backupset from the server, otherwise only files you own are restored.)
- From portable media on a device attached to your client workstation
- From a backupset file on the 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.

#### **Notes:**

1. If you cannot restore a backupset from portable media, 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 device.
2. There is no support in the Tivoli Storage Manager API for the backupset format.
3. To enable the GUI client to restore a backupset on an attached device on a UNIX and Linux standalone workstation, without requiring a server connection, use the *localbackupset* option. See “Localbackupset” on page 282 for more information.
4. Note that the **restore backupset** command supports restore of local backupsets from local media without using the *localbackupset* option.

## **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. The types of data can be



file data or Image data. If you have upgraded from Tivoli Storage Manager Express, some application data is also supported.

**Note:** Image data and application data restore is only available when restoring from the server. It is not supported from a client local backupset restore.

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 nodename 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 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 plus sign (+) 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:** or **Tape name:** from the dropdown list and then enter the tape or 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.

It is now possible to select files from within a backupset if that backupset is located at the server and was generated with a table of contents.
- Click **Restore**. The Restore Destination window appears. Enter the appropriate information.
- Click on **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 such as tape 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. Image restore
  - b. 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
  - c. Application data restore if the server was migrated from the Tivoli Storage Manager Express product

## Restoring backupsets using the command line client

The `restore backupset` command restores a backupset from the server, a local file, or a local tape device. Refer to the `backupsetname` option on the `restore` and `query` commands for more information.

Use the `location` option with the `restore backupset` and `query backupset` 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. Tapes that are generated on the server can be used locally by specifying the `location` option and either the file name or the tape device. Refer to the `backupsetname` option on the `restore` and `query` commands for more information.

Use the `query backupset` command to query a backupset from a local file, tape device, or the Tivoli Storage Manager server. The `query backupset` command displays the backupset name, generation date, retention, and description. See “Query Backupset” on page 447 for more information.

---

## Restoring data to a GPFS 3.1 file system with multiple storage pools

GPFS 3.1 allows you to manage your data using storage pools. A storage pool is a collection of disks or RAID configurations with similar properties that are managed together as a group. The group under which the storage pools are managed together is the file system. The automated placement and management of files on the storage pool level is done by policies. A policy is a set of rules that describes the life cycle of user data, based on the file’s attributes.

When a file is created, the placement policy determines the initial location of the file’s data and assigns the file to a storage pool. All data written to that file is placed in the assigned storage pool. The management policy determines file management operation, such as migration and deletion. The files within a GPFS 3.1 file system are distributed over different storage pools, depending on the enabled placement and migration policies.

During restore, the files are placed on the correct storage pool. Tivoli Storage Manager is not aware of pool-to-pool migrations, so the files are placed on the storage pool from where the backup has taken place. The policy engine replaces the files based on migration policies.

GPFS handles the placement of files after a restore as follows:

- The file is placed according to the storage pool ID that is stored in the extended attributes
- The file is placed based on the current placement policy
- The file is placed in the system pool

If a storage pool ID is stored in the extended attributes of the file, and that storage pool is available, the file is always placed in that storage pool. If the storage pool is not available, the file is placed according to the placement policy. If the placement policy does not match the file, the file is placed in the system pool.

The following restrictions apply:

- Correct placement to storage pools only works with GPFS 3.1 PTF 6 or higher
- The restore of stub files does not work with multiple storage pools, or with files that have ACLs
- Filesets are not supported
- Unlink of Filesets are not allowed
- `ctime` option of GPFS should be set to *no* (default), to prevent unwanted Backup-Archive backups of files after GPFS file migration from pool to pool

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

---

## Restoring NAS file systems

You restore NAS file system images using the Web client or command line interface. For information on how to install and configure the Web client, see “Configuring the Web client” on page 50.

You can restore full or differential NAS file system images that were backed up previously. If you restore a differential image, Tivoli Storage Manager automatically restores the full backup image first, followed by the differential image. It is not necessary for a client node to mount a NAS file system to perform backup or restore operations on that file system.

### Restoring NAS file systems using the Web client

For information on how to install and configure the Web client, see “Configuring the Web client” on page 50. To restore NAS file systems using the Web client GUI:

1. Click the **Restore** from the main window. The Restore window appears.
2. Expand the directory tree if necessary. To expand a node in the tree, click the plus sign (+) next to an object in the tree.

**Notes:**

- a. Nodes shown are those that have been backed up and to which your administrator has authority.
- b. The root node called **Nodes** is not selectable. This node only appears if a NAS plug-in is present on the client machine.

- c. NAS nodes display on the same level as the client workstation's node. Only nodes to which the administrator has authority appear.
3. Expand the NAS node to reveal the Image object.
4. Expand the Image object to display volumes that you can restore. You cannot expand Volume objects.
5. Click the selection boxes next to the volumes under the Image object that you want to restore.

**Notes:**

- a. If you want to restore a NAS image that was backed up on a particular date, click the **Point In Time** button. After you select a date, the last object that was backed up on or prior to that date appears, including any inactive objects.
- b. If you want to display all images (including active images and inactive images), before you select them, select **View** → **Display active/inactive files** from the menu bar.
6. Click **Restore**. The Restore Destination window appears. Enter the information in the Restore Destination window. If you choose to restore to a different destination, you can only restore one volume at a time to a different destination.

**Note:** You can restore NAS file system images to any volume on the NAS file server from which they were backed up. You cannot restore images to another NAS file server.

7. Click **Restore**. The NAS Restore **Task List** window displays the restore processing status and progress bar. If there is a number next to the progress bar, it indicates the size of the restore, if known. After the restore completes, the NAS Restore Report window displays processing details.

**Note:** If it is necessary to close the Web browser session, current NAS operations will continue after disconnect. You can use the **Dismiss** button on the NAS Restore **Task List** window to quit monitoring processes without ending the current operation.

8. (Optional) To monitor processing of an operation, select the **Actions** → **TSM Activities** from the main window.

Considerations:

- Workstation and remote (NAS) backups are mutually exclusive in a Restore window. After selecting an item for restore, the next item you select must be of the same type (either NAS or non NAS).
- Details will not appear in the right-frame of the Restore window for NAS nodes or images. To view information about a NAS image, highlight the NAS image and select **View** → **File Details** from the menu.
- To delete NAS file spaces, select **Utilities** → **Delete Filespaces**. You can delete both workstation and remote objects.

## Restoring NAS file systems using the command line

Table 41 on page 136 lists the commands and options you can use to restore NAS file system images from the command line.

Table 41. NAS options and commands

| Option or command        | Definition                                                                                                                                                                                                                                                                                                         | Page |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <b>query node</b>        | Displays all the nodes for which a particular administrative user ID has authority to perform operations. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the Web client. | 457  |
| <b>query backup</b>      | Use the <b>query backup</b> command with the <i>class</i> option to display information about file system images backed up for a NAS file server.                                                                                                                                                                  | 444  |
| <b>query filesystem</b>  | Use the <b>query filesystem</b> command with the <i>class</i> option to display a list of file spaces belonging to a NAS node.                                                                                                                                                                                     | 449  |
| <b>restore nas</b>       | Restores the image of a file system belonging to a Network Attached Storage (NAS) file server.                                                                                                                                                                                                                     | 479  |
| <b>monitor process</b>   | Displays current backup and restore processes for all NAS nodes for which an administrative user has authority. The administrative user can then select one process to monitor.                                                                                                                                    | 438  |
| <b>cancel process</b>    | Displays current backup and restore processes for all NAS nodes for which an administrative user has authority. From the display, the administrative user can select one process to cancel.                                                                                                                        | 415  |
| <b>delete filesystem</b> | Use the <b>delete filesystem</b> with the <i>class</i> option to display a list of file spaces belonging to a NAS node so that you can choose one to delete.                                                                                                                                                       | 423  |

Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.

**Note:** When you initiate a NAS restore operation using the command line client or the Web client, the server starts a process to initiate, control, and monitor the operation. It might take several moments before you notice progress at the command line client interface because the server must perform mount and other necessary tasks before data movement occurs. The Tivoli Storage Manager command line client might display an Interrupted ... message when the mount occurs. You can ignore this message.

## Restoring the WebSphere Application Server (WAS)

Use the Web client GUI or command line client to restore full or differential image backups of a WebSphere Application Server if the Data Protection for WebSphere Application Server is installed. When you back up a WAS, file spaces are created on the Tivoli Storage Manager server with the naming convention WAS\_INSTANCENAME (for the Application Server file space) and WAS\_ND\_INSTANCENAME (for the Network Deployment Manager file space).

Before you begin:

- It is strongly recommended that you restore data at the Network Deployment Manager node or Application Server node level only, rather than individual files. Restoring data other than at these levels can corrupt your WebSphere installation. See *IBM Tivoli Storage Manager for Application Servers 5.2: Data Protection for WebSphere Application Server Installation and User's Guide*, SC32-9075, for more information.

- You need to take WAS offline before you begin the restore.

To restore the WebSphere Application Server:

1. Click **Restore** from the Web client GUI main window. The Restore window appears.
2. Expand the directory tree if necessary.
3. Expand the WebSphere Application Server node to reveal the WAS\_<INSTANCENAME> and WAS\_ND\_<INSTANCENAME> file spaces.

**Notes:**

- a. The most recent active full or differential backup appears under the WebSphere Application Server tree.
  - b. If only a single instance of WAS was backed up, the WebSphere Application Server tree displays only the WAS\_<NODENAME> and WAS\_ND\_<NODENAME> file spaces.
  - c. If multiple instances of WAS were backed up, the WebSphere Application Server tree displays multiple instances as WAS\_<INSTANCENAME> and WAS\_ND\_<INSTANCENAME>. The INSTANCENAME is a variation of the NODENAME.
4. Click the selection box next to the Application Server (INSTANCENAME) node or Network Deployment Manager (INSTANCENAME) node that you want to restore.

**Notes:**

- a. If you want to restore objects that were backed on a particular date, click the **Point In Time** button. After you select a date, the last object that was backed up on or prior to that date appears, including inactive objects.
  - b. If you want to see all active and inactive WAS backups before you select them, select **View** → **Display active/inactive files** from the menu bar. Inactive WAS backup objects are displayed with an X next to the object in the tree.
5. Click **Restore**. The Restore Task List window displays the restore processing status. If you want to re-access the Task List window after exiting and restarting the Web client, select **Actions** → **TSM Activities** from the menu.

Considerations:

- Workstation, NAS, and WAS restores are mutually exclusive in the Restore window. After selecting an item for restore, the next item you select must be of the same type.
- To display information about a WAS object (from the Restore window), select a WAS object, click **View** → **File Details**.
- If you need to recover a WAS file space to a new machine, you must install WebSphere Application Server, the client, and the Data Protection for WebSphere Application Server on the new machine before you can restore your WAS file space. See “Restore WAS” on page 481 for more information.
- To delete WAS file spaces, select **Utilities** → **Delete Filespaces** from the Restore window the menu bar.

To restore WAS objects from the Tivoli Storage Manager command line, use the **restore was** command. See “Restore WAS” on page 481 for more information.



---

## Restore: Additional considerations

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

### Authorizing another user to restore or retrieve your files

You can authorize another user on the same workstation or a different workstation to restore backup versions or retrieve archive copies of your files. This permits you to share files with other people or with other workstations that you use with a different node name. To authorize a user on another workstation to restore or retrieve your files, the other workstation must be running one of the UNIX clients and must be registered with your Tivoli Storage Manager server.

To authorize another user to restore or retrieve your files:

1. Click **Utilities** → **Node Access List** from the main window. The Node Access List window appears.
2. Click the **Add** button. The Add Access Rule window appears.
3. In the Add Access Rule window, select an item in the Permit Access to field to specify the type of data that the other user can access. You can select either Backed up Objects or Archived Objects.
4. In the Grant Access to Node field, type the node name of the user's host machine that can access your data.
5. In the User field, type the name of the user on a node who can access your data.
6. In the Filespace and Directory field, select the file space and the directory that the user can access. You can select one file space and one directory at a time. If you want to give the user access to another file space or directory, you must create another access rule.
7. If you want to limit the user to specific files in the directory, type the name or pattern of the files on the server that the other user can access in the Filename field. You can make only one entry in the Filename field. It can either be a single file name or a pattern which matches one or more files. You can use a wildcard character as part of the pattern. Your entry must match files that have been stored on the server.
8. For the Java GUI: If you want to give access to all files that match the file name specification within the selected directory including its subdirectories, click **Include subdirectories**.
9. Click the **OK** button to save the access rule and close the Add Access Rule window.
10. The access rule that you created is displayed in the list box in the Node Access List window. When you have finished working with the Node Access List window, click the **OK** button. If you do not want to save your changes, click **Cancel** or close the window.

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 491
- "Query Access" on page 441
- "Delete Access" on page 417



## 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 server, restore the other user's backup versions, or retrieve the other user's archive copies to your local file system.

To display another user's file spaces on the server, restore the other user's backup versions, or retrieve the other user's archive copies to your local drives:

1. Click **Utilities** from the main window.
2. Click **Access Another Node**. The Access Another Node window appears.
3. Type the node name of the user's host machine in the Node name field. Type the user name in the User name field.
4. Click the **Set** button.

If you are using commands, use the *fromnode* and *fromowner* options to indicate the node name and the name of the user who owns the files.

For example, to restore files to one of your own file systems that were backed up from a workstation named Node1 and owned by a user named Ann, enter:

```
dsmc restore -fromn=node1 -fromo=ann "/home/proj/*" /home/gillis/
```

Use the **query filesystem** command to get a list of file spaces (see "Query Filespace" on page 449). For example, to get a list of file spaces owned by Ann on Node1, enter:

```
dsmc query filesystem -fromn=node1 -fromo=ann
```

See "Fromnode" on page 256 for more information about the *fromnode* option. See "Restore" on page 466 for more information about using the *fromnode* and *fromowner* options with the **restore** command. Also see "Retrieve" on page 484 for more information about the **retrieve** command.

## Restore or retrieve files to another workstation

From a different workstation, you can restore or retrieve files you have already backed up from your own workstation. You must know the Tivoli Storage Manager password assigned to your node.

To restore or retrieve files to another workstation, use the *virtualnodename* option to specify the node name of the workstation from which you backed up the files. *Virtualnodename* cannot be set to the hostname of the machine. You can use the *virtualnodename* option when you start Tivoli Storage Manager or you can add the *virtualnodename* option to your client user options file `dsm.opt`. Use the *virtualnodename* option on the **dsmj** command if you are borrowing another user's machine and you do not want to update their client user options file.

Tivoli Storage Manager prompts you for the password for your original node. After you enter the correct password, all file systems from your original workstation appear in the Restore or Retrieve window. You can restore or retrieve files as if you were working on your own workstation.

**Attention:** When you use this method to access files, you have access to all files backed up and archived from your workstation. You are considered a virtual root user.

You can use the *virtualnodename* option in a command. For example, to restore your *projx* files, enter:

```
dsmc restore -virtualnodename=nodeone "/home/monnett/projx/*"
```

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

The considerations for retrieving files are the same as restoring files.

## Restoring a disk in case of disk loss

Tivoli Storage Manager can recover your files only if you can run the client. If the file system that contains the client is lost, you must reinstall the client before you can recover your files. If you also lose the file system that contains the operating system and communication software, you must recover them before you can connect to the server.

To protect yourself against these kinds of losses, you need to put together a set of installation media that you can use to restore your system to a state that lets you contact the server and begin recovering data. The installation media should contain:

1. A bootable operating system that lets you perform basic functions.
2. A correctly configured communication program that lets you establish communications with the server.
3. A client with appropriate customized options files. You can use the command line client to complete this task.

The communication package you use determines what files you need. Consult your operating system and communication software manuals to set up your installation media.

If you also have the Tivoli Space Manager installed on your workstation, your installation media should include the space manager command line client. For information about restoring migrated files, see *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide*, SC32-0148.

**Note:** Your administrator can schedule restore operations which can be very useful when you need to restore a large number of files.

## Deleting file spaces

### 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 420 for more information.

You can delete file spaces using the Tivoli Storage Manager GUI or command line clients. To delete NAS file spaces, use the Web client or command line client.

To delete a file space using the GUI, perform the following steps:

1. Select **Utilities** → **Delete Filespaces** from the main window.
2. Click the selection boxes next to the file spaces you want to delete.
3. Click the **Delete** button. Tivoli Storage Manager prompts you for confirmation before deleting the file space.

You can also delete a file space using the **delete filesystem** command. See “Delete Filespace” on page 423 for more information. Use the *class* option with the **delete filesystem** command to delete NAS file spaces. See “Class” on page 202 for more information.



---

## Chapter 6. Archiving and retrieving your data

Archiving and retrieving files is similar to backing up and restoring files. Many of the windows and concepts are similar. In this chapter, we cover the main archive and retrieve tasks, but where windows and concepts are the same, as for backup and restore, see Chapter 4, “Backing up your data,” on page 79.

All client archive and retrieve procedures in this chapter also apply to the Web client, except the following:

- Estimate
- Preferences editor

See “Starting a Web client session” on page 72 for information on starting the Web client.

Table 42 identifies tasks described in this chapter:

*Table 42. Archive and retrieve: Primary tasks*

| Task                                   | Page |
|----------------------------------------|------|
| Archiving data using the GUI           | 143  |
| Archiving data using the command line  | 144  |
| Deleting archive data                  | 147  |
| Retrieving data using the GUI          | 149  |
| Retrieving data using the command line | 149  |

---

### 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 151 for information on checking and running the schedules available to you. The following sections cover how to archive files without using a schedule.

#### 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 on **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. See “Archive” on page 402 for more information about the **archive** command.

### Associating a local snapshot with a server file space

Use the *snapshotroot* option with the **archive** command in conjunction with a third-party application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. The *snapshotroot* option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot. See “Snapshotroot” on page 351 for more information.

Table 43 shows examples of using the **archive** command to archive objects. See “Archive” on page 402 additional examples.

Table 43. Command line archive examples

| Task                                                                                           | Command                                                 | Considerations                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------------------------------------------------------------------|---------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Archive all files in the /home/proj1 directory with a file extension of .txt.                  | dsmc archive "/home/proj1/*.txt"                        | Use wildcards to archive more than one file at a time.                                                                                                                                                                                                                                                                                                    |
| Archive all files in the /home/jones/proj/ directory and delete the files on your workstation. | dsmc archive /home/jones/proj/<br>-deletefiles          | Retrieve the archived files to your workstation whenever you need them again. See “Deletefiles” on page 218 for more information about the <i>deletefiles</i> option.                                                                                                                                                                                     |
| Archive the /home/jones/h1.doc and /home/jones/test.doc files.                                 | dsmc archive /home/jones/h1.doc<br>/home/jones/test.doc | If you specify the <i>removeoperandlimit</i> option with the <b>archive</b> 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 324 for more information about this option. |

Table 43. Command line archive examples (continued)

| Task                                                                                                                                                                                                                                                            | Command                                                                               | Considerations                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Archive a list of files in the /home/avi/filelist.txt file.                                                                                                                                                                                                     | <code>dsmc archive -filelist=/home/avi/filelist.txt</code>                            | Use the <i>filelist</i> option to process a list of files. See “Filelist” on page 250 for more information.                                                                                                                                          |
| Archive the /home/jones/ch1.doc file and assign a description to the archive.                                                                                                                                                                                   | <code>dsmc archive /home/jones/ch1.doc -description="Chapter 1, first version"</code> | If you do not specify a description with the <b>archive</b> command, the default is <code>Archive Date:x</code> , where <i>x</i> is the current system date. See “Description” on page 219 for more information about the <i>description</i> option. |
| Archive all of the files in the /home/jones/proj/ directory and its subdirectories.                                                                                                                                                                             | <code>dsmc archive /home/jones/proj/ -subdir=yes</code>                               | See “Subdir” on page 356 for more information about the <i>subdir</i> option.                                                                                                                                                                        |
| Use the <i>v2archive</i> option with the <b>archive</b> command to archive only files in the /home/relx/dir1 directory, but not the relx or dir1 directories.                                                                                                   | <code>dsmc archive "/home/relx/dir1/" -v2archive</code>                               | Tivoli Storage Manager archives only files in the /home/relx/dir1 directory. Directories that exist in the path are not processed. See “V2archive” on page 378 for more information about the <i>v2archive</i> option.                               |
| Use the <i>archmc</i> option with the <b>archive</b> 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 /home/plan/proj1/budget.jan</code>                 | See “Archmc” on page 195 for more information about the <i>archmc</i> option. See Chapter 8, “Understanding storage management policies,” on page 157 for more information about management classes.                                                 |
| Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, archive the /usr/dir1/sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name /usr. | <code>dsmc archive /usr/dir1/sub1/ -subdir=yes -snapshotroot=/snapshot/day1</code>    | Tivoli Storage Manager considers the <i>snapshotroot</i> value as a file space name. See “Snapshotroot” on page 351 for more information.                                                                                                            |

## 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 machine responsible for performing the archive can change over time, such as with a cluster. The *asnodename* option also allows data to be restored from a different system than the one which performed the backup. Use the *asnodename* option with the appropriate command to back up, archive, restore, and retrieve data under the target node name on the Tivoli Storage Manager server. This support is only available with Tivoli Storage Manager Version 5.3 and higher. To enable this option, follow these steps:

- Install the backup-archive client on all nodes in a 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.

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

1. Specify *encryptkey=save* in the options file.
2. Back up at least one file with *asnode=ProxyNodeName* to create a local encryption key on each agent node in the multiple node environment.

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

1. Specify *encryptkey=prompt* in the options file.
  2. Ensure that users of the agent nodes in the multiple node environment are using the same encryption key.
- If you change the encryption key, you must repeat the previous steps.
  - Use the same encryption key for all files backed up in the shared node environment.

Follow these steps to enable multinode operation from the GUI:

1. Verify that the client node has proxy authority to a target node (or authorized to act as the target node) using the **query proxynode** administrative client command.
2. Select **Edit** → **Preferences** to open the preferences window.
3. Select the **General** tab and fill in the **As Node Name** field with the name of the proxy authorized target node.
4. Click **Apply** and then **OK** to close the preferences window.

Follow these steps to verify that your client node is now accessing the server as the target node:

1. Open the tree window and check that the target node name specified by the **As Node Name** field appears, or
2. Verify the target node name in the **Accessing As Node** field in the **Connection Information** window.

To return to single node operation, delete the **As Node Name** from the **Accessing As Node** field in the **General** → **Preferences** tab.

#### **Suggestions:**

- All agent nodes in the multiple node environment should be of the same platform type.
- Do not use target nodes as traditional nodes. Use them only for multiple node processing.

#### **Restrictions enforced within a proxied session:**

- You cannot perform a system object, system state or system services backup or restore.

- You cannot access another node (either from GUI drop down or use of the *fromnode* option).
- You cannot use the *clusternode* option.
- You cannot perform NAS backup or restore.

See “Asnodename” on page 197 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 or Web client main menu. Your authority status is provided in the **Delete Archive Files** field.

To delete archive copies from the Tivoli Storage Manager GUI or Web client:

1. Select **Delete Archive Data** from the **Utilities** menu. The Archive Delete window appears.
2. Expand the Directory tree by clicking the plus sign (+) or folder icon next to the object you want to expand. Objects on the tree are grouped by archive package description.
3. Click the selection boxes next to objects that you want to delete. If you want to estimate the amount of time it takes to process your files and directories, click the **Estimate** button.
4. Click the **Delete** button. Tivoli Storage Manager prompts you for confirmation before deleting archive selections. The Archive Delete Task List window displays the progress of the delete.

To delete archive copies using the Tivoli Storage Manager command line client, use the **delete archive** command. See “Delete Archive” on page 418 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 a file, Tivoli Storage Manager saves standard UNIX access permissions assigned to the file. Depending on your operating system, it also saves extended permissions. For example, for files on an AIX workstation, Tivoli Storage Manager saves access control lists.

If you are a user, and you archive a file to which you have read access, you own the archived copy of the file. You are the only user who can retrieve the archived file unless you grant access to another user.

### Understanding how symbolic links are handled

When you archive a symbolic link, Tivoli Storage Manager archives the file to which the symbolic link points. It does not archive path information for the directory.

If you archive a symbolic link that points to a directory, Tivoli Storage Manager archives the files contained in the directory (and its subdirectories if the *subdir* option is set to yes) under the name of the symbolic link.

Use the *archsymblinkasfile* option to specify whether Tivoli Storage Manager archives the symbolic link and the file or directory it points to, or the symbolic link only. See “Archsymblinkasfile” on page 196 for more information.

Table 44 shows symbolic link archive and retrieve functions and the action taken:

*Table 44. Symbolic link management table for archive and retrieve*

| Function                                                                                              | Action taken                                                                                                                                                                                        |
|-------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Archive of a file link.                                                                               | Archives the file to which the symbolic link points.                                                                                                                                                |
| Archive of a directory link.                                                                          | Archives the directory and its contents.                                                                                                                                                            |
| Archive of a file with <i>subdir=yes</i> .                                                            | Archives the file, directory path and all like-named files in the subtree.                                                                                                                          |
| Archive of a directory with <i>subdir=yes</i> .                                                       | Archives the directory, its contents, and contents of subdirectories.                                                                                                                               |
| Archive of a symbolic link that points to a file or directory that does not exist.                    | Archives the symbolic link.                                                                                                                                                                         |
| Retrieve a symbolic link that points to file; the file and link exist.                                | Replaces the file if <i>replace=y</i> is set.                                                                                                                                                       |
| Retrieve a symbolic link that points to file; the symbolic link no longer exists.                     | Retrieves the file replacing the file name with the symbolic link name and places it in the directory where the symbolic link resided.                                                              |
| Retrieve a symbolic link that points to a directory; the symbolic link and directory no longer exist. | A directory is created in the directory where the symbolic link resides, and all files and subdirectories are restored to that directory. The symbolic link name is used as the new directory name. |
| Retrieve a symbolic link that points to a directory; the symbolic link and directory still exist.     | Tivoli Storage Manager will not retrieve as long as the symbolic link exists.                                                                                                                       |

## 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. See “Authorizing another user to restore or retrieve your files” on page 138, “Restoring or retrieving files from another client node” on page 139, and “Restore or retrieve files to another workstation” on page 139.

## Retrieving data using the GUI

To retrieve an archived file:

1. Click **Retrieve** from the client Java GUI main window. The Retrieve window appears.
2. Expand the directory tree by clicking the plus sign (+) or the folder icon next to an object you want to expand. To search or filter files, click the **Search** icon from the tool bar.

### 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*. For the Motif GUI: To estimate the transfer time for your archived selections, click the **Estimate** button. The Java GUI does not have an **Estimate** button.
  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.

## Retrieving data using the command line

You *retrieve* a file when you want to return an archive copy from the server to your workstation. You can retrieve a single file, a group of files, or all the files in a directory or subdirectory. When you retrieve a file, Tivoli Storage Manager sends you a copy of that file. The archived file remains in storage.

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

Table 45 shows examples of using the **retrieve** command. See “Retrieve” on page 484 for additional examples, and detailed information about the **retrieve** command.

Table 45. Command line examples of retrieving archives

| Task                                                                                                              | Command                                             | Considerations                                                                           |
|-------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|------------------------------------------------------------------------------------------|
| Retrieve the /home/jones/h1.doc file to its original directory.                                                   | dsmc retrieve /home/jones/h1.doc                    | If you do not specify a destination, the files are retrieved to their original location. |
| Retrieve the /home/jones/h1.doc file under a new name and directory.                                              | dsmc retrieve /home/jones/h1.doc /home/smith/h2.doc | None                                                                                     |
| Retrieve all files from the /home/jones directory that end with the characters .bak to the /home/smith directory. | dsmc retrieve "/home/jones/*.bak" /home/smith/      | None.                                                                                    |

Table 45. Command line examples of retrieving archives (continued)

| Task                                                                                                  | Command                                                                                | Considerations                                                                                                                                                                                                                    |
|-------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Retrieve the /home/jones/ch1.doc file and assign a description to the retrieve.                       | <code>dsmc retrieve /home/jones/ch1.doc -description="Chapter 1, first version"</code> | If you do not specify a description with the <b>retrieve</b> command, the default is Retrieve Date:x, where x is the current system date. See "Description" on page 219 for more information about the <i>description</i> option. |
| Use the <i>pick</i> option to display a list of archives from which you can select files to retrieve. | <code>dsmc retrieve "/home/jones/*" -pick</code>                                       | See "Pick" on page 308 for more information about the <i>pick</i> option.                                                                                                                                                         |
| Retrieve a list of files specified in the retrievelist.txt file to their original directory.          | <code>dsmc retrieve -filelist=/home/dir2/retrievelist.txt</code>                       | See "Filelist" on page 250 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 Java GUI or Web client directory tree. 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 439.

For information on the various management class attributes used to manage your archives, see "Displaying information about management classes and copy groups" on page 159. See "Assigning a management class to files" on page 163 for information about using the include-exclude options list.

---

## Chapter 7. Automating tasks

This chapter applies to the Authorized User only. Root authorization is only required when updating the `/etc/inittab` and `/etc/rc` files.

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 182 for more information.

### Notes:

1. The schedule start time is based on the time of day of the server and the client.
2. As part of the initial sign-on greeting, the command line client displays the current time at both the server and the client.
3. Install the command line client and ensure the communication software is running before you start the client scheduler.

---

## Specifying scheduling options

You can modify scheduling options in the client system options file (`dsm.sys`) or in the graphical user interface. However, if your administrator specifies a value for these options, that value overrides the value in your client.

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 182.

---

## 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 152 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, **WEEKLY\_INC**, starts a weekly incremental backup in the /proj file system.

The schedule name, **DAILY\_INC**, starts a daily incremental backup. The next incremental backup 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 153 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

When you run the **schedule** command in the foreground, your screen displays output from the scheduled commands. Output is also directed to the `dsmsched.log` file in the current directory unless you change the path and file name using the *schedlogname* option.

When you run the **schedule** command in the background, output from scheduled commands is directed to the `dsmsched.log` file in the current directory, or to the path and file name that you specified. Please note that the `dsmsched.log` cannot be a symbolic link.

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 *shedlogretention* option or specify a maximum size using the *shedlogmax* option. See "Specifying scheduling options" on page 151 for more information.

---

## 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 *shedcmduser* server option.

The *shedcmduser* 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 *shedcmduser* 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 *shedcmddisabled* 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 *shedcmddisabled* 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 331 for more information.

You can use the *shedrestretrdisabled* option to prevent the Tivoli Storage Manager Server administrator from executing restore or retrieve schedule operations. See "Schedrestretrdisabled" on page 338 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 354 for more information.

You can use the *srvprepostsnapdisabled* option to prevent the Tivoli Storage Manager Server administrator from executing pre-snapshot and post-snapshot

commands when performing scheduled image snapshot backup operations. See “Srvprepostsnapdisabled” on page 355 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 dserror.log file (and, for scheduled events, the dsmsched.log file).

For a description of the return codes and their meanings, see Table 46

Table 46. 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"><li>• The file satisfies an entry in an exclude list.</li><li>• The file was in use by another application and could not be accessed by the client.</li><li>• The file changed during the operation to an extent prohibited by the copy serialization attribute. See “Copy serialization” on page 161.</li></ul>                                                           |
| 8    | The operation completed with at least one warning message. For scheduled events, the status will be <i>Completed</i> . Review dserror.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 dserror.log file (and dsmsched.log file for scheduled events) to determine what error messages were issued and to assess their impact on the operation. As a general rule, this return code means that the error was severe enough to prevent the successful completion of the operation. For example, an error that prevents an entire file system from being processed yields return code 12. When a file is not found the operation yields return code 12. |

Table 46. Client return codes and their meanings (continued)

| Code         | Explanation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>other</i> | <p>For scheduled operations where the scheduled action is <code>COMMAND</code>, 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> <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 "/home/devel/*" -subdir=yes
incremental "/home/devel/TestDriver/*" -subdir=yes
archive "/home/plan/proj1/*" -subdir=yes
```

If the first command completes with return code 0; the second command completed with return code 8; and the third command completed with return code 4, the return code for the macro will be 8.

Also see "Handling return codes from `preschedulecmd` and `postschedulecmd` scripts" on page 502.

---

## 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. You can associate (or *bind*) your data to these policies; then when the data is backed up or archived, it is managed according to policy criteria. Policy criteria include a policy domain, a policy set, a management class, and a copy group .

Policies determine:

- Whether a file is eligible for backup or archive services.
- How many backup versions to keep.
- How long to keep inactive backup versions and archive copies.
- Where to place the copies in storage.
- For incremental backup, policies also determine:
  - How frequently a file can be backed up.
  - Whether a file must change before it is backed up again.

If you have the Tivoli Space Manager client installed, your administrator also defines rules that determine whether files are eligible for migration from your local file systems to storage.

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.

**Note:** If you have the Tivoli Space Manager client installed, it can also contain specific requirements for migrating files to storage.

Most of the work you do with storage management policies is with management classes. You must associate (or *bind*) each file and directory that you back up and each file that you archive with a management class:

- If you do not associate 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 162 for more information. In your client system options file (*dsm.sys*), you can associate directories with a management class, using the *dirmc* option. See “Selecting a management class for directories” on page 164 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:

- Whether a file that has changed since the last backup is backed up again.
- 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:

- Where the server stores backup versions of your files and directories.
- How many backup versions the server keeps of your files and directories.
- How long the server keeps backup versions of your files and directories.
- How long the server keeps inactive backup versions.

- How long the last remaining inactive version of a file is kept.

An *archive copy group* contains attributes that control:

- Whether a file is archived if it is in use
- Where 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 166.

---

## 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 47 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 47. 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  |



Table 47. Default values in the standard management class (continued)

| Attribute             | Backup default | Archive default |
|-----------------------|----------------|-----------------|
| 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. This management class attribute is ignored during a journal-based backup.

Copy frequency works with the *mode* parameter. For example, if frequency is *zero* (0) and mode is *modified*, a file or directory is backed up *only if* it changed since the last incremental backup. If frequency is *zero* (0) and mode is *absolute*, a file is backed up every time you run an incremental backup against it. This attribute is not checked for selective backups.

For archive copy groups, copy frequency is always CMD (command). There is no restriction on how often you archive a file.

## Versions data exists

The *Versions Data Exists* attribute specifies the maximum number of different backup versions retained for files and directories currently on your workstation. 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 directories that you erased from your workstation. This parameter is ignored as long as the file or directory remains on your workstation.

If you erase the file or directory, 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 directory is retained. If *Nolimit* is specified, the last version is retained indefinitely.

This parameter goes into effect during the next incremental backup after a file is deleted from the client machine. 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 directory must not be modified during a backup or archive. If the object is changed during a backup or archive attempt, it is not backed up or archived.  
**Note:** During an image backup, the static copy serialization value is no longer controlled by the server management class, but is instead controlled directly from the client, using the *imagetype* option. See “Imagetype” on page 265 for more information.
- **Shared static.** A file or directory must not be modified during backup or archive. Tivoli Storage Manager attempts to perform a backup or archive as many as four additional times, depending on the value specified on the *changingretries* option in your *dsm.sys* file. If the object is changed during every backup or archive attempt, it is not backed up or archived.
- **Dynamic.** A file or directory is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.  
**Note:** During an image backup, the dynamic copy serialization value is no longer controlled by the server management class, but is instead controlled directly from the client, using the *imagetype* option. See “Imagetype” on page 265 for more information.
- **Shared dynamic.** A file or directory is backed up or archived regardless of whether it changes during a backup or archive. Tivoli Storage Manager attempts to perform a backup or archive as many as four additional times, depending on the value specified on the *changingretries* option in your client system options file without the file changing during the attempt. The file is backed up or archived on the last try even if it has changed.

**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 directory is considered for incremental backup regardless of whether it changed or not since the last backup. Tivoli Storage Manager does not check the mode for selective backups. The value for this parameter can be one of the following:

- **Modified.** The file is considered for incremental backup *only if* it has changed since the last backup. A file is considered changed if any of the following are true:
  - The date or time of the last modification is different.
  - The file size is different.
  - If only the file meta-data changes (such as access permissions), but the file data does not change, Tivoli Storage Manager might back up only the meta-data.
  - The file owner is different.
  - The file permissions are different.
- **Absolute.** The file is considered for incremental backup regardless of whether it changed since the last backup. For archive copy groups, the mode is always *absolute*, indicating that a file is archived regardless of whether it changed since the last archive request.

## Copy destination

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 directories, use the *dirmc* option in your *dsm.sys*. See “Dirmc” on page 222 for more information.

You can assign a management class for a file or file group by using an **include** statement in your client systems options (*dsm.sys*) file or the include-exclude file specified by the *incl excl* option. 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 /home/jones/costs/* budget
```

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

```
include /* managall
```

The example below demonstrates how to use a management class:

```
exclude /*.sno
include /home/winter/*.*.ice mcweekly
include /home/winter/december/*.*.ice mcdaily
include /home/winter/january/*.*.ice mcmonthly
include /home/winter/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 with an extension of `ice` in the `/home/winter/january` directory is assigned to the management class, **mcmonthly**.
3. Any file with an extension of `ice` in the `/home/winter/december` directory is assigned to the management class, **mcdaily**.
4. Any other files with an extension of `ice` in any directory under `/home/winter` are assigned to the management class, **mcweekly**.
5. Any file with an extension of `sno` (except `/home/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 269 for more information about the *include* option.

When you archive a file using the graphical user interface, you can select a different management class to override the management class assigned to the file.

---

## 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 `/home/jones/budget.jan` with the management class **ret2yrs**, you would enter:

```
dsmc archive -archmc=ret2yrs /home/jones/budget.jan
```

---

## Selecting a management class for directories

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

If there is more than one management class with the longest RETONLY setting, 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 directories.

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 or Web Client: Select **View Policy Information** from the **Utilities** menu.
- Command line client: Run **dsmc query mgmtclass -detail**.

**Note:** During expiration processing on a Tivoli Storage Manager server, if an archived directory is eligible for expiration, the server checks if any existing archived files require the archived directory to remain. If so, the archived directory is not expired and the Tivoli Storage Manager client updates the insert date on the archived directory to ensure that the directory is not expired before the files under it.

---

## 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 493 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 187 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`) or on the command line 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
- Command processing options
- Authorization options
- Error processing options
- Transaction processing option
- Web client options
- Diagnostics options

See Chapter 2, “Configuring Tivoli Storage Manager,” on page 41 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 61 on page 188.

---

### Communication options

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

For UNIX and Linux use one of the following communication protocols:

- TCP/IP
- Shared Memory (AIX, HP-UX, Linux, and Solaris global and non-global zones only)

Use the *commmethod* option to specify the communication protocol. For more information, see “Commmethod” on page 206.

You can also use the *lanfreecommmethod* option to specify the communication protocol in a SAN environment. See “Lanfreecommmethod” on page 276 for more information.

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.

If you plan to back up an NFS system, see “Nfstimeout” on page 295.

Table 48. TCP/IP options

| Option                          | Description                                                                                                                                                                                        | Page |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>httpport</i>                 | Specifies a TCP/IP port address for the Tivoli Storage Manager Web client.                                                                                                                         | 262  |
| <i>lanfreectpport</i>           | Specifies the TCP/IP port number where the Tivoli Storage Manager storage agent is listening.                                                                                                      | 279  |
| <i>lanfreectcpserveraddress</i> | Specifies the TCP/IP address for the Tivoli Storage Manager storage agent.                                                                                                                         | 361  |
| <i>tcpbuffsize</i>              | Specifies the size, in kilobytes, of the Tivoli Storage Manager internal TCP/IP communication buffer.                                                                                              | 361  |
| <i>tcpnodelay</i>               | Specifies whether the server or client disables the delay of sending successive small packets on the network. This option is for all UNIX clients, except UNIX System Services client for z/OS.    | 365  |
| <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.          | 360  |
| <i>tcpcadaddress</i>            | Specifies a TCP/IP address for <i>dsmcad</i> .                                                                                                                                                     | 362  |
| <i>tcpport</i>                  | Specifies the TCP/IP port address for a Tivoli Storage Manager server.                                                                                                                             | 366  |
| <i>tcpserveraddress</i>         | Specifies the TCP/IP address for a Tivoli Storage Manager server.                                                                                                                                  | 367  |
| <i>tcpwindowsize</i>            | Specifies the size, in kilobytes, of the TCP/IP sliding window for your client node.                                                                                                               | 368  |
| <i>webports</i>                 | Enables the use of the Web client outside a firewall by specifying the TCP/IP port number used by the Client Acceptor daemon and the Web Client Agent service for communications with the Web GUI. | 392  |

## Shared Memory options

Table 49. Shared Memory communication options

| Option               | Description                                                                                                                          | Page |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>lanfreshmport</i> | Specifies the unique number that is used by the client and the storage agent to identify shared memory area used for communications. | 278  |

Table 49. Shared Memory communication options (continued)

| Option         | Description                                                                                                                   | Page |
|----------------|-------------------------------------------------------------------------------------------------------------------------------|------|
| <i>shmport</i> | Specifies the unique number that is used by the client and the server to identify shared memory area used for communications. | 278  |

---

## Server and Node options

### Authorized User

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 217 for more information.

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

**Note:** You cannot override the migration server specified in the client system options file.

Figure 3 on page 172 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

You can specify the following node options in your dsm.sys:

Table 50. 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.<br><br>Also specifies the server to contact for space management services if you have the HSM client installed and do not specify a server with the <i>migrateserver</i> option. See <i>IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide</i> , SC32-0148, for more information. | 217  |
| <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.                                                                                                                                                                                                                 | 197  |
| <i>clusternode</i>   | Specifies whether Tivoli Storage Manager participates in a High Availability Cluster Multi Processing (HACMP) environment.                                                                                                                                                                                                                                                                                                                                             | 203  |
| <i>nasnodename</i>   | Specifies the node name for the NAS file server when processing NAS file systems.                                                                                                                                                                                                                                                                                                                                                                                      | 294  |
| <i>nodename</i>      | Use the <i>nodename</i> option in your dsm.sys file to identify your workstation to the server to establish communications.                                                                                                                                                                                                                                                                                                                                            | 297  |
| <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.                                                                                                                                                                                                                                                | 342  |

Table 50. Server and Node Options (continued)

| Option                 | Description                                                                                                                                          | Page |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <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. | 384  |

## Backup and archive processing options

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

Table 51. Backup and archive processing options

| Option                   | Description                                                                                                                                                                                                                                                                                                                                                                                                        | Page |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>archmc</i>            | Use the <i>archmc</i> option with the <b>archive</b> command to specify the available management class for your policy domain to which you want to bind your archived files.                                                                                                                                                                                                                                       | 195  |
| <i>archsymbinkasfile</i> | Specifies whether you want Tivoli Storage Manager to follow a symbolic link and archive the file or directory to which it points, or archive the symbolic link only.                                                                                                                                                                                                                                               | 196  |
| <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.                                                                                                                                                             | 197  |
| <i>automount</i>         | Use this option with the <i>domain</i> option to specify all automounted file systems the Tivoli Storage Manager client tries to mount at the following points in time: <ul style="list-style-type: none"> <li>• When Tivoli Storage Manager client starts</li> <li>• When the backup is started</li> <li>• When the Tivoli Storage Manager client has reached an automounted file system during backup</li> </ul> | 199  |
| <i>changingretries</i>   | Specifies the number of retries when attempting to back up or archive a file that is in use.                                                                                                                                                                                                                                                                                                                       | 201  |
| <i>class</i>             | Specifies whether to display a list of NAS, client, or WebSphere Application Server (WAS) objects during a <b>query backup</b> , <b>query filesystem</b> , or <b>delete filesystem</b> operation.                                                                                                                                                                                                                  | 202  |
| <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.                                                                                                                                                                                                                                           | 210  |
| <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.                                                                                                                                                                                                                        | 211  |

**Note:** The *compression* option also applies to migrated files if you install the Tivoli Storage Manager HSM client on your workstation.



Table 51. Backup and archive processing options (continued)

| Option                                                                                       | Description                                                                                                                                                                                                                                                                                                                    | Page |
|----------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>deletefiles</i>                                                                           | Use the <i>deletefiles</i> option with the <b>archive</b> command to delete files from your workstation after you archive them. You can also use this option with the <b>restore image</b> command and the <i>incremental</i> option to delete files from the restored image if they were deleted after the image was created. | 218  |
| <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.                                                                                                                                                             | 219  |
| <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.                                                                                                                                                                          | 221  |
| <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.                                                                                                               | 222  |
| <i>dirsonly</i>                                                                              | Backs up, restores, archives, retrieves, or queries directories <i>only</i> .                                                                                                                                                                                                                                                  | 223  |
| <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.                                                                                                                                                                 | 225  |
| <i>domain</i>                                                                                | Specifies the file systems to include in your default client domain for an incremental backup.                                                                                                                                                                                                                                 | 227  |
| <i>domain.image</i>                                                                          | Specifies the mounted file systems and raw logical volumes that you want to include in your client domain for an image backup. This option is for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris <i>only</i> .                                                                                                      | 231  |
| <i>domain.nas</i>                                                                            | Specifies the volumes to include in your default domain for NAS image backups.                                                                                                                                                                                                                                                 | 232  |
| <i>enablearchiveretentionprotection</i>                                                      | Allows the client to connect to a data retention server.                                                                                                                                                                                                                                                                       | 235  |
| <i>enablelanfree</i>                                                                         | Specifies whether to enable an available LAN-free path to a storage area network (SAN) attached storage device.                                                                                                                                                                                                                | 237  |
| <i>exclude</i><br><i>exclude.backup</i><br><i>exclude.file</i><br><i>exclude.file.backup</i> | Use these options to exclude a file or group of files from backup services and space management services (if the HSM client is installed). The <i>exclude.backup</i> option only excludes files from normal backup, but not from HSM.                                                                                          | 246  |
| <i>exclude.archive</i>                                                                       | Excludes a file or a group of files that match the pattern from archive services <i>only</i> .                                                                                                                                                                                                                                 | 246  |
| <i>exclude.attribute.symlink</i>                                                             | Excludes a file or a group of files that are symbolic links from backup processing only.                                                                                                                                                                                                                                       | 246  |

Table 51. Backup and archive processing options (continued)

| Option                                                         | Description                                                                                                                                                                                                                                                               | Page |
|----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>exclude.compression</i>                                     | Excludes files from compression processing if you set the <b>compression</b> option to <i>yes</i> . This option applies to backups and archives.                                                                                                                          | 246  |
| <i>exclude.dir</i>                                             | Excludes a directory, its files, and all its subdirectories and their files from backup processing.                                                                                                                                                                       | 246  |
| <i>exclude.encrypt</i>                                         | Excludes specified files from encryption processing.                                                                                                                                                                                                                      | 246  |
| <i>exclude.fs</i>                                              | Excludes file spaces matching a pattern. This option is valid for all UNIX clients.                                                                                                                                                                                       | 246  |
| <i>exclude.fs.nas</i>                                          | Excludes file systems on the NAS file server from an image backup when used with the <b>backup nas</b> command. This option is for AIX and Solaris clients <i>only</i> .                                                                                                  | 246  |
| <i>exclude.image</i>                                           | Excludes mounted file systems and raw logical volumes that match the specified pattern from full image backup operations. This option is valid for AIX, HP-UX, Solaris, and all Linux clients <i>only</i> .                                                               | 246  |
| <i>encryptiontype</i>                                          | Specifies whether to use AES 128-bit or DES 56-bit data encryption during backup and archive operations. DES 56-bit encryption is the default.                                                                                                                            | 239  |
| <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.                                                                                           | 250  |
| <i>filesonly</i>                                               | Backs up, restores, retrieves, or queries files <i>only</i> .                                                                                                                                                                                                             | 253  |
| <i>groupname</i>                                               | Use this option with the <b>backup group</b> command to specify the fully qualified name of the group leader for a group.                                                                                                                                                 | 259  |
| <i>guitreeviewafterbackup</i>                                  | Specifies whether the GUI client is returned to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.                                                                                                                                  | 261  |
| <i>imagetype</i>                                               | Use the <b>imagetype</b> option with the <b>backup image</b> command or the <b>include.image</b> option to specify the type of image backup you want to perform. This option is valid for AIX, Solaris, HP-UX, Linux x86/x86_64, and Linux on POWER clients <i>only</i> . | 265  |
| <i>includexcl</i>                                              | Specifies the path and file name of an include-exclude options file.                                                                                                                                                                                                      | 268  |
| <i>include</i><br><i>include.backup</i><br><i>include.file</i> | Use these options to include files or assign management classes for backup processing.                                                                                                                                                                                    | 269  |
| <i>include.archive</i>                                         | Includes files or assigns management classes for archive processing.                                                                                                                                                                                                      | 269  |
| <i>include.attribute.symlink</i>                               | Includes a file or a group of files that are symbolic links within broad group of excluded files for backup processing <i>only</i> .                                                                                                                                      | 269  |

Table 51. Backup and archive processing options (continued)

| Option                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Page |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>include.compression</i>   | Includes files for compression processing if you set the <b>compression</b> option to <i>yes</i> . This option applies to backups and archives.                                                                                                                                                                                                                                                                                                                | 269  |
| <i>include.encrypt</i>       | Includes the specified files for encryption processing. By default, Tivoli Storage Manager does not perform encryption processing.                                                                                                                                                                                                                                                                                                                             | 269  |
| <i>include.fs</i>            | Use the <b>include.fs</b> option to control how Tivoli Storage Manager processes your file space for incremental backup.                                                                                                                                                                                                                                                                                                                                       | 269  |
| <i>include.fs.nas</i>        | Use the <b>include.fs.nas</b> option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system image backup, using the <b>toc</b> option with the <b>include.fs.nas</b> option in your <i>dsm.sys</i> file. See “Toc” on page 371 for more information. This option is valid for AIX and Solaris clients <i>only</i> . | 269  |
| <i>include.image</i>         | Specifies a file system or logical volume to be included for image backup processing. This option also provides a way to specify an explicit management class assignment for a specified file system or logical volume. The backup image command ignores all other include options. This option is valid for AIX, HP-UX, Solaris, and all Linux clients.                                                                                                       | 269  |
| <i>incrbydate</i>            | Use with the <b>incremental</b> command to request an incremental backup by date.                                                                                                                                                                                                                                                                                                                                                                              | 274  |
| <i>incremental</i>           | Use with the <b>restore image</b> command to ensure that any changes that were made to the base image are also applied to the restored image. This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris <i>only</i> .                                                                                                                                                                                                                 | 275  |
| <i>memoryefficientbackup</i> | Specifies a memory-saving backup algorithm for incremental backups when used with the <b>incremental</b> command.                                                                                                                                                                                                                                                                                                                                              | 289  |

Table 51. Backup and archive processing options (continued)

| Option                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Page |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>mode</i>                   | Use the <i>mode</i> option with these commands, as follows:<br><br><b>backup image</b><br>To specify whether to perform a selective or incremental image backup of client file systems.<br><br><b>backup nas</b><br>To specify whether to perform a full or differential image backup of NAS file systems.<br><br><b>backup was</b><br>To specify whether to perform a full or differential backup of the WebSphere Application Server (WAS) Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server to the Tivoli Storage Manager server.<br><br><b>backup group</b><br>To specify whether to perform a full or differential group backup containing a list of files from one or more file space origins. | 291  |
| <i>monitor</i>                | Specifies whether you want to monitor an image backup of file systems belonging to a Network Attached Storage (NAS) file server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 293  |
| <i>noprompt</i>               | suppresses the confirmation prompt that is presented by the <b>delete group</b> , <b>delete archive</b> , <b>expire</b> , <b>restore image</b> , and <b>set event</b> commands.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 300  |
| <i>nojournal</i>              | Use this option with the <b>incremental</b> command to specify that you want to perform the traditional full incremental backup, instead of the default journal-based backup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 299  |
| <i>optfile</i>                | Specifies the client user options file you want to use when you start a Tivoli Storage Manager session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 303  |
| <i>postsnapshotcmd</i>        | During an online image backup, this option allows you to manually bring up an application after the LVM starts a snapshot. This option is only valid if the LVM is installed and configured on your system allowing you to perform an online image backup. This option is for Linux x86/x86_64 clients <i>only</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 313  |
| <i>preservelastaccessdate</i> | Use this option during a backup or archive operation to specify whether to reset the last access date of any specified files to their original value following a backup or archive operation. By default, the Tivoli Storage Manager client <i>will not</i> reset the last access date of any backed up or archived files to their original value prior to the backup or archive operation.                                                                                                                                                                                                                                                                                                                                                                                | 316  |

Table 51. Backup and archive processing options (continued)

| Option                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Page |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>presnapshotcmd</i>     | During an online image backup operation, this option allows you to manually quiesce an application before the LVM starts a snapshot. This option is only valid if the LVM is installed and configured on your system allowing you to perform an online image backup operation. This option is for Linux x86/x86_64 clients <i>only</i> .                                                                                                                                                                                                                                                               | 321  |
| <i>removeoperandlimit</i> | Specifies that Tivoli Storage Manager removes the 20-operand limit for Unix-family platforms. If you specify the <i>removeoperandlimit</i> option with the <b>incremental</b> , <b>selective</b> , or <b>archive</b> commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits.                                                                                                                                                                                                                                                    | 324  |
| <i>skipacl</i>            | Specifies whether to skip ACL processing completely.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 348  |
| <i>skipaclupdatecheck</i> | Specifies whether to perform checksum and size comparisons before and after backup and during incremental processing.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 349  |
| <i>snapshotcachesize</i>  | Use this option to specify an appropriate snapshot size so that all old data blocks can be stored. A snapshot size of 100 percent will ensure a valid snapshot. This option is valid for Linux x86/x86_64 clients <i>only</i> .                                                                                                                                                                                                                                                                                                                                                                        | 350  |
| <i>snapshotroot</i>       | Use the <i>snapshotroot</i> option with the <b>incremental</b> , <b>selective</b> , or <b>archive</b> commands in conjunction with a third-party application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. This option is valid for all UNIX and Linux clients.                                                                                                                                                                                                         | 351  |
| <i>subdir</i>             | Specifies whether to include subdirectories of a named directory.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 356  |
| <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.                                                                                                                                                                                                                                                                                                                                                                                                                      | 358  |
| <i>toc</i>                | Use the <i>toc</i> option with the <b>backup nas</b> command or the <i>include.fs.nas</i> option to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup. If you save TOC information, you can use the <b>query toc</b> server command to determine the contents of a file system backup in conjunction with the <b>restore node</b> server command to restore individual files or directory trees. You can also use the Tivoli Storage Manager Web client to examine the entire file system tree and select files and directories to restore. | 371  |
| <i>type</i>               | Use the <i>type</i> option with the <b>query node</b> command to specify the type of node to query.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 376  |

Table 51. Backup and archive processing options (continued)

| Option                   | Description                                                                                                                                                                                                                                                                                                                                                                                           | Page |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>v2archive</i>         | Use the <i>v2archive</i> option with the <b>archive</b> 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.                                                                                                                                                                                     | 378  |
| <i>virtualfsname</i>     | Use this option with the <b>backup group</b> command to specify the name of the container for the group on which you want to perform the operation.                                                                                                                                                                                                                                                   | 381  |
| <i>virtualmountpoint</i> | Defines a virtual mount point for a file system if you want to consider files for backup that begin with a specific directory within that file system.                                                                                                                                                                                                                                                | 382  |
| <i>wasnode</i>           | If WAS security is enabled, use the <i>wasnode</i> option with the <b>set waspassword</b> command to specify the node name when setting the user name and password for each installation of WAS on your machine.                                                                                                                                                                                      | 389  |
| <i>wasexphome</i>        | To back up the WebSphere Application Server-Express, use the <i>wasexphome</i> option to specify the fully qualified installation path of the WebSphere Application Server-Express.                                                                                                                                                                                                                   | 386  |
| <i>washome</i>           | Use the <i>washome</i> option in your client user options file (dsm.opt) to specify an override base install path for the Application Server                                                                                                                                                                                                                                                          | 387  |
| <i>wasndhome</i>         | Use the <i>wasndhome</i> option in your client user options file (dsm.opt) to specify an override base install path for the Network Deployment Manager.                                                                                                                                                                                                                                               | 388  |
| <i>wastype</i>           | Use the <i>wastype</i> option with the <b>backup was</b> command to back up the WAS Network Deployment Manager or the Application Server. If WAS security is enabled, use the <i>wastype</i> option with the <b>set waspassword</b> command to specify the WAS Network Deployment Manager or Application Server when setting the user name and password for each installation of WAS on your machine. | 390  |
| <i>wasuser</i>           | If WAS security is enabled, use the <i>wasuser</i> option with the <b>set waspassword</b> command to set the user name for each installation of WAS on your machine.                                                                                                                                                                                                                                  | 391  |

## Restore and retrieve processing options

The following options relate to restore and retrieve processing.

Table 52. Restore and retrieve processing options

| Option          | Description                                                                                | Page |
|-----------------|--------------------------------------------------------------------------------------------|------|
| <i>dirsonly</i> | Qualifies the operation (backup, archive, restore, retrieve) to process directories alone. | 223  |
| <i>filelist</i> | Specifies a file containing a list of files to be processed by the specified command.      | 250  |

Table 52. Restore and retrieve processing options (continued)

| Option                        | Description                                                                                                                                                                                                                                                                                                                                                                                                   | Page |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>filesonly</i>              | Qualifies the operation (backup, archive, restore, retrieve) to process files alone.                                                                                                                                                                                                                                                                                                                          | 253  |
| <i>followsymbolic</i>         | Specifies whether you want to restore files to symbolic links or use a symbolic link as a virtual mount point.                                                                                                                                                                                                                                                                                                | 254  |
| <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.                                                                                                                                                                                                               | 255  |
| <i>fromnode</i>               | Permits one node to perform commands for another node. A user on another node must use the <b>set access</b> command to permit you to query, restore, or retrieve files or images for the other node.                                                                                                                                                                                                         | 256  |
| <i>fromowner</i>              | Displays file spaces for an alternate owner. Also specifies an alternate owner from which to restore or retrieve files.                                                                                                                                                                                                                                                                                       | 257  |
| <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.                                                                                                                                                                                                               | 258  |
| <i>guitreeviewafterbackup</i> | Specifies whether the GUI client is returned to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.                                                                                                                                                                                                                                                                      | 261  |
| <i>ifnewer</i>                | Replaces an existing file with the latest backup version only if the backup version is newer than the existing file.                                                                                                                                                                                                                                                                                          | 263  |
| <i>imagetofile</i>            | Use the <i>imagetofile</i> option with the <b>restore image</b> command to specify that you want to restore the source image to a file. You might need to restore the image to a file in the event of bad sectors present on the target volume, or if you want to do some manipulations with the image data. This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris <i>only</i> . | 264  |
| <i>inactive</i>               | Displays a list of active and inactive files when used with the <i>pick</i> option.                                                                                                                                                                                                                                                                                                                           | 267  |
| <i>latest</i>                 | Restores the most recent backup version of a file whether it is active or inactive.                                                                                                                                                                                                                                                                                                                           | 281  |
| <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.                                                                                                                                                                                                                                                            | 282  |
| <i>location</i>               | Specifies where Tivoli Storage Manager searches for the backup set during a query or restore operation.                                                                                                                                                                                                                                                                                                       | 283  |
| <i>makesparsefile</i>         | Use the <i>makesparsefile</i> option with the <b>restore</b> or <b>retrieve</b> commands to specify how sparse files are recreated.                                                                                                                                                                                                                                                                           | 285  |



Table 52. Restore and retrieve processing options (continued)

| Option              | Description                                                                                                                                                                                                                                   | Page |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>monitor</i>      | Specifies whether you want to monitor an image restore of one or more file systems belonging to a Network Attached Storage (NAS) file server.                                                                                                 | 293  |
| <i>noprompt</i>     | suppresses the confirmation prompt that is presented by the <b>delete group</b> , <b>delete archive</b> , <b>expire</b> , <b>restore image</b> , and <b>set event</b> commands.                                                               | 300  |
| <i>optfile</i>      | Specifies the client user options file you want to use when you start a Tivoli Storage Manager session.                                                                                                                                       | 303  |
| <i>pick</i>         | Creates a list of backup versions, images, or archive copies that match the file specification you enter. From the list, you can select the versions to process. Include the <i>inactive</i> option to view both active and inactive objects. | 308  |
| <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.                                                                        | 309  |
| <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.                                                                        | 310  |
| <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.                                                                                                 | 318  |
| <i>replace</i>      | Specifies whether to overwrite an existing file, or to prompt you for your selection when you restore or retrieve files.                                                                                                                      | 325  |
| <i>showmembers</i>  | Displays all members of a group.                                                                                                                                                                                                              | 347  |
| <i>subdir</i>       | Specifies whether you want to include subdirectories of a named directory.                                                                                                                                                                    | 356  |
| <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.                                                                                      | 358  |
| <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.                                             | 373  |
| <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.                                             | 373  |
| <i>type</i>         | Use the <i>type</i> option with the <b>query node</b> command to specify the type of node to query.                                                                                                                                           | 376  |

Table 52. Restore and retrieve processing options (continued)

| Option             | Description                                                                                                                                                                                                                                                                                                                                                                                                   | Page |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>verifyimage</i> | Use the <i>verifyimage</i> option with the <b>restore image</b> command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log. This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris <i>only</i> . | 380  |
| <i>washome</i>     | Use the <i>washome</i> option in your client user options file (dsm.opt) to specify an override base install path for the Application Server                                                                                                                                                                                                                                                                  | 387  |
| <i>wasndhome</i>   | Use the <i>wasndhome</i> option in your client user options file (dsm.opt) to specify an override base install path for the Network Deployment Manager.                                                                                                                                                                                                                                                       | 388  |
| <i>wastype</i>     | Use the <i>wastype</i> option with the <b>restore was</b> command to restore the WAS Network Deployment Manager or Application Server.                                                                                                                                                                                                                                                                        | 390  |

## 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 53. Scheduling options

| Option                                             | Description                                                                                                                                                                             | Page |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>managedservices</i>                             | Specifies whether the Storage Manager Client Acceptor daemon manages the Web client, the scheduler, or both.                                                                            | 286  |
| <i>maxcmdretries</i>                               | Specifies the maximum number of times the client scheduler attempts to process a scheduled command that fails.                                                                          | 288  |
| <i>postschedulecmd,</i><br><i>postnschedulecmd</i> | Specifies a command to process after running a schedule.                                                                                                                                | 311  |
| <i>preschedulecmd,</i><br><i>prenschedulecmd</i>   | Specifies a command to process before running a schedule.                                                                                                                               | 314  |
| <i>queryschedperiod</i>                            | Specifies the number of hours the client scheduler waits between attempts to contact the server for scheduled work.                                                                     | 322  |
| <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. | 329  |
| <i>schedcmddisabled</i>                            | Specifies whether to disable the scheduling of generic commands specified by your Tivoli Storage Manager administrator.                                                                 | 331  |

Table 53. Scheduling options (continued)

| Option                                    | Description                                                                                                                                                                                                                                                                                                                                               | Page |
|-------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>schedcmduser</i> (server defined only) | The scheduler executes commands under a uid of 0, however, there might be some users who have a different user ID. In this case, your Tivoli Storage Manager administrator can define schedules and allow these schedules to be executed under a uid other than 0, using this option. The Tivoli Storage Manager Client API does not support this option. | 154  |
| <i>schedlogmax</i>                        | Specifies the maximum size of the scheduler log, in megabytes.                                                                                                                                                                                                                                                                                            | 332  |
| <i>schedlogname</i>                       | Specifies the path and file name where you want to store schedule log information.                                                                                                                                                                                                                                                                        | 333  |
| <i>schedlogretention</i>                  | Specifies the number of days to keep log file entries in the schedule log, and whether to save pruned entries.                                                                                                                                                                                                                                            | 334  |
| <i>schedmode</i>                          | Specifies which schedule mode to use, <i>polling</i> or <i>prompted</i> .                                                                                                                                                                                                                                                                                 | 336  |
| <i>schedrestretrdisabled</i>              | Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing restore or retrieve schedule operations.                                                                                                                                                                                                                      | 338  |
| <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.                                                                                                                                                                              | 344  |
| <i>srvprepostscheddisabled</i>            | Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing pre-schedule and post-schedule commands when performing scheduled operations.                                                                                                                                                                                 | 354  |
| <i>srvprepostsnapdisabled</i>             | Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing pre-snapshot and post-snapshot commands when performing scheduled image snapshot backup operations.                                                                                                                                                           | 355  |
| <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 336) for details.  | 363  |
| <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 336) for details.                                                                                                                                                  | 364  |

---

## Format options

You can use the following options to select different formats for date, time, and numbers.

Table 54. Format options

| Option              | Description                                  | Page |
|---------------------|----------------------------------------------|------|
| <i>dateformat</i>   | Specifies the format for displaying dates.   | 214  |
| <i>numberformat</i> | Specifies the format for displaying numbers. | 301  |
| <i>timeformat</i>   | Specifies the format for displaying time.    | 369  |

---

## Command processing options

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

Table 55. 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.                                                                                                                                     | 234  |
| <i>quiet</i>        | Limits the number of messages that are displayed on your screen during processing. This option can be overridden by the server.                                                                                                         | 323  |
| <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> .                                                                        | 339  |
| <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. | 340  |
| <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.                                                                                | 379  |

---

## Authorization options

These options control access to a Tivoli Storage Manager server.

Table 56. Authorization options

| Option                | Description                                                                                                                                                    | Page |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <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.                                | 239  |
| <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. | 240  |
| <i>groups</i>         | Specifies the groups on your workstation that you want to authorize to request Tivoli Storage Manager services from the server.                                | 260  |

Table 56. Authorization options (continued)

| Option                    | Description                                                                                                                                                                                 | Page |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>mailprog</i>           | Specifies the program and user ID where you want to send a newly-generated password when the old one expires.                                                                               | 284  |
| <i>password</i>           | Specifies a Tivoli Storage Manager password.                                                                                                                                                | 304  |
| <i>passwordaccess</i>     | Specifies whether you want to use a generated password or be prompted for a password each time you start the client.                                                                        | 305  |
| <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. | 307  |
| <i>revokeremoteaccess</i> | Restricts an administrator with client access privileges from accessing your workstation through the Web client.                                                                            | 330  |
| <i>users</i>              | Authorizes specific users on your workstation to request services from a server.                                                                                                            | 377  |

## 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 57. Error processing options

| Option                   | Description                                                                                                                                 | Page |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>errorlogmax</i>       | Specifies the maximum size of the error log, in megabytes.                                                                                  | 242  |
| <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. | 243  |
| <i>errorlogretention</i> | Specifies how many days to maintain error log entries before pruning, and whether to save the pruned entries.                               | 244  |

## Transaction processing options

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

Table 58. 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). | 204  |
| <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.                                                                                                                                                                                                                                                                                                                                                                  | 208  |
| <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.                                                                                                                                                                                                                                                                                                                                                        | 209  |
| <i>diskbuffsize</i>        | Specifies the maximum disk I/O buffer size (in kilobytes) that the client can use when reading files.                                                                                                                                                                                                                                                                                                                                                                                                                   | 224  |
| <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.<br><br><b>Recommendation:</b> Discontinue the use of <i>largecommbuffers</i> because future releases of Tivoli Storage Manager might not accept this option.           |      |
| <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.                                                                                                                                                                                                                                                                                                                                                                                                    | 295  |
| <i>resourceutilization</i> | Use the <i>resourceutilization</i> option in your <i>dsm.sys</i> file to regulate the level of resources the Tivoli Storage Manager server and client can use during processing.                                                                                                                                                                                                                                                                                                                                        | 327  |
| <i>txnbytelimit</i>        | Specifies the number of kilobytes the client program buffers before it sends a transaction to the server.                                                                                                                                                                                                                                                                                                                                                                                                               | 375  |

---

## Web client options

The following are options for the Tivoli Storage Manager Web Client.

Table 59. Web client options

| Option                    | Description                                                                                                                                                                                           | Page |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>httpport</i>           | Specifies a TCP/IP port address for the Web client.                                                                                                                                                   | 262  |
| <i>managedservices</i>    | Specifies whether the Storage Manager Client Acceptor daemon manages the Web client, the scheduler, or both.                                                                                          | 286  |
| <i>revokeremoteaccess</i> | Restricts administrator access on a client workstation through the Web client.                                                                                                                        | 330  |
| <i>webports</i>           | Enables the use of the Web client outside a firewall by specifying the TCP/IP port number used by the Client Acceptor daemon and the Web Client Agent service for communications with the Web client. | 392  |

---

## 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 462 for more information.

Table 60. Diagnostics options

| Option          | Description                                                                                                                              | Page |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------|------|
| <i>console</i>  | Use the <i>console</i> option with the <b>query systeminfo</b> command to output system information to the console.                      | 213  |
| <i>filename</i> | Use the <i>filename</i> option with the <b>query systeminfo</b> command to specify a file name in which to store the system information. | 252  |

---

## 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 61.

## 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 1999" /home/
```

- For options that do not include parameters, enter a command, a dash (-) and the option name. For example,

```
dsmc incremental -quiet
```

**Note:** Use a leading dash (-) to indicate that the following text is the name of an option. If an object name begins with a dash, you must surround it in either single quotes (') or double quotes ("). 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 double quotes (").

- 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 xiii.
- Enter options before or after command parameters. For example, you can enter the option before or after a file specification:

```
dsmc selective -subdir=yes "/home/devel/proj1/*"
dsmc selective "/home/devel/proj1/*" -subdir=yes
```

- When entering several options on a command, separate them with a blank space.
  - Enclose the value in quotes (" ") if the option value that you enter contains a blank space. For example:
- ```
dsmc archive -description="Project A" "/home/devel/proj1/*"
```
- Most options that you enter on the command line override the value set in the preferences file. However, when you use the *domain* option with the **incremental** command, it adds to the domain specified in your client options file rather than overriding the current value.
 - The maximum number of characters for a file name is 256. The maximum combined length of the file name and path name is 1024 characters.

Table 61 lists client command options that you can enter only on the command line with specific commands.

Table 61. 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	195

Table 61. Client command options (continued)

Command option	Description	Commands	Page
<i>class</i>	Specifies whether to display a list of NAS objects or client objects when using the following commands:	query backup delete filespace query filespace	202
<i>console</i>	Use the <i>console</i> option with the query systeminfo command to output system information to the console.	query systeminfo	213
<i>deletefiles</i>	Deletes the local copy of files from your workstation after they are archived on the server. Can also be used with the restore image command and the <i>incremental</i> option to delete files from the restored image that are deleted from the file space after the image is created.	archive restore image	218
<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	219
<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	221
<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	223
<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 backup group delete archive delete backup expire incremental query archive query backup restore retrieve selective	250
<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	252
<i>filesonly</i>	Backs up, restores, retrieves, or queries files <i>only</i> .	archive incremental query archive query backup restore restore backupset retrieve selective	253

Table 61. Client command options (continued)

Command option	Description	Commands	Page
<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 restore group restore was	255
<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 or images for the other node.	query archive query backup query filespace query group query image query mgmtclass query was restore restore group restore image restore was retrieve	256
<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 query group query image query was restore restore image restore group restore was retrieve	257
<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 restore group retrieve restore was	258
<i>groupname</i>	Specifies the fully qualified name for a group.	backup group	259
<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 restore group retrieve restore was	263
<i>imagetofile</i>	Use the <i>imagetofile</i> option with the restore image command to specify that you want to restore the source image to a file. You might need to restore the image to a file in the event of bad sectors present on the target volume, or if you want to do some manipulations with the image data. This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris <i>only</i> .	restore image	264

Table 61. Client command options (continued)

Command option	Description	Commands	Page
<i>inactive</i>	Displays a list of active and inactive files when used with the <i>pick</i> option.	delete group query backup query group query image query nas query was restore restore group restore image restore nas restore was	267
<i>incrbydate</i>	Requests an incremental backup by date.	incremental	274
<i>incremental</i>	Applies changes to the base image using information from incremental backups made after the original image backup. This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris <i>only</i> .	restore image	275
<i>latest</i>	Restores the most recent backup version of a file whether it is active or inactive.	restore restore group restore was	281
<i>location</i>	Specifies whether Tivoli Storage Manager searches for a backup set on the server, in local files, or on a tape device during a query or restore operation.	query backupset restore backupset	283
<i>mode</i>	Use the <i>mode</i> option with these commands, as follows: backup image To specify whether to perform a selective or incremental image backup of client file systems. backup nas To specify whether to perform a full or differential image backup of NAS file systems. backup was To specify whether to perform a full or differential backup of the WebSphere Application Server (WAS) Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server to the Tivoli Storage Manager server. backup group To specify whether to perform a full or differential group backup containing a list of files from one or more file space origins.	backup group backup nas backup image backup was restore nas	291
<i>monitor</i>	Specifies whether you want to monitor an image backup or restore of one or more file systems belonging to a Network Attached Storage (NAS) file server.	backup nas restore nas	293

Table 61. Client command options (continued)

Command option	Description	Commands	Page
<i>nojournal</i>	Use this option with the with the incremental command to specify that you want to perform the traditional full incremental backup, instead of the default journal-based backup.	incremental	299
<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 delete group expire restore image	300
<i>optfile</i>	Specifies the client user options file you want to use when you start a Tivoli Storage Manager session.	dsmc	303
<i>pick</i>	Creates a list of backup versions, images, or archive copies that match the file specification you enter. From the list, you can select the versions to process. Include the <i>inactive</i> option to view both active and inactive objects.	delete archive delete group expire query nas restore restore group restore image restore nas restore was retrieve	308
<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 query group query image query nas query was restore restore group restore image restore nas restore was	309
<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 query image query nas restore restore image restore nas	310
<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 restore group retrieve restore was	318
<i>removeoperandlimit</i>	Specifies that Tivoli Storage Manager removes the 20-operand limit for Unix-family platforms. 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	324
<i>showmembers</i>	Displays all members of a group.	query group query was restore group restore was	347

Table 61. Client command options (continued)

Command option	Description	Commands	Page
<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 restore group restore was	373
<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 restore group restore was	374
<i>type</i>	Use the <i>type</i> option with the query node command to specify the type of node to query.	query node	376
<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	378
<i>verifyimage</i>	Use the <i>verifyimage</i> option with the restore image command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log. This option is valid for Windows 32-bit platforms AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris <i>only</i> .	restore image	380
<i>virtualfsname</i>	Specifies the name of the virtual file space for the group on which you want to perform the operation.	backup group	381
<i>wasnode</i>	Use the <i>wasnode</i> option with the set waspassword commands to specify the WAS node name when performing the operation on the WAS Network Deployment Manager or Application Server.	set waspassword	389
<i>wastype</i>	Use the <i>wastype</i> option with the backup was , query was , restore was , or set waspassword commands to perform the operation on the WAS Network Deployment Manager or Application Server.	backup was query was restore was set waspassword	390
<i>wasuser</i>	If WAS security is enabled, use the <i>wasuser</i> option with the set waspassword command to set the user name for each installation of WAS on your machine.	set waspassword	391

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

Options valid on the initial command line

<i>clusternode</i>	<i>preschedulecmd</i> , <i>prenschedulecmd</i> (can be included in the schedule definition)
<i>commmethod</i>	
<i>diskbuffsize</i>	<i>queryschedperiod</i>
<i>editor</i>	<i>resourceutilization</i>
<i>enablelanfree</i>	<i>retryperiod</i>
<i>errorlogmax</i>	<i>schedlogmax</i>
<i>errorlogname</i>	<i>schedlogname</i>
<i>errorlogretention</i>	<i>schedlogretention</i>
<i>lanfreecommmethod</i>	<i>schedmode</i>
<i>lanfreeshimport</i>	<i>servername</i>
<i>lanfreetcpport</i>	<i>sessioninitiation</i>
<i>language</i>	<i>tcpbuffsize</i>
<i>maxcmdretries</i>	<i>tcpadaddress</i>
<i>nfstimeout</i>	<i>tcpclientaddress</i>
<i>nodename</i>	<i>tcpclientport</i>
<i>optfile</i>	<i>tcpwindowsize</i>
<i>password</i>	<i>txnbytelimit</i>
<i>postschedulecmd</i> , <i>postnschedulecmd</i> (can be included in the schedule definition)	<i>virtualnodename</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 xiii 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 client options file (if applicable).
- Examples of using the option on the command line (if applicable). Options with a command line example of **Does not apply** cannot be used with command line or scheduled commands.

Note: 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**.

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.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—ARCHMc *--managementclass*—◄◄

Parameters

managementclass

Specifies an available management class in the active policy set of your policy domain. This management class overrides the default management class and any *include* statements for the files and directories you are archiving.

Examples

Command line:

```
dsmc archive -archmc=ret2yrs /home/plan/proj1/budget.jan
```

Archsymlinkasfile

The *archsymlinkasfile* option specifies whether Tivoli Storage Manager follows a symbolic link and archives the file or directory to which it points, or archives the symbolic link only. Use this option with the **archive** command.


Supported Clients

This option is valid for all UNIX clients. The server can also define this option.

Options File

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

Syntax



Parameters

- Yes* Specifies that Tivoli Storage Manager follows a symbolic link and archives the associated file or directory. This is the default.
- No* Specifies that Tivoli Storage Manager archives the symbolic link and not the associated file or directory.

Examples

Options file:
archsymlinkasfile no

Command line:
-archsymb=no

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 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.

For example, you can use the following command to back up shared data for file space /cluster1/mydata stored under the node name MyCluster:

```
dsmc incremental c:/cluster1/mydata -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* and you cannot perform NAS backup using *asnodename*.. Also, *asnodename* can be used for clustered systems, although no specific cluster software is supported.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

Place this option in the *dsm.sys* file *within* a server stanza. You can set this option on the **General** tab of the Preferences editor.

Syntax

▶▶—ASNODENAME— *targetnode* —————▶▶

Parameters

targetnode

Specifies the node name on the Tivoli Storage Manager server under which you want to back up or restore data.

Examples

Options file:

```
asnodename mycluster
```

Command line:

```
-asnodename=mycluster
```

This option is not valid in interactive mode, but it can be defined in the options portion of a schedule definition.

Automount

Authorized User

The *automount* option adds an automounted file system into the domain by mounting it. Use this option with the *domain* option.

Use this option to specify all automounted file systems the Tivoli Storage Manager client tries to mount at the following points in time:

- When Tivoli Storage Manager client starts
- When the backup is started
- When the Tivoli Storage Manager client has reached an automounted file system during backup

It is unnecessary to explicitly specify an automounted file system in the *automount* statement if you use the keywords *all-auto-nfs* or *all-auto-lofs* in the domain statement and the file system is already mounted. However, you should add this file system in the *automount* statement to ensure the file system has been mounted at all the points in time mentioned above. The automounted file systems are remounted if they have gone offline in the meantime during a backup.

See “Domain” on page 227 for more information about working with automounted file systems and the *domain* option.

Supported Clients

This option is valid for all UNIX platforms. The Tivoli Storage Manager client API does not support this option.

Options File

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

Syntax



Parameters

file spacename

Specifies one or more automounted file systems that are mounted and added into the domain.

Examples

Options file:

```
automount fs1 fs2
```

Command line:

Does not apply.

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**
- **query image**
- **query systemobject**
- **query systemservices**
- **query systemstate**
- **restore**
- **restore backupset**
- **restore image**
- **restore systemobject**
- **restore systemservices**
- **restore systemstate**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

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

Syntax



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*.

tape device

Specifies the name of the tape device where the tape containing the start of the backup set is located. The *location* option should be set to *tape*.

Examples

Command line:

```
dsmc restore /bkSets/file.1 -backupsetname=YEAR_END_ACCOUNTING.1234  
dsmc query image -backupsetname=WEEKLY_BSET.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.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

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

Class

The *class* option specifies whether to display a list of NAS or client objects when using the following commands:

- **query backup**
- **delete filesystem**
- **query filesystem**

For example, to display a list of the file spaces belonging to a NAS node, enter the following command:

```
query filesystem -class=nas
```

Supported Clients

This option is valid for AIX and Solaris clients *only*. The Tivoli Storage Manager client API does not support this option.

Syntax



Parameters

client

Specifies that you want to display a list of file spaces for a client node. This is the default.

nas

Specifies that you want to display a list of file spaces for a NAS node.

Examples

Command line:

```
q backup -nasnodename=nodename -class=nas
```

Clusternode

The *clusternode* option specifies whether Tivoli Storage Manager participates in a High Availability Cluster Multi Processing (HACMP) environment. For information on how to configure a Tivoli Storage Manager server to manage a cluster configured client, see Appendix B, “Configuring the backup-archive client in an HACMP takeover environment,” on page 511.

Note: You must set the *clusternode* option to *yes* for all Tivoli Storage Manager-managed cluster operations. Inconsistent use of the *clusternode* option for a given Tivoli Storage Manager cluster node name can cause Tivoli Storage Manager to invalidate the cluster node name encrypted password, and prompt the user to reenter the password during the next Tivoli Storage Manager program invocation.

Use the *optfile* option to properly call the correct (cluster) dsm.opt for all Tivoli Storage Manager programs to ensure proper Tivoli Storage Manager functionality for cluster related operations. See “Optfile” on page 303 for more information.

Supported Clients

This option is valid for the AIX clients.

Options File

Place this option in the client system options file (dsm.sys).

Syntax



Parameters

Yes Specifies that you want Tivoli Storage Manager to back up cluster resources and participate in cluster failover for high availability. The default node name is the HACMP-defined cluster name. .

No Specifies that you do not want the Tivoli Storage Manager client to participate in cluster failover. This is the default.

Examples

Options file:
cluster no

Command line:
-cluster=yes

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

Collocatebyfilespec

Use the *collocatebyfilespec* option to specify whether the Tivoli Storage Manager client uses only one server session to send objects generated from one file specification.

Setting the *collocatebyfilespec* option to *yes* attempts to eliminate interspersing of files from different file specifications, by limiting the client to one server session per file specification. Therefore, if you store the data to tape, files for each file specification are stored together on one tape (unless another tape is required for more capacity).

Considerations:

- Use the *collocatebyfilespec* option only if the storage pool is going directly to tape. If you use this option going to a disk storage pool, you could affect some load balancing, and therefore, performance.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option.

Options File

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

Syntax



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.


Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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.

Syntax



Parameters

TCPip

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

SHAREdmem

Use the Shared Memory communication method when the client and server are running on the same system. This provides better performance than the TCP/IP protocol. This communication method is valid for AIX, HP-UX, Linux, and Solaris clients *only*.

When specifying this communication method on AIX, the client can be logged in as root or non-root, as long as the server is running as root. If the server is not running as root, the user ID running the client must match the user ID running the server.

Attention: When using **COMMMethod** *SHAREdmem* on Linux, you might receive error message: ANR8294W Shared Memory Session unable to initialize on the server or storage agent console. By default, Linux is not set up with sufficient system resources to create the message queues. You must increase the kernel parameter, *MSGMNI*, to 128 (the default is 16). You can modify this parameter by performing the following command:

```
echo 128 > /proc/sys/kernel/msgmni
```

To enable this parameter to remain persistent after rebooting the system, you can instead add the following line to the file */etc/sysctl.conf*, then reboot the machine:

```
kernel.msgmni=128
```

To view the current ipc settings, run this command:

```
ipcs -l
```

Now look at the max queues system wide value. The default is 16.

Examples

Options file:

```
comm tcpip
```

Command line:

```
-comm=tcpip
```

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.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option.

Options File

Place this option in the `dsm.sys` file. 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. API behavior depends on the application. Application backups might fail.

Examples

Options file:
`compressalways yes`

Command line:
`-compressa=no`

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

Compression

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. The **backup image** command uses the *compression* option value specified in the *dsm.sys*. This option is valid on the initial command line and in interactive mode. The server can also define this option which overrides the client value.

Tivoli Storage Manager backs up a sparse file as a regular file if client compression is off. Set the *compression* option to *yes* to enable file compression when backing up sparse files to minimize network transaction time and maximize server storage space.

If you set the *compressalways* option to *yes*, compression continues even if the file size increases. To stop compression if the file size grows, and resend the file uncompressed, set the *compressalways* option to *no*.

If you set the *compression* option to *yes*, you can control compression processing in the following ways:

- Use the *exclude.compression* option in your client system options file (*dsm.sys*) to exclude specific files or groups of files from compression processing. See “Exclude options” on page 246 for more information.
- Use the *include.compression* option in your client system options file (*dsm.sys*) to include files within a broad group of excluded files for compression processing. See “Include options” on page 269 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.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option.

Options File

Place this option in the *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 (includes ULIMIT information for UNIX and Linux).
- POLICY - Policy set dump.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually dsmsched.log).
- CLUSTER - AIX cluster information.

Note: The **query systeminfo** command is intended primarily as an aid for IBM support to assist in diagnosing problems, although users who are familiar with the concepts addressed by this information might also find it useful. If you use the *console* option, no special formatting of the output is performed to accommodate screen height or width. Therefore, the console output might be difficult to read due to length and line-wrapping. In this case, 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 252 for more information.

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

►►—CONSOLE—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
query systeminfo dsmoptfile errorlog -console
```

Dateformat

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

The AIX, Solaris, and HP-UX clients support locales other than English that describe every user interface that varies with location or language. See Table 18 on page 47 for supported locales.

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 asr**
- **query backup**
- **query filespace**
- **query image**
- **restore**
- **restore image**
- **restore nas**
- **retrieve**
- **set event**

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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:

0 Use the locale-specified date format. For AIX, HP-UX, and Solaris: This is the default if the locale-specified date format consists of digits and separator characters.

1 MM/DD/YYYY

For AIX, HP-UX, and Solaris: This is the default if the locale-specified date format consists of anything but digits and separator characters.

This is the default for the following available translations:

- US English
- Chinese (Traditional)
- Korean

2 DD-MM-YYYY

This is the default for the following available translations:

- Brazilian Portuguese
- Italian

3 YYYY-MM-DD

This is the default for the following available translations:

- Japanese
- Chinese (Simplified)
- Polish

4 DD.MM.YYYY

This is the default for the following available translations:

- German
- French
- Spanish
- Czech
- Russian

5 YYYY.MM.DD

This is the default for the following available translations:

- Hungarian

For AIX, HP-UX, and Solaris: To set a particular date format, edit the source file for your locale and modify the **d_fmt** line to support your needs. Whatever date format you select applies both to output and to input; however, the input year can be either 2 or 4 digits.

"%m/%d/%y"

Displays the date in the form MM/DD/YY

"%d.%m.%Y"

Displays the date in the form DD.MM.YYYY

When you include the *dateformat* option with a command, it must precede the *fromdate*, *pitdate*, and *todate* options.

Examples

Options file:

```
dateformat 3
```

Command line:

```
-date=3
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This will be the value from the `dsm.opt` file unless overridden by the initial command line or by an option forced by the server.

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 client user options file (`dsm.opt`).

If you have the HSM client installed on your workstation, and you do not specify a migration server with the *migrateserver* option, use this option to specify the server to which you want to migrate files. See *IBM Tivoli Storage Manager for Space Management for UNIX and Linux User's Guide* for more information.

Supported Clients

This option is valid for all UNIX clients.

Options File

Place this option *at the beginning* of the `dsm.sys` file *before* any server stanzas.

Syntax

▶▶—DEFAULTServer— *servername*—————▶▶

Parameters

servername

Specifies the name of the default server to which you back up or archive files. The server to which files are migrated from your local file systems can also be specified with this option.

Examples

Options file:

```
defaults server_a
```

Command line:

Does not apply.

Deletefiles

Use the *deletefiles* option with the **archive** command to delete files from your workstation after you archive them.

You can also use this option with the **restore image** command and the *incremental* option to delete files from the restored image if they were deleted after the image was created. Deletion of files will be performed correctly if the Tivoli Storage Manager server's backup copy group has enough versions for existing and deleted files.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—DELetefiles—◀◀

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc archive "/home/foo/*.c" -deletefiles
dsmc restore image /local/data -incremental -deletefiles
```

Description

The *description* option assigns or specifies a description for files when performing archive, delete archive, retrieve, query archive, or query backupset.

For example, if you want to archive a file named budget.jan and assign to it the description **2002 Budget for Proj 1**, you would enter:

```
dsmc archive -des="2003 Budget for Proj 1" /home/plan/  
proj1/budget.jan
```

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**

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—DEscription =- *description*—►►

Parameters

description

Assigns a description to the file you are archiving. If you do not specify a description with the **archive** command, the default is Archive Date:*x*, where *x* is the current system date. Note that the date is always 10 characters long. If your date format uses a two digit year, there 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" "/home/  
allproj/*"
```

You can then use the description to retrieve all of the files.

Examples

Command line:

```
dsmc archive "/home/foo/*.prj" -des="2003 Budget for Proj 1"  
dsmc query backupset -loc=server -descr="My Laptop"
```

Detail

Use the *detail* option to display management class, file space, backup, and archive information depending on the command with which it is used.

Use the *detail* option with the **query mgmtclass** command to display detailed information about each management class in your active policy set. If you do not use the *detail* option, only the management class name and a brief description are displayed on the screen. If you specify the *detail* option, information about attributes in each copy group contained in each management class is displayed on the screen. A management class can contain a backup copy group, an archive copy group, both, or neither.

Use the *detail* option with the **delete filespace** and **query filespace** commands to determine the fsID of a file space. The fsID also appears in the file information dialog in the backup-archive client and Web client GUIs.

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

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

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—DETail—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

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

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 163 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.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option.

Options File

Place this option in the *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**

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶—Dirsonly—▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

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

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, or HSM migration client performance can usually be achieved if the value for this option is equal to or smaller than the amount of file read ahead provided by the client file system. A larger buffer will require more memory and might not improve performance.

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

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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. For AIX: If the *enablelanfree* option is set to *no*, the default setting for *diskbuffsize* is 256.

Examples

Options file:
diskbuffsize 64

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 269 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 2 gigabytes of disk space for each million files or directories being backed up.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

Place this option in the *dsm.sys* file within a server stanza.

Syntax

►►—DISKCACHELocation— *path* —————►►

Parameters

path Specifies the location where the disk cache database 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 file systems 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 file systems to process during an incremental backup:

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

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

Note: You can include a virtual mount point in your client domain. For information about defining a virtual mount point, see “Virtualmountpoint” on page 382.

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

```
domain /home /usr /datasave
```

and the following on the command line:

```
dsmc incremental -domain="/fs1 /fs2"
```

Tivoli Storage Manager performs an incremental backup of the `/home`, `/usr`, `/datasave`, `/fs1`, and `/fs2` file systems.

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 file systems that you specify in the file specification. For example, if you enter:

```
dsmc incremental /fs1 /fs2 -domain="/fs3 /fs4"
```

Tivoli Storage Manager performs an incremental backup for the `/fs1` and `/fs2` file systems only.

You can also exclude file systems by specifying the dash (-) operator before the file systems. For example, in the following option Tivoli Storage Manager will process all local file systems except for the `/home` file system:

```
domain ALL-LOCAL -/home
```

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

Attention: If you are running GPFS for AIX or GPFS for Linux x86/x86_64 in a multi-node cluster, and all nodes share a mounted GPFS file system, Tivoli Storage Manager processes this file system as a local file system. Tivoli Storage Manager backs up the file system on each node during an incremental backup. To avoid this, you can do *one* of the following:

- Explicitly configure the *domain* statement in the client user options file (dsm.opt) to list the file systems you want that node to back up.
- Set the *exclude.fs* option in the client system options file to exclude the GPFS file system from backup services.

Automounted file systems

When performing a backup with the *domain* option set to *all-local*, files handled by automounter and loopback file systems are not backed up.

If you back up a file system with the *domain* option set to *all-local*, any subdirectories that are mount points for an automounted file system (autofs) are excluded from backup. Any files that exist on the server for the automounted subdirectory are expired.

When performing a backup with the *domain* option set to *all-lofs*, all explicit loopback file systems (LOFS) are backed up and all automounted file systems are excluded. For loop devices and local file systems handled by automounter, set the *domain* option to *all-auto-lofs*.

You should use the *automount* option with the *domain* option to specify one or more automounted file systems to be mounted and added into the domain. If you specify the *automount* option, automounted file systems are remounted if they have gone offline during the execution of the **incremental** command. See “Automount” on page 199 for more information.

Virtual mount points cannot be used with automounted file systems.

For HP-UX: The *domain* option is enhanced with the new keywords *all-auto-lofs* and *all-auto-nfs* to support automounted file systems. To use this enhancement, set the *autofs* parameter to 1 in the `/etc/rc.config.d/nfsconf` file. Changing this parameter requires a reboot.

The following UNIX and Linux platforms support automounter, LOFS, or LOFS through automounter, as indicated:

Platform	automounter	LOFS	LOFS through automounter
AIX	yes	yes	yes
HP-UX	yes	yes	no
All Linux clients	yes	yes	yes
z/OS UNIX	yes	no	no
Solaris	yes	yes	yes

Note: For HP-UX 11i, you must use the new automounter, AutoFS. To activate AutoFS, set the *autofs* parameter to 1 in the `/etc/rc.config.d/nfsconf` file. Changing this parameter requires a reboot. If you encounter problems with

NFS automounted file systems, install patches PHCO_24777 and PHNE_26388 (or later patches). For further information, refer to the HP-UX documentation.


Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the `dsm.sys` file 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 file systems except LOFS file systems and LOFS through automounter. This is the default.

The `/tmp` directory is not included.

domain

Defines the file systems to include in your default client domain.

When you use *domain* with the **incremental** command, it processes these file systems in addition to those you specify in your default client domain.

-domain

Defines the file systems to exclude in your default client domain.

all-lofs

Backs up all loopback file systems, except those handled by automounter.

all-nfs

Backs up all network file systems, except those handled by automounter.

all-auto-nfs

Backs up all network file systems which are handled by automounter.

all-auto-lofs

Backs up all loop devices and local file systems which are handled through automounter.

Examples

Options file:

```
domain /tst /datasave /joe  
"domain all-local"  
domain ALL-LOCAL -/home  
domain ALL-NFS -/mount/nfs1
```

Command line:

```
-domain="/fs1 /fs2"  
-domain=/tmp  
-domain="ALL-LOCAL -/home"
```

Domain.image

The *domain.image* option specifies the mounted file systems and raw logical volumes that you want to include in your client domain for an image backup. Raw logical volumes must be named explicitly.

If you do not specify a file system with the **backup image** command, the file systems you specify with the *domain.image* option are backed up.

When you specify a file system with the **backup image** command, the *domain.image* option is ignored.

If you do not use the *domain.image* option to specify file systems in your client options file, and you do not specify a file system with the **backup image** command, a message is issued and no backup occurs.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, Solaris *only*. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client system options file (dsm.sys) or dsm.opt. You can set this option on the **Backup** category → **Domain for Backup** box of the Preferences editor.

Syntax



Parameters

domain

Defines the file systems or raw logical volumes to include in your default client image domain.

Examples

Options file:

```
domain.image /fs1 /fs2
```

Command line:

Does not apply.

Domain.nas

The *domain.nas* option specifies the volumes to include in your NAS image backups. You can specify *all-nas* to include all the mounted file systems on the NAS file server, except those you exclude with the *exclude.fs.nas* option. When you use this option in your client system options file (*dsm.sys*), the *domain.nas* option defines your default domain for NAS image backups.

Tivoli Storage Manager uses your domain for NAS image backups when you run a **backup nas** command and you do not specify which volumes to process.

When you perform a NAS file system image backup using the **backup nas** command, Tivoli Storage Manager adds volumes that you specify on the command line to the volumes defined in your *dsm.sys* file. For example, if you enter the following in your *dsm.sys* file:

```
domain.nas nas1/vol/vol0 nas1/vol/vol1
```

and you enter the following on the command line

```
dsmc backup nas -nasnodename=nas1 /vol/vol2
```

Tivoli Storage Manager backs up the *vol/vol0*, *vol/vol1*, and *vol/vol2* volumes on node *nas1*.

If you set the *domain.nas* option to *all-nas* in the *dsm.opt* file, Tivoli Storage Manager backs up all mounted volumes on the NAS file server. When performing a backup, if you use a file specification and set the *domain.nas* option to *all-nas* in the *dsm.sys* file, *all-nas* takes precedence.

Supported Clients

This option is valid for AIX and Solaris clients *only*. The server can also define this option.

Options File

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

Syntax



Parameters

domain

Defines the volumes you want to process. You cannot exclude volumes by specifying the dash (-) operator.

all-nas

Processes all mounted volumes on the NAS file server, except those you exclude with the *exclude.fs.nas* option. This is the default. If there is no *domain.nas* statement in the *dsm.opt* file and no volumes specified on the command line, Tivoli Storage Manager backs up all mounted volumes on the NAS server.

Examples

Options file:

```
domain.nas nas1/vol/vol0 nas1/vol/vol1  
domain.nas all-nas
```

Command line:

Does not apply.

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* the client user options file (*dsm.opt*).

The *dsmc* command line editor only supports Latin characters. If you want to enter any other characters (for example, DBCS), you must disable the Tivoli Storage Manager command line editor by setting the *editor* option to *no*.

Supported Clients

This option is valid for all UNIX and Linux clients (except z/OS UNIX System Services) clients.

Options 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.

Note: The editor is not supported on z/OS UNIX System Services, so the *editor* option should be set to *No*.

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.

Enable archiver retention protection

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
 - query TSA


Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

This option is valid only in the `dsm.sys` file *within* a server stanza and is not valid in a client option set from the server. It is not valid on any command line.

Syntax



Parameters

No The data retention server connection is refused. This is the default.

Yes The client connects to a data retention server.

Enablelanfree

Authorized User

The *enablelanfree* option specifies whether to enable an available LAN-free path to a storage area network (SAN) attached storage device. A LAN-free path allows backup, restore, archive, and retrieve processing between the Tivoli Storage Manager client and the SAN-attached storage device.

To support LAN-free data movement you must install and configure the Tivoli Storage Manager Managed System for SAN Storage Agent on the client workstation. For more information, refer to the following publications:

- *IBM Tivoli Storage Manager for AIX Storage Agent User's Guide*, SC32-0129
- *IBM Tivoli Storage Manager for Sun Solaris Storage Agent User's Guide*, SC32-0132
- *IBM Tivoli Storage Manager for HP-UX Storage Agent User's Guide*, SC32-0130
- *IBM Tivoli Storage Manager for Linux Storage Agent User's Guide*, SC32-0131

Notes:

1. If you place the *enablelanfree* option in the client option file (*dsm.opt*), but zero (0) bytes were transferred through the SAN during an operation, ensure that you bind the data to a LAN-free enabled management class.
2. To restore backup sets in a SAN environment, see "Restore Backupset" on page 470 for more information.

To specify a communication protocol between the Tivoli Storage Manager client and Storage Agent, see "Lanfreecommmethod" on page 276 for more information.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris clients.

Options File

Place this option in the *dsm.sys* file *within* a server stanza. You can set this option on the **General** category → **Enable Lanfree** check box of the Preferences editor.

Syntax



Parameters

- Yes* Specifies that you want to enable an available LAN-free path to a SAN-attached storage device.
- No* Specifies that you do not want to enable a LAN-free path to a SAN-attached storage device. This is the default.

Examples

Options file:
enablelanfree yes

Command line:

`-enablelanfree=yes`

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

Encryptiontype

The *encryptiontype* option allows you to use AES 128-bit data encryption, (not supported on z/OS), 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.

Notes:

1. 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.
2. HSM data migration only supports 56-bit DES encryption. Therefore, even though you enabled AES 128-bit encryption for the backup-archive client, migrated files will be backed up using 56-bit DES encryption.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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

Authorized User

The *encryptkey* option specifies whether to save the encryption key password locally when performing a backup-archive operation or whether to prompt for the encryption key password. The encryption key password is saved to the TSM.PWD file in encrypted format. 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.

If you set the *encryptkey* option to *save*, you are only prompted the first time you perform an operation. Thereafter, Tivoli Storage Manager does not prompt for the password.

If you do not save the encryption key password, you are prompted for the initial encryption key password when you begin encryption processing.

You can encrypt the data that is sent to the server during a backup or archive operation using standard encryption. If you use the encryption feature to encrypt your data during backup or archive, you *must* have the encryption key in order to restore or retrieve the data. If the encryption key is not available on the client machine (via the *encryptkey* option) and you forgot the encryption key, then the data *cannot* be restored or retrieved under any circumstances.

Table 63 shows how both Authorized Users and non-Authorized Users can encrypt or decrypt data during a backup or restore operation depending on the value specified for the *passwordaccess* option. The encryption key is saved when an encrypted backup operation is run by the Authorized User, *passwordaccess* is set to generate and *encryptkey* is set to *save*.

Note: When both encryption and compression are enabled, Tivoli Storage Manager compresses the data first, and then encrypts the data.

Table 63. Encrypting or decrypting data

Operation	Passwordaccess option	Encryptkey option	Result
Authorized user backup	generate	prompt	prompted for encryptkey password, data encrypted
	generate	save	data encrypted using the saved key
	prompt	prompt	prompted for encryptkey password, data encrypted
	prompt	save	data not encrypted
Non-authorized user backup	generate	prompt	encryption not authorized
	generate	save	data encrypted using the saved key
	prompt	prompt	encryption not authorized
	prompt	save	encryption not authorized

Table 63. Encrypting or decrypting data (continued)

Operation	Passwordaccess option	Encryptkey option	Result
Authorized or non-authorized user restore	generate	prompt	prompted for encryptkey password, data decrypted
	generate	save	data decrypted using the saved key
	prompt	prompt	prompted for encryptkey password, data decrypted
	prompt	save	prompted for encryptkey password, data decrypted

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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

Specifies that you want to save the encryption key password to the local TSM.PWD file. If you set the *encryptkey* option to *save*, you are only prompted the first time you perform an operation. Thereafter, Tivoli Storage Manager does not prompt for the encryption password. This is the default.

prompt

Tivoli Storage Manager prompts for the encryption password for each backup, archive, and restore operation.

Examples

Options file:

```
encryptkey prompt
```

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*.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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. The value for this option overrides the DSM_LOG environment variable. The dsmwebcl.log and dsmsched.log files are created in the same directory as the error log file you specify with the *errorlogname* option.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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 dsmerror.log file *cannot* be a symbolic link.

Examples

Options file:

```
errorlogname /tmp/tsmerror.log
```

Command line:

```
-errorlogname=/tmp/tsmerror.log
```

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

Errorlogretention

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.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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, image, 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. With the exception of *exclude.fs*, when you exclude a file that was previously included, existing backup versions become inactive during the next incremental backup.
2. The server can define exclude options with the *incl excl* option.

Exclude any system files or images that could corrupt the operating system when recovered. You should also exclude the directory containing the Tivoli Storage Manager client files.

Use wildcard characters to exclude a broad range of files. See “Including and excluding groups of files” on page 62 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 */any/test*, enter the following:

```
exclude.dir /any/test
```

To exclude subdirectories that begin with *test* under the */any* directory, enter the following:

```
exclude.dir /any/test*
```

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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 and space management services (if the HSM client is installed). The *exclude.backup* option only excludes files from normal backup, but not from HSM.

exclude.archive

Excludes a file or a group of files that match the pattern from archive services *only*.

exclude.attribute.symlink

Excludes a file or a group of files that are symbolic links 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 /test/dan/data1` excludes the `/test/dan/data1` directory, its files, and all its subdirectories and their files.

However, selective backup of a single file overrides `exclude.dir`. For example, you can still back up the `FILE1` file from `/test/dan/data1` using a selective backup, as follows:

```
dsmc sel -subdir=yes /test/dan/data1/
```

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.

exclude.fs

Excludes file systems that match the specified pattern from file system incremental backup operations and incremental image backup operations. If files from the excluded file systems were ever backed up, then management class rebinding and deleted file expiration will not occur. However, existing backup versions will remain on the server subject to associated management class settings.

Use **exclude.image** to exclude file systems from full image backup operations.

exclude.fs.nas

Excludes file systems on the NAS file server from an image backup when used with the **backup nas** command. The NAS node name must be prefixed to the file system name, for example: `netappsj1/vol/vol1`. To apply this exclude to all NAS nodes, replace the NAS node name with a wildcard, for example: `*/vol/vol1`. The **backup nas** command ignores all other exclude statements including *exclude.fs* and *exclude.dir* statements. This option is valid for AIX and Solaris clients *only* all Windows clients.

exclude.image

Excludes mounted file systems and raw logical volumes that match the specified pattern from full image backup operations.

Use *exclude.fs* to exclude file systems from incremental image backup operations.

Parameters

pattern

Specifies the file or group of files that you want to exclude.

Note: For NAS file systems: You must prefix the NAS node name to the file specification to specify the file server to which the exclude statement applies. If you do not specify a NAS node name, the file system identified refers to the NAS nodename specified in the client system options file (dsm.sys) or on the command line.

If the pattern begins with a single or double quote or contains any embedded blanks or equal signs, you must surround the value in either single (') or double (") quotation marks. The opening and closing quotation marks must be the same type of quotation marks.

For the *exclude.image* option, the pattern is the name of a mounted file system or raw logical volume.

Examples

Options file:

```
exclude /unix/  
exclude /.../core  
exclude /home/jones/proj1/*  
exclude.archive /.../core  
exclude.backup /home/jones/proj1/devplan/  
exclude.dir /home/jones/tmp  
exclude.backup /users/home1/file1  
exclude.image /usr/*/*  
exclude.encrypt /users/home2/file1  
exclude.compression /home/gordon/proj1/*  
exclude.fs.nas netappsj/vol/vol0  
exclude.attribute.symlink /.../*
```

Command line:

Does not apply.

Controlling symbolic link processing

Tivoli Storage Manager treats symbolic links as actual files and backs them up. However, the file referenced by the symbolic link is not backed up. In some cases symbolic links can be easily recreated and need not be backed up. In addition, backing up these symbolic links can increase backup processing time and occupy a substantial amount of space on the Tivoli Storage Manager server. You can use the *exclude.attribute.symlink* option to exclude a file or a group of files that are symbolic links from backup processing. If necessary, you can use the *include.attribute.symlink* option to include symbolic links within broad group of excluded files for backup processing. For example, to exclude all symbolic links from backup processing, except those that exist under the /home/spike directory, enter these statements in your dsm.sys file:

```
exclude.attribute.symlink /.../*  
include.attribute.symlink /home/spike/.../*
```

See "Include options" on page 269 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 211 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.fs*, *exclude.dir*, and other include-exclude statements first. Tivoli Storage Manager then considers any *exclude.compression* statements. For example, consider the following include-exclude list:

```
exclude /home/jones/proj1/*.*
exclude.compression /home/jones/proj1/file.txt
include /home/jones/proj1/file.txt
```

Tivoli Storage Manager examines the statements (reading from bottom to top) and determines that the `/home/jones/proj1/file.txt` file is a candidate for backup, but is not a candidate for compression processing.

- Include-exclude compression processing is valid for backup and archive processing *only*. The *exclude.compression* option does not affect whether files are excluded from backup or archive processing, only whether they are excluded from compression processing.

Processing NAS file systems

Use the *exclude.fs.nas* option to exclude file systems from NAS image backup processing.

A NAS file system specification uses the following conventions:

- NAS nodes represent a unique node type. The NAS node name uniquely identifies a NAS file server and its data to Tivoli Storage Manager. You can prefix the NAS node name to the file specification to specify the file server to which the exclude statement applies. If you do not specify a NAS node name, the file system identified applies to all NAS file servers.
- Regardless of the client platform, NAS file system specifications use the forward slash (/) separator, as in this example: `/vol/vol0`.

For example, to exclude `/vol/vol1` from backup services on all NAS nodes, specify the following exclude statement:

```
exclude.fs.nas */vol/vol1
```

Filelist

Use the *filelist* option with the following commands to process a list of files:

- **archive**
- **backup group**
- **delete archive**
- **delete backup**
- **expire**
- **incremental**
- **query archive**
- **query backup**
- **restore**
- **retrieve**
- **selective**

The Tivoli Storage Manager client opens the file you specify with this option and processes the list of files within according to the specific command. With the exception of the **restore** and **retrieve** commands, when you use the *filelist* option, Tivoli Storage Manager ignores all other file specifications on the command line.

The files (entries) listed in the filelist must adhere to the following rules:

- Each entry must be a fully or partially qualified path to a file or directory or a relative path.
- Each entry must be on a new line.
- Do *not* use wildcard characters.
- Each entry results in the processing of only one object (file or directory).
- If the file name contains any spaces, enclose the file name with quotes.
- Tivoli Storage Manager ignores any entry that is not valid.

The following is an example of a list of files within a filelist:

```
/home/dir/file1
/usr/tivoli/file2
/usr/avi/dir1
/fs1/dir2/file3
"/fs2/Ha Ha Ha/file.txt"
"/fs3/file.txt"
```

If an entry in the filelist indicates a directory, only that directory 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=/home/dir/file3 /usr/record/
```

the file specification `/usr/record/` represents the restore destination for all entries in the filelist. However, in the **selective** command:

```
selective -filelist=/home/dir/file3 /usr/record/
```

the file specification `/usr/record/` 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 `/home/dir/file1` and `/home/dir`. If the `/dir` directory does not exist on the server, the `/home/dir` directory is sent to the server a second time.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—FILEList == *filelistspec*—————▶▶

Parameters

filelistspec

Specifies the location and name of the file that contains the list of files to process with the command.

Note: When you specify the *filelist* option on the command line, the *subdir* option is ignored.

Examples

Command line:

```
sel -filelist=/home/avi/filelist.txt
```

Filename

Use the *filename* option with the **query systeminfo** command to specify a file name in which to store information gathered from one or more of the following items:

- DSMOPTFILE - The contents of dsm.opt file.
- DSMSYSFILE - The contents of the dsm.sys file.
- ENV - Environment variables.
- ERRORLOG - The Tivoli Storage Manager error log file.
- FILE - Attributes for the file name that you specify.
- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- OPTIONS - Compiled options.
- OSINFO - Name and version of the client operating system (includes ULIMIT information for UNIX and Linux).
- POLICY - Policy set dump.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually dsmsched.log).
- CLUSTER - AIX cluster information.

Note: The **query systeminfo** command is intended primarily as an aid for IBM support to assist in diagnosing problems, although users who are familiar with the concepts addressed by this information might also find it useful. If you use the *console* option, no special formatting of the output is performed to accommodate screen height or width. Therefore, the console output might be difficult to read due to length and line-wrapping. In this case, 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 213 for more information.

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

►►—FILENAME == *outputfilename*—◀◀

Parameters

outputfilename

Specifies a file name in which to store the information. If you do not specify a file name, by default the information is stored in the dsminfo.txt file.

Examples

Command line:

```
query systeminfo dsmoptfile errorlog -filename=tsminfo.txt
```

Filesonly

The *filesonly* option restricts backup, restore, retrieve, or query processing to files *only*. You cannot restore or retrieve directories from the Tivoli Storage Manager server when using the *filesonly* option with the **restore** or **retrieve** commands. However, directories with default attributes are created, if required, as placeholders for files that you restore or retrieve.

You can also use the *filesonly* option with the following commands:

- **archive**
- **incremental**
- **query archive**
- **query backup**
- **restore**
- **restore backupset**
- **restore group**
- **retrieve**
- **selective**

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶—FILESOnly—▶▶

Parameters

There are no parameters for this option.

Examples

Command line:
dsmc incremental -filesonly

Followsymbolic

The *followsymbolic* option specifies whether you want to restore files to symbolic links or use a symbolic link as a virtual mount point. Use this option with the **restore** and **retrieve** commands, or in the client user options file (dsm.opt).

The *followsymbolic* option does not determine whether Tivoli Storage Manager follows symbolic links during backup or archive operations. During a backup operation, symbolic links are never followed. During an archive operation, you can use the *archsymb linkasfile* option to specify whether Tivoli Storage Manager follows a symbolic link and archives the file or directory it points to, or archives the symbolic link only. See “Archsymlinkasfile” on page 196 for more information about the *archsymb linkasfile* option. See “Understanding how symbolic links are handled” on page 116 for more information about how Tivoli Storage Manager handles symbolic links.

Supported Clients

This option is valid for all UNIX clients.

Options File

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

Syntax



Parameters

No Specifies that you do not want to restore to symbolic links, or to use symbolic links as virtual mount points. This is the default.

Yes Specifies that you want to restore to symbolic links, or to use a symbolic link as a virtual mount point.

Examples

Options file:

```
followsymbolic Yes
```

Command line:

```
-fo1=Yes
```

Fromdate

Use the *fromdate* option with the *fromtime* option to specify a date and time from which you want to search for backups or archives during a restore, retrieve, or query operation. Files that were backed up or archived before this date and time are not included, although older directories might be included, if necessary, to restore or retrieve the files.

Use the *fromdate* option with the following commands:

- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**
- **restore was**
- **retrieve**

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—FROMDate =- *date*—————►►

Parameters

date

Specifies the date from which you want to search for backup copies or archived files. Enter the date in the format you selected with the *dateformat* option.

When you include *dateformat* with a command, it must precede the *fromdate*, *pitdate*, and *todate* options.

Examples

Command line:

```
dsmc query backup -fromdate=12/11/2003 /home/dilbert/*
```

Fromnode

The *fromnode* option permits one node to perform commands for another node. A user on another node must use the **set access** command to permit you to query, restore, or retrieve files or images for the other node.

Use the *fromnode* option with the following commands:

- **query archive**
- **query backup**
- **query filespace**
- **query group**
- **query image**
- **query mgmtclass**
- **query was**
- **restore**
- **restore group**
- **restore image**
- **restore was**
- **retrieve**

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

►►—FROMNode == *node*—————▶▶

Parameters

node

Specifies the node name on a workstation or a file server whose backup copies or archived files you want to access.

Examples

Command line:

```
dsmc query archive -fromnode=bob -subdir=yes "/home/jones/*"
```

Fromowner

The *fromowner* option specifies an alternate owner from which to restore backup versions or archived files or images. The owner must give access to another to use the files or images. For example, to restore files from the `/home/devel/proja` directory belonging to *usermike* on system *puma*, and place the restored files in a directory you own named `/home/id/proja`, enter:

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

Non-root users can specify *-fromowner=root* to access files owned by the root user if the root user has granted them access.

Note: If you specify the *fromowner* option without the *fromnode* option, the active user must be on the same node as the *fromowner* user.

See “How Tivoli Storage Manager handles long user and group names” on page 86 for information about the character limits for group names.

Use the *fromowner* option with the following commands:

- **query archive**
- **query backup**
- **query group**
- **query image**
- **query was**
- **restore**
- **restore image**
- **restore group**
- **restore was**
- **retrieve**

Supported Clients

This option is valid for all UNIX clients.

Syntax

►►—FROMowner == *owner*—◀◀

Parameters

owner

Name of an alternate owner.

Examples

Command line:

```
dsmc query archive "/home/id/proja/*" -fromowner=mark
```

Fromtime

Use the *fromtime* option with the *fromdate* option to specify a beginning time from which you want to search for backups or archives during a restore, retrieve, or query operation. Tivoli Storage Manager ignores this option if you do not specify the *fromdate* option.

Use the *fromtime* option with the following commands:

- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**
- **restore was**
- **retrieve**

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—FROMTime =- *time*—◄◄

Parameters

time

Specifies a beginning time on a specific date from which you want to search for backed up or archived files. If you do not specify a time, the time defaults to 00:00:00. Specify the time in the format you selected with the *timeformat* option.

When you include the *timeformat* option in a command, it must precede the *fromtime*, *pittime*, and *totime* options.

Examples

Command line:

```
dsmc q b -timeformat=4 -fromt=11:59AM -fromd=06/30/2003 -tot=11:59PM  
-tod=06/30/2003 /home/*
```

Groupname

Use the *groupname* option with the **backup group** command to specify the name for a group. You can only perform operations on new groups or the current active version of the group.

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

▶▶—GROUPName =- *name*—▶▶

Parameters

name

Specifies the name of the group which will contain the files backed up using the *filelist* option. Directory delimiters are not allowed in the group name since the group name is not a file specification, but a name field.

Examples

Command line:

```
backup group -filelist=/home/dir1/filelist1 -groupname=group1  
-virtualfsname=virtfs -mode=full
```

Groups

Authorized User

The *groups* option specifies UNIX and Linux system groups on your workstation that you want to authorize to request Tivoli Storage Manager services from the server. You can use the *groups* option more than once to specify several group names.

If you do not specify group names with the *groups* option, or user IDs with the *users* option, *all* users can request Tivoli Storage Manager services. If you use both the *groups* option and the *users* option, only users specified with these options can request Tivoli Storage Manager services. A root user is always authorized to request services.

See “How Tivoli Storage Manager handles long user and group names” on page 86 for information about the character limits for group names.

Supported Clients

This option is valid for all UNIX clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the *dsm.sys* file *within* a server stanza.

Syntax



Parameters

groupname

Specifies the name of a group you want to authorize to request Tivoli Storage Manager services.

Examples

Options file:

```
groups dsmcdev group1 test1 test2 design1
groups endicott almaden qadev qadev1 tools23
```

Command line:

Does not apply.

Guitreeviewafterbackup

The *guitreeviewafterbackup* option specifies whether the client returns to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.


Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

Place this option in the client user options file (*dsm.opt*) or the *dsm.sys* file. You can set this option on the **General** 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

Authorized User

The *httpport* option specifies a TCP/IP port address for the Web client.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the 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**
- **restore group**
- **restore was**
- **retrieve**

Note: This option is ignored if the *replace* option is set to *No*.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶—IFNewer—————▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/home/grover/*" -sub=y -rep=y -ifnewer
```

Imagetofile

Use the *imagetofile* option with the **restore image** command to specify that you want to restore the source image to a file. You might need to restore the image to a file if bad sectors are present on the target volume, or if you want to manipulate the image data. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume.

Linux supports mounting an image file as a logical volume, so you can get access to file data within the image. The following are some examples:

- The file system /usr has been backed up by Tivoli Storage Manager. The following command will restore the file system image to the file /home/usr.img:

```
# dsmc restore image /usr /home/usr.img -imagetofile
```

- To mount the image file at the /mnt/usr directory, the following mount command can be executed:

```
# mount /home/usr.img /mnt/usr -o loop=/dev/loop0
```

Now the image contents are available from /mnt/usr as if a regular file system was mounted at that directory.

Supported Clients

This option is valid for AIX, HP-UX, all Linux clients and Solaris. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶—IMAGETOfile—▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore image /usr /home/usr.img -imagetofile
```

Imagetype

Authorized User

Use the *imagetype* option with the **backup image** command or the *include.image* option to specify the type of image backup you want to perform. This option replaces the dependency on the copy serialization value in the management class to perform an image backup.

Attention: Tivoli Storage Manager dynamic image backups should not be used for certain file systems such as the AIX JFS2 file system, because the file system's internal design could provide inconsistent data even when there is no write activity. Dynamic image backup can easily result in a fuzzy image, and that image cannot be guaranteed to be valid or complete when restored. If the Tivoli Storage Manager client fails to mount the file system after restoring an image, you can run **fsck** to attempt to restore the file system's integrity. However, be aware that running **fsck** can affect the integrity of large amounts of data.

Place the *include.image* statement containing the *imagetype* value in your client system options file *dsm.sys*.

Supported Clients

This option is valid for AIX, Solaris, HP-UX, and all Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client system options file (*dsm.sys*). You can set this option on the **Include-Exclude** category of the Preferences editor.

Syntax

▶▶—IMAGEType==value—————▶▶

Parameters

value

Specifies one of the following values:

snapshot

Specifies that you want to perform a snapshot image backup during which the volume is available to other system applications. This is the default for file systems residing on a logical volume created by the Linux Logical Volume Manager. Valid for Linux clients *only*.

dynamic

Use this option only if the volume cannot be unmounted and remounted read-only. Tivoli Storage Manager backs up the volume *as is* without remounting it read-only. Corruption of the backup can occur if applications write to the volume while the backup is in progress. In this case, run **fsck** after a restore and manually mount the filesystem in order to regain access to the volume. This option is valid for AIX, Solaris, HP-UX, and all Linux clients.

static

Specifies that you want to perform an image backup during which the

volume is unmounted and remounted read-only. This is the default for AIX, HP-UX, and Solaris. This option is valid for AIX, Solaris, HP-UX, and all Linux clients

Examples

Options file:

```
include.image /home MYMC imagetype=static
```

Command line:

```
-imagetype=static
```

Inactive

Use the *inactive* option with the following commands to display both active and inactive objects:

- **delete group**
- **query asr**
- **query backup**
- **query group**
- **query image**
- **query nas**
- **query was**
- **restore**
- **restore group**
- **restore image**
- **restore nas**
- **restore was**

| **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.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶—INActive—▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/home/zoe/*" -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, such as */etc*.

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.

If you have the HSM client installed on your workstation, you can use an include-exclude options file to exclude files from backup and space management, from backup only or from space management only.

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 59.

Supported Clients

This option is valid for all all UNIX and Linux clients.

Options File

Place this option in the *dsm.sys* file *within* a server stanza. You can set this option on the **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:

```
incl excl /usr/dsm/backup.excl
incl excl /etc/incl excl.def
```

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, archive, image, and space management services.
- Files that are included for backup or archive processing that you want to include for encryption processing.
- Files that are included for backup or archive processing that you also want to include for compression processing.
- Objects to which you want to assign a specific management class.
- A management class to assign to all objects to which you do not explicitly assign a management class.
- File spaces to which you want to assign memory-efficient backup processing 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.fs* and *exclude.dir* statements override all include statements that match the pattern.
2. The server can also define these options with the *inclexcl* option.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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* —┐
 └ *mgmtclassname* —┘

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 /home/test/* archmc
include.backup /home/test/*
```

include.archive

Includes files or assigns management classes for archive processing.

include.attribute.symlink

Includes a file or a group of files that are symbolic links 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, your *dsm.opt* 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 *include.fs*, *memoryefficientbackup* and *diskcachelocation* below 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
```

include.fs.nas

Use the *include.fs.nas* option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system image backup, using the *toc* option with the *include.fs.nas* option in your *dsm.sys* file. See "Toc" on page 371 for more information. This option is valid for AIX and Solaris clients *only*.

include.image

Includes a file space or logical volume, or assigns a management class when used with the **backup image** command. The **backup image** command ignores all other include options.

You can use the *imagetype* option with the **backup image** command or the *include.image* option to specify whether to perform a static, dynamic, or snapshot (Linux x86/x86_64 only) image backup.

For Linux x86/x86_64 clients: Use the *snapshotcachesize* option with the **backup image** command, in the *dsm.opt* file, or with the *include.image* option to specify an appropriate snapshot size so that all old data blocks can be stored while the image backup occurs. A snapshot size of 100 percent will ensure a valid snapshot. See “Snapshotcachesize” on page 350 for more information.

This option is valid for AIX, HP-UX, Solaris, and all Linux clients.

Parameters

pattern

Specifies the objects to include for backup or archive processing or to assign a specific management class.

Note: For NAS file systems: You must prefix the NAS node name to the file specification to specify the file server to which the include statement applies. If you do not specify a NAS node name, the file system identified refers to the NAS node name specified in the client system options file (*dsm.sys*) or on the command line.

If the pattern begins with a single or double quote or contains any embedded blanks or equal signs, you must surround the value in either single (') or double (") quotation marks. The opening and closing quotation marks must be the same type of quotation marks.

For the *include.image* option, the pattern is the name of a mounted file system or raw logical volume.

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.

Examples

Options file:

```
include /home/proj/text/devel.*
include /home/proj/text/* textfiles
include * managall
include /WAS_ND_NDNODE mgmtclass
include /WAS_APPNODE mgmtclass
include.image /home/*/*
include.archive /home/proj/text/* myarchiveclass
include.backup /home/proj/text/* mybackupclass
include.compression /home/proj/text/devel.*
include.encrypt /home/proj/gordon/*
include.fs.nas netappsj/vol/vol0 homemgmtclass
include.image /home MGMTCLASSNAME type=snapshot
    snapshotcachesize=40
include.image /home imagetype=static
include.image /home imagetype=snapshot
include.image /home MGMTCLASSNAME imagetype=static
include.attribute.symlink /home/spike/.../*
include.fs /usr memoryefficientbackup=diskcachemethod
```

Command line:

Does not apply.

Controlling symbolic link processing

Tivoli Storage Manager treats symbolic links as actual files and backs them up. However, the file referenced by the symbolic link is not backed up. In some cases symbolic links can be easily recreated and need not be backed up. In addition, backing up these symbolic links can increase backup processing time and occupy a

substantial amount of space on the Tivoli Storage Manager server. You can use the *exclude.attribute.symlink* option to exclude a file or a group of files that are symbolic links from backup processing. If necessary, you can use the *include.attribute.symlink* option to include symbolic links within broad group of excluded files for backup processing. For example, to exclude all symbolic links from backup processing, except those that exist under the /home/spike directory, enter these statements in your dsm.sys file:

```
exclude.attribute.symlink /.../*
include.attribute.symlink /home/spike/.../*
```

See “Exclude options” on page 246 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 211 for more information.
- Tivoli Storage Manager processes *exclude.fs*, *exclude.dir*, and other include-exclude statements first. Tivoli Storage Manager then considers any *include.compression* and *include.encrypt* statements. For example, consider the following include-exclude list:

```
exclude /home/jones/proj1/file.txt
include.compression /home/jones/proj1/file.txt
include.encrypt /home/jones/proj1/file.txt
```

Tivoli Storage Manager examines the `exclude /home/jones/proj1/file.txt` statement first and determines that `/home/jones/proj1/file.txt` is excluded from backup processing and is, therefore, not a candidate for compression and encryption processing.

- Include-exclude compression and encryption processing is valid for backup and archive processing *only*.

Processing NAS file systems

Use the *include.fs.nas* option to bind a management class to NAS file systems and to control if Table of Contents information is saved for the file system backup.

A NAS file system specification uses the following conventions:

- NAS nodes represent a new node type. The NAS node name uniquely identifies a NAS file server and its data to Tivoli Storage Manager. You can prefix the NAS node name to the file specification to specify the file server to which the include statement applies. If you do not specify a NAS node name, the file system you specify applies to all NAS file servers.
- Regardless of the client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.

Use the following syntax:

►► *pattern*– *mgmtclassname*– *toc=value* ◀◀

Parameters

pattern

Specifies the objects to include for backup services, to assign a specific management class, or to control TOC creation. You can use wildcards in the pattern.

mgmtclassname

Specifies the name of the management class to assign to the objects. If a management class is not specified, the default management class is used.

toc=value

See “Toc” on page 371 for more information.

For example, to assign a management class to the /vol/vol1 file system of a NAS node called netappsj, specify the following include statement:

```
include.fs.nas netappsj/vol/vol1 nasMgmtClass toc=yes
```

Processing WebSphere Application Server (WAS) file systems

You can use the *include* option in your dsm.sys file to assign a management class to a WAS group backup. For example:

- For the Network Deployment Manager: `include /WAS_ND_NDNODE mgmtclass`
- For the Application Server: `include /WAS_APPNODE mgmtclass`

See “Creating an include-exclude list (optional root user or authorized user task)” on page 59 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, such as Access control list (ACL) data, unless the modification dates and times have also changed.
- It ignores the copy group frequency attribute of management classes.

Note: If you have limited time during the week to perform backups, but extra time on weekends, you can maintain current server storage of your workstation files by performing an incremental backup with the *incrbydate* option on weekdays and a full incremental backup on weekends.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—INCRbydate—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:
dsmc incremental -incrbydate

Incremental

Use the *incremental* option with the **restore image** command to ensure that any changes that were made to the base image are also applied to the restored image.

If you also use the *deletefiles* option, changes include the deletion of files and directories that were in the original image but later deleted from the workstation.

Note: Using the *incremental* option with the **restore image** command to perform a dynamic image backup is not supported.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris *only*. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—INCREmental—————►►

Examples

Command line:

```
res i "/home/devel/projecta/*" -incremental
```

Lanfreecommmethod

Authorized User

The *lanfreecommmethod* option specifies the communications protocol between the Tivoli Storage Manager client and Storage Agent. This enables processing between the client and the SAN-attached storage device.

If you are using LAN failover, you must have *lanfreecommmethod TCPip* in the *dsm.sys* file *within* a server stanza.

For AIX, HP-UX, Linux and Solaris, use the *lanfreeshmport* to specify the Shared Memory port number where the Storage Agent is listening. See “Lanfreeshmport” on page 278 for more information.

Supported Clients

This option is valid for AIX, HP-UX, Linux, and Solaris clients *only*.

Options File

Place this option in the *dsm.sys* file *within* a server stanza.

Syntax

▶▶—LANFREECOMMmethod— *commmethod* —————▶▶

Parameters

commmethod

Specifies the supported protocol for your Tivoli Storage Manager client:

TCPip

The Transmission Control Protocol/Internet Protocol (TCP/IP) communication method.

Use the *lanfreetcport* option to specify the TCP/IP port number where the Storage Agent is listening. See “Lanfreetcport” on page 279 for more information. The TCP/IP communication method is the default for *non-root* users on all supported platforms. On Linux, this is the only communication method you can use.

SHAREdmem

Use the Shared Memory communication method when the client and Storage Agent are running on the same system. Shared Memory provides better performance than the TCP/IP protocol. This is the default communication method for AIX, HP-UX, and Solaris *root* users. When specifying this communication method on AIX, the backup-archive client user can be logged in as root or non-root, as long as the Storage Agent is running as root. If the Storage Agent is not running as root, the user ID running the backup-archive client must match the user ID running the Storage Agent.

Examples

Options file:

```
lanfreec tcp
```

Command line:

`-lanfreec=tcp`

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

Lanfreeshmport

Authorized User

Use the *lanfreeshmport* option when *lanfreecommmethod*=*SHAREdmem* is specified for communication between the Tivoli Storage Manager client and Storage Agent. This enables processing between the client and the SAN-attached storage device. See “Lanfreecommmethod” on page 276 for more information about the *lanfreecommmethod* option.

Supported Clients

This option is valid for AIX, HP-UX, Linux and Solaris clients *only*.

Options File

Place this option in the client system options file (*dsm.sys*) *within* a server stanza.

Syntax

▶▶—LANFREESHmport— *port_address* —————▶▶

Parameters

port_address

Specifies the number that is used to connect to the storage agent. The range of values is 1 through 32767; the default is 1.

Examples

Options file:

```
lanfrees 1520
```

Command line:

```
-lanfrees=1520
```

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

Lanfreetcport

Authorized User

The *lanfreetcport* option specifies the TCP/IP port number where the Tivoli Storage Manager Storage Agent is listening.

Use this option when you specify *lanfreecommmethod=TCPIP* for communication between the Tivoli Storage Manager client and Storage Agent. Do not specify the *lanfreetcport* option if you want to use the NAMEDpipes communication method for LAN-free communication. See “Lanfreecommmethod” on page 276 for more information about the *lanfreecommmethod* option.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris clients *only*.

Options File

Place this option in the dsm.sys file *within* a server stanza.

Syntax

▶▶—LANFREETCport— *port_address* —————▶▶

Parameters

port_address

Specifies the TCP/IP port number where the Storage Agent is listening. The range of values is 1000 through 32767; the default is 1500.

Note: The client *lanfreetcport* value must match Storage Agent *tcport* value for communications with the Storage Agent (virtual server). The client *tcport* value must match the server *tcport* value for communications with the actual server.

Examples

Options file:

```
lanfreetcp 1520
```

Command line:

```
-lanfreetcp=1520
```

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

Lanfreetcpserveraddress

Authorized User

The *lanfreetcpserveraddress* option specifies the TCP/IP address for a Tivoli Storage Manager Storage Agent. Use this option when you specify *lanfreecommmethod=TCPIP* for communication between the Tivoli Storage Manager client and Storage Agent. Overriding the default for this option is useful when configuring LAN-free in an environment where the client and storage agent are running on different systems. You can obtain this Storage Agent address from your administrator.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, and Solaris clients *only*.

Options File

Place this option in the client system options file.

Syntax

▶▶—LANFREETCPServeraddress— *stagent_address* —————▶▶

Parameters

stagent_address

Specifies a 1 to 64 character TCP/IP address for a server. Specify a TCP/IP domain name or a dot address. Default value is 127.0.0.1 (localhost).

Examples

Options file:

```
LANFREETCPServeraddress stagent.bayamon.ibm.com
```

```
LANFREETCPServeraddress 192.168.1.50
```

Command line:

Does not apply.

Latest

Use the *latest* option with the following commands to restore the most recent backup version of a file, even if the backup is inactive:

- **restore**
- **restore group**
- **restore was**

If you are performing a point-in-time restore (using the *pitdate* option), it is not necessary to specify *latest* since this option is implicit when *pitdate* is used.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶—LATEST—▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/home/devel/projecta/*" -latest
```

Localbackupset

The *localbackupset* option specifies whether the Tivoli Storage Manager GUI bypasses initial logon with the Tivoli Storage Manager server to restore a local backup set on a standalone workstation. 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 470 for more information.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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`

This option is not valid with the `dsmsc` command line client.

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 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.

Use the *location* option with the **query** and **restore** commands.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

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.

tape

Specifies that Tivoli Storage Manager searches for the backup set on a local tape device. Specifying *location=tape* covers all tape device types. This parameter is valid for Solaris, AIX, and HP-UX clients.

Note: For Solaris, AIX, and HP-UX, if you want to restore a backup set from a 3570 or 3590 tape device, but you do not have the 3570 or 3590 generic device driver on your client workstation, you can download these device drivers from the following Web site:

<ftp://ftp.software.ibm.com/storage/devdrv/>

For Solaris: Use tapes that are fully compliant with Sun standards.

Examples

Command line:

```
restore backupset "/dev/rmt0" -loc=tape
restore backupset mybackupsetname -loc=server
restore backupset /home/budget/backupsetfile.name
-loc=file
```

Mailprog

Authorized User

The *mailprog* option passes the information that is specified in the option value to the operating system to execute.

Supported Clients

This option is valid for all UNIX clients.

Options File

Place this option in the *dsm.sys* file *within* a server stanza.

Syntax

▶▶—MAILprog *command*—————▶▶

Parameters

command

Indicates how you want to invoke a mail program.

Examples

Options file:

```
MAILPROG /usr/bin/mailx -c cc@host1 -b bcc@host2 root@host3
```

On UNIX, this example would invoke the mailx program to send email with CC and BCC copies, because any number of parameters are acceptable. On z/OS and OS/400, only two parameters would be accepted: program name and recipient's email.

Command line:

Does not apply.

Makesparsefile

Use the *makesparsefile* option with the **restore** or **retrieve** commands to specify how sparse files are recreated. Sparse files do not have disk space allocated for every block in the whole address space, leading to holes within the file. The Tivoli Storage Manager client detects sparse files during a backup operation and marks them as sparse on the Tivoli Storage Manager server. Holes are detected by their content, which is always zeros.

If you set the *makesparsefile* option to *yes* (default), holes within the file are not written to disk so no additional disk space is allocated during a restore.

If you set the *makesparsefile* option to *no*, holes are not recreated, leading to disk blocks allocated for the whole address space. This might result in a larger amount of used disk space. Ensure that you have enough disk space to restore all data.

On some UNIX and Linux systems, it might be necessary to back up system specific files as non-sparse files. Use the *makesparsefile* option for files where the existence of physical disk blocks is required, such as *ufsboot* on Solaris, which is executed during boot time. The boot file loader of the operating system accesses physical disk blocks directly and does not support sparse files.

Supported Clients

This option is valid for all UNIX clients.

Options File

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

Syntax



Parameters

- Yes* Specifies that holes within the file are not written so that no additional disk space is allocated during a restore. This is the default.
- No* Specifies that holes are not recreated leading to disk blocks allocated for the whole address space.

Examples

Options file:
makesparsefile no

Command line:
-makesparsefile=no

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

Managedservices

Authorized User

The *managedservices* option specifies whether the Tivoli Storage Manager Client Acceptor daemon (CAD) manages the scheduler, the Web client, or both.

See “Configuring the client scheduler” on page 51 for instructions to set up the CAD 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 344 for more information.

The CAD serves as an external timer for the scheduler. When the scheduler is started, it queries the server for the next scheduled event. The event is either executed immediately or the scheduler exits. The CAD restarts the scheduler when it is time to execute the scheduled event.

Notes:

1. If you set the *shedmode* option to *prompt*, the server prompts the CAD when it is time to run the schedule. The scheduler will connect and disconnect to the server when the CAD is first started.
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 305 for more information.

Using the CAD to manage the scheduler service can provide the following benefits:

- Memory retention problems that can occur when using traditional methods of running the scheduler are resolved. Using the CAD to manage the scheduler requires very little memory between scheduled operations.
- The CAD 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 CAD manages the Web client only.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the 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

►►—MANAGEDServices—mode—webclient
schedule—►►

Parameters

mode

Specifies whether the CAD manages the scheduler, the Web client, or both.

webclient

Specifies that the CAD manages the Web client. This is the default.

schedule

Specifies that the CAD manages the scheduler.

Examples

Options file:

The following are examples of how you might specify the *managementservices* option in your client system options file (dsm.sys).

Task Specify that the CAD manages the Web client *only*.
managementservices webclient

Task Specify that the CAD manages the scheduler *only*.
managementservices schedule

Task Specify that the CAD manages both the Web client and the scheduler.
managementservices 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.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file. 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 269 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 2 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.


Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

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

Mode

Use the *mode* option with these commands, as follows:

backup image

To specify whether to perform a selective or incremental image backup of client file systems.

backup nas

To specify whether to perform a full or differential image backup of NAS file systems.

backup was

To specify whether to perform a full or differential backup of the WebSphere Application Server (WAS) Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server (also contains setup, application files, and configuration information) to the Tivoli Storage Manager server.

backup group

To specify whether to perform a full or differential group backup containing a list of files from one or more file space origins.

The *mode* option has no effect on a raw logical device backup.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86/x86_64, Linux on POWER, Linux zSeries and Solaris *only*. The Tivoli Storage Manager client API does not support this option.

Syntax

For image backup of client file systems (AIX, Solaris, HP-UX, Linux x86/x86_64, and Linux on POWER)



For image backup of NAS file systems (AIX and Solaris)



For group and WAS backups (AIX, Solaris, Linux x86/x86_64)



Parameters

selective

Specifies that you want to perform a full (selective) image backup. This is the default for image backup of client file systems.

incremental

Specifies that you want to back up only new and changed files after the most recent image backup (full or incremental). If an image backup has not yet been made, then the first backup will be a full image backup, regardless of the mode option.

Deleted files are not marked inactive on the server.

full

Specifies that you want to perform a full backup of NAS, WAS, or group objects. This is the default for WAS and group backups.

differential

This is the default for NAS objects. Specifies that you want to perform a NAS, WAS, or group backup of files that changed since the last full backup. If there is no copy of a full image stored on the Tivoli Storage Manager server, a full backup occurs. If a full image exists, whether it is restorable, or expired and being maintained because of dependent differential images, specifying **MODE=differential** sends a differential image backup. If a full image is sent during a *differential* backup, it is reflected as a full image using the **query nasbackup** server command.

A full image can be eligible for expiration based on versioning or retention (verexists reextra), but still be maintained on the Tivoli Storage Manager server to allow for restoring dependent differential images. A full image that is eligible for expiration *cannot* be selected for restore, so it is not displayed using the **query nasbackup** server command. The differential image backups that depend on an "expired" full image can be restored.

Examples

Task Perform the NAS image backup of the entire file system.

Command: `dsmc backup nas -mode=full -nasnodename=nas1 /vol/vol0 /vol/vol1`

Task Back up the /home/test file space using an image incremental backup that backs up only new and changed files after the last full image backup.

Command: `dsmc backup image /home/test -mode=incremental`

Task Perform a full backup of all the files in filelist /home/dir1/filelist1 to the virtual file space name /virtfs containing the group leader /home/group1 file.

Command:

`backup group -filelist=/home/dir1/filelist1 -groupname=group1 -virtualfsname=virtfs -mode=full`

Monitor

The *monitor* option specifies whether to monitor an image backup or restore of file systems belonging to a Network Attached Storage (NAS) file server.

If you specify *monitor=yes*, Tivoli Storage Manager monitors the current NAS image backup or restore operation and displays processing information on your screen. This is the default.


If you specify *monitor=no*, Tivoli Storage Manager does not monitor the current NAS image backup or restore operation and is available to process the next command.

Use this option with the **backup nas** or **restore nas** commands.

Supported Clients

This option is valid for AIX and Solaris clients *only*.

Syntax



Parameters

- Yes* Specifies that you want to monitor the current NAS image backup or restore operation and display processing information on your screen. This is the default.
- No* Specifies that you do not want to monitor the current NAS image backup or restore operation.

Examples

Command line:

```
backup nas -mode=full -nasnodename=nas1 -monitor=yes  
/vol/vol0 /vol/vol1
```

Nasnodename

The *nasnodename* option specifies the node name for the NAS file server when processing NAS file systems. The node name identifies the NAS file server to the Tivoli Storage Manager server. The server must register the NAS file server.

You can specify this option on the command line or in the client system options file (*dsm.sys*).

You can override the default value in the *dsm.sys* file by entering a different value on the command line. If you do not specify the *nasnodename* option in the *dsm.sys* file, you *must* specify this option on the command line when processing NAS file systems.

You can use the *nasnodename* option with the following commands:

- **backup nas**
- **delete filesystem**
- **query backup**
- **query filesystem**
- **restore nas**

Supported Clients

This option is valid for the AIX and Solaris clients *only*. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client system options file (*dsm.sys*). You can set this option on the **General** category of the Preferences editor.

Syntax

▶▶—NASNodename— *nodename* —————▶▶

Parameters

nodename

Specifies the node name for the NAS file server.

Examples

Options file:

```
nasnodename nas2
```

Command line:

```
-nasnodename=nas2
```

Nfstimeout

Authorized User

The *nfstimeout* option specifies the number of seconds the client waits for a status system call on an NFS file system before it times out.

You can use this option to mitigate the default behavior of status calls on file systems. For example, if an NFS file system is stale, a status system call 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.

Note: On Solaris and HP-UX, the *nfstimeout* option can fail if the NFS mount is hard. If a hang occurs, deactivate the *nfstimeout* option and mount the NFS file system soft mounted, as follows:

```
mount -o soft,timeo=5,retry=5 machine:/filesystem /mountpoint
```

The parameters are defined as follows:

soft Generates a soft mount of the NFS file system. If an error occurs, the `stat()` function returns with an error. If the option `hard` is used, `stat()` does not return until the file system is available.

timeo=n
Sets the time out for a soft mount error to `n` seconds

retry=n
Set the internal retries and the mount retries to `5`, the default is `10000`.

See “Understanding how NFS hard and soft mounts are handled” on page 118, SC32-0145, for a discussion of how NFS hard and soft mounts are handled.

Supported Clients

This option is for all UNIX clients. The server can also define this option.

Options File

Place this option in the `dsm.sys` file within a server stanza *or* the client options file (`dsm.opt`).

Syntax

►►—NFSTIMEout *number*—————►►

Parameters

number
Specifies the number of seconds the client waits for a status system call on a file system before timing out. The range of values is 0 through 120; the default is 0 seconds.

Examples

Options file:

```
nfstimeout 10
```

Command line:

```
-nfstimeout=10
```

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

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 384 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.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

Place this option in the `dsm.sys` file *within* a server stanza. You can set this option on the **General** 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. If you set the *clusternode* option to *yes*, the default is the cluster name.

Examples

Options file:
nodename cougar

Nojournal

Use the *nojournal* option with the **incremental** command to specify that you want to perform a traditional full incremental backup, instead of the default journal-based backup.

Journal-based incremental backup differs from the traditional full incremental backup in the following ways:

- Tivoli Storage Manager does not enforce non-default copy frequencies (other than 0).
- UNIX special file changes are not detected by the Journal daemon and are not, therefore, backed up.

For these reasons, you want to use the *nojournal* option periodically to perform a traditional full incremental backup.

See “Journal-based backup” on page 431 for more information about journal-based incremental backup.

Supported Clients

This option is valid for the AIX Backup-Archive client.

Syntax

▶▶—NOJournal—◀◀

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc incr /home -nojournal
```

Noprompt

The *noprompt* option suppresses the confirmation prompt that is presented by the **delete group**, **delete archive**, **expire**, **restore image**, and **set event** commands.

- **delete archive**
- **delete backup**
- **delete group**
- **expire**
- **restore image**

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶—NOPrompt—————▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc delete archive -noprompt "/home/project/*"
```

Numberformat

The *numberformat* option specifies the format you want to use to display numbers.

The AIX, Solaris, and HP-UX clients support locales other than English that describe every user interface that varies with location or language. See Table 18 on page 47 for supported locales.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time the client is called. Consult the documentation on your local system for details about setting up your locale definition.

Note: The *numberformat* option does not affect the Web client. The Web client uses the number format for the locale that the browser is running in. If the browser is not running in a supported locale, the Web client uses the number format for US English.

You can use the *numberformat* option with the following commands:

- **delete archive**
- **delete backup**
- **expire**
- **query archive**
- **query asr**
- **query backup**
- **query image**
- **query nas**
- **restore**
- **restore image**
- **restore nas**
- **retrieve**
- **set event**

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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

This is the default for the following available translations:

- US English

```

|           • Japanese
|           • Chinese (Traditional)
|           • Chinese (Simplified)
|           • Korean
|           2 1,000,00
|           3 1 000,00
|
|           This is the default for the following available translations:
|           • French
|           • Czech
|           • Hungarian
|           • Polish
|           • Russian
|           4 1 000.00
|           5 1.000,00
|
|           This is the default for the following available translations:
|           • Brazilian Portuguese
|           • German
|           • Italian
|           • Spanish
|           6 1'000,00

```

For AIX, HP-UX, and Solaris: To define number formats, modify the following lines in the source file of your locale. Whatever format you select applies both to output and to input.

- decimal_point**
The character that separates the whole number from its fractional part.
- thousands_sep**
The character that separates the hundreds from the thousands from the millions.
- grouping**
The number of digits in each group that is separated by the thousands_sep character.

Examples

Options file:
num 4

Command line:
-numberformat=4

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This will be the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

Optfile

The *optfile* option specifies the client user options file you want to use when you start a Tivoli Storage Manager session.

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

▶▶—OPTFILE =- *file_name*—————▶▶

Parameters

file_name

Specifies an alternate client options file, if you use the fully qualified path name. If you specify only the file name, Tivoli Storage Manager assumes the file name specified is located in the current working directory. The default is dsm.opt.

Restriction: Specify the full path when you use this option with the Client Acceptor Daemon (dsmcad), because the CAD changes its working directory to root ("/") after initialization.

Examples

Command line:

```
dsmc query session -optfile=myopts.opt
```

CAD:

```
dsmcad -optfile=/usr/tivoli/tsm/client/ba/bin/myopts.opt
```

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

Password

The *password* option specifies a Tivoli Storage Manager password. If you do not specify this option and your administrator has set authentication to *On*, you are prompted for a password when you start a Tivoli Storage Manager session.

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*.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

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

Syntax

▶—PASsword— *password* —▶

Parameters

password

Specifies a 1 to 63 character password. A password is not case-sensitive. Valid characters include:

Characters

	Description
A–Z	Any letter, A through Z, uppercase or lowercase
0–9	Any number, 0 through 9
+	Plus
.	Period
_	Underscore
-	Hyphen
&	Ampersand

Examples

Options file:

```
password secretword
```

Command line:

```
-password=secretword
```

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

Passwordaccess

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 307.

When the *passwordaccess* option is set to *generate* and you specify the *password* option, the *password* option is ignored.

Setting the *passwordaccess* option to *generate* is required in the following situations:

- When using the HSM client.
- When using the Web client.
- When performing NAS operations.
- When using Tivoli Continuous Data Protection for Files.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

Place this option in the *dsm.sys* file *within* a server stanza. You can set this option on the **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.

API applications must supply the password when a session is initiated. The application is responsible for obtaining the password.

generate

Encrypts and stores your password locally and generates a new password when the old password expires.

A password prompt is displayed when registering a workstation with a server using open registration or if your administrator changes your password manually.

You can use the *mailprog* option to specify the program and user ID where you want to send the new password each time the old password expires.

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. . Regardless of where it is stored, the password file created by Tivoli Storage Manager is always named TSM.PWD.

Supported Clients

This option is valid for all UNIX clients.

Options File

Place this option in the client system options file (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.

The default directory for AIX is /etc/security/adsm and for other UNIX and Linux platforms it is /etc/adsm.

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 /etc/security/tsm
```

```
passworddir "/Users/user1/Library/Preferences/Tivoli Storage Manager/"
```

Command line:

Does not apply.

Pick

The *pick* option creates a list of backup versions, images, or archive copies that match the file specification you enter. From the list, you can select the versions to process. Include the *inactive* option to view both active and inactive objects.

For images, if you do not specify a source file space and destination file space, the pick list contains all backed up images. In this case, the images selected from the pick list are restored to their original location. If you specify the source file space and the destination file space, you can select only one entry from the pick list.

Use the *pick* option with the following commands:

- **delete archive**
- **delete backup**
- **delete group**
- **expire**
- **restore**
- **restore group**
- **restore image**
- **restore nas**
- **restore was**
- **retrieve**

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶—Pick—————▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/home/project/*" -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 or images 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 asr**
- **query backup**
- **query group**
- **query image**
- **query nas**
- **query was**
- **restore**
- **restore group**
- **restore was**
- **restore image**
- **restore nas**

When *pitdate* is used, the *inactive* and *latest* options are implicit.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—PITDate == *date*—————▶▶

Parameters

date

Specifies the appropriate date. Enter the date in the format you selected with the *dateformat* option.

When you include *dateformat* with a command, it must precede the *fromdate*, *pitdate*, and *todate* options.

Examples

Command line:

```
dsmc restore "/fs1/*" -sub=y -pitdate=08/01/2003 -pittime=06:00:00
```

Pittime

Use the *pittime* option with the *pitdate* option to establish a point in time for which you want to display or restore the latest version of your backups. Files or images that were backed up *on or before* the date and time you specify, and which were not deleted *before* the date and time you specify, are processed. Backup versions that you create after this date and time are ignored. This option is ignored if you do not specify *pitdate* option.

Use the *pittime* option with the following commands:

- **delete backup**
- **query asr**
- **query backup**
- **query image**
- **query nas**
- **restore**
- **restore image**
- **restore nas**
- **restore was**

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—PITtime == *time*—————▶▶

Parameters

time

Specifies a time on a specified date. If you do not specify a time, the time defaults to 23:59:59. Specify the time in the format you selected with the *timeformat* option.

When you include the *timeformat* option in a command, it must precede the *fromtime*, *pittime*, and *totime* options.

Examples

Command line:

```
dsmc q b "/fs1/*" -pitt=06:00:00 -pitd=08/01/2003
```

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).

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the *dsm.sys* file. 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 double quotes. If you placed double quotes 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.

Examples

Options file:

```
postschedulecmd "restart database"
```

The command string is a valid command for restarting your database.

Command line:

```
-postschedulecmd="'restart database'"
```

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

Postsnapshotcmd

The *postsnapshotcmd* option allows you to run operating system shell commands after the LVM starts a snapshot. This allows you to bring up an application after the LVM starts a snapshot during an online image backup operation. This option is only valid if the LVM is installed and configured on your system allowing you to perform an online image backup operation.

For an online image backup, use this option with the **backup image** command, the *include.image* option, or in the *dsm.sys* file.

If the *postsnapshotcmd* fails the volume image backup continues, but appropriate warnings are logged.

Supported Clients

This option is valid for Linux x86/x86_64 clients only. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (*dsm.opt*). You can set this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

►—POSTSNAPSHOTCMD— "*cmdstring*" —————►

Parameters

"cmdstring"

Specifies a command to process. Use the *svoprepostsnapdisabled* option to prevent the Tivoli Storage Manager server administrator from executing operating system commands on the client system. See "Srvprepostscheddisabled" on page 354 for details on this option.

If the command string contains blanks, enclose the command string in double quotes:

```
"net stop someservice"
```

If you placed double quotes within the command string, then enclose the entire command string in single quotes:

```
'net stop "some service"'
```

Examples

Options file:

```
postsnapshotcmd "any command"
```

The command string is a valid command for restarting your application.

Command line:

```
backup image -postsnapshotcmd="any command"
```

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

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).

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the `dsm.sys` file. You can set this option on the **Scheduler** category, **Schedule Command** button of the Preferences editor.

Syntax

►► — PRESchedulecmd — "cmdstring" ————— ◀◀
 └─ Prenschedulecmd ─┘

Parameters

"cmdstring"

Specifies the command to process. Use only one *preschedulecmd* option. You can enter a command to be executed before a schedule using this option.

If the command string contains blanks, enclose the command string in double quotes. If you placed double quotes 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.

Examples

Options file:

```
preschedulecmd "<your database product's quiesce command>  
database"
```

The command string is a valid command for quiescing your database.

Command line:

```
-preschedulecmd="'quiesce database'"
```

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

Preservelastaccessdate

Any application that touches a file can implicitly cause that file's last access date to change to the time that the application touches it. This is a function of the file system, not the application. Because of this, when the client backs up or archives a file, it can trigger an update to the file's last access date. This can cause problems for other applications such as Storage Resource Management (SRM) or Hierarchical Storage Management, whose processing relies on accurate last access dates.

Use the *preservelastaccessdate* option during a backup or archive operation to specify whether to reset the last access date of any specified files to their original value following the backup or archive operation. By default, the Tivoli Storage Manager client *will not* reset the last access date of any backed up or archived files to their original value following the backup or archive operation.

Use this option with the **incremental**, **selective**, or **archive** commands.

Notes:

1. This option only applies to files; it does not apply to directories.
2. Resetting the last access date incurs additional overhead that can impact backup and archive performance. The last access date should be reset only if you are using another application, such as a Storage Resource Management (SRM) or Hierarchical Storage Management that relies on accurate last access dates.
3. You cannot reset the last access date of read-only files. The *preservelastaccessdate* option ignores read-only files and does not change their date.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

Place this option in the client user options file (dsm.opt). You can set this option on the **Backup** category of the Preferences editor.

Syntax



Parameters

- No* Specifies that the Tivoli Storage Manager client *will not* reset the last access date of any backed up or archived files to their original value following the backup or archive operation. This is the default.
- Yes* Specifies that the Tivoli Storage Manager *will* reset the last access date of any backed up or archived files to their original value following the backup or archive operation.

Examples

Options file:

```
preservelastaccessdate yes
```

Command line:

```
Incremental /proj/test/test_file -preservelastaccessdate=yes
```

Preservepath

The *preservepath* option specifies how much of the source path to reproduce as part of the target directory path when you restore or retrieve files to a new location. Use the *-subdir=yes* option to include the entire subtree of the source directory (directories and files below the lowest-level source directory) as source to be restored. If a required target directory does not exist, it is created. If a target file has the same name as a source file, it is overwritten. Use the *-replace=prompt* option to have Tivoli Storage Manager prompt you before files are overwritten.

Use the *preservepath* option with the following commands:

- **restore**
- **restore backupset**
- **restore group**
- **retrieve**
- **restore was**

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

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:

```
/fs/h1/m1/file.a
/fs/h1/m1/file.b
/fs/h1/m1/l1/file.x
/fs/h1/m1/l1/file.y
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -preser=complete
```

Restores these directories and files:

```
/u/ann/h1/m1/file.a
/u/ann/h1/m1/file.b
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -preser=nobase
```

Restores these directories and files:

```
/u/ann/file.a
/u/ann/file.b
```

This command:

```
dsmc res backupset /fs/h1/m1/ /u/ann/ -su=yes
-preser=nobase -loc=file
```

Restores these directories and files:

```
/u/ann/file.a
/u/ann/file.b
/u/ann/file.x
/u/ann/file.y
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -preser=subtree
```

Restores these directories and files:

```
/u/ann/m1/file.a
/u/ann/m1/file.b
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -preser=none
```

Restores these directories and files:

```
/u/ann/file.a
/u/ann/file.b
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -su=yes -preser=complete
```

Restores these directories and files:

```
/u/ann/h1/m1/file.a
/u/ann/h1/m1/file.b
/u/ann/h1/m1/l1/file.x
/u/ann/h1/m1/l1/file.y
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -su=yes -preser=nobase
```

Restores these directories and files:

```
/u/ann/file.a
/u/ann/file.b
/u/ann/l1/file.x
/u/ann/l1/file.y
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -su=yes -preser=subtree
```

Restores these directories and files:

```
/u/ann/m1/file.a  
/u/ann/m1/file.b  
/u/ann/m1/l1/file.x  
/u/ann/m1/l1/file.y
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -su=yes -preser=none
```

Restores these directories and files:

```
/u/ann/file.a  
/u/ann/file.b  
/u/ann/file.x  
/u/ann/file.y
```

Presnapshotcmd

The *presnapshotcmd* option allows you to run operating system shell commands before the LVM starts a snapshot. This allows you to quiesce an application before the LVM starts the snapshot during an online image backup. This option is only valid if the LVM is installed and configured on your system allowing you to perform an online image backup.

For an online image backup, use this option with the **backup image** command, the *include.image* option, or in the dsm.opt or dsm.sys file.

If the *presnapshotcmd* fails it is assumed that the application is not in a consistent state and the client will stop the backup and display the appropriate error message.

Supported Clients

This option is valid for the Linux x86/x86_64 clients only. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (dsm.opt) or in the client system option file (dsm.sys). You can set this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

►—PRESNapshotcmd— "*cmdstring*" —◄

Parameters

"cmdstring"

Specifies a command to process.

Use the *srvprepostsnapdisabled* option to prevent the Tivoli Storage Manager server administrator from executing operating system commands on the client system. See "Srvprepostscheddisabled" on page 354 for details on this option.

If the command string contains blanks, enclose the command string in double quotes:

```
"net stop someservice"
```

If you placed double quotes within the command string, then enclose the entire command string in single quotes:

```
'net stop "some service"'
```

Examples

Options file:

```
presnapshotcmd "any shell command or script"
```

Command line:

```
backup image -presnapshotcmd="any shell command or script"
```

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

Queryschedperiod

Authorized User

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.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

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

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define the **quiet** option, overriding the client setting. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client user options file (dsm.opt). You can set this option on the **Command Line** 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 Unix-family platforms. 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 Unix-family platforms.

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 Unix-family platforms 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.

Supported Clients

This option is valid for all UNIX clients.

Syntax

▶—REMOVEOperandlimit—▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
-removeoperandlimit
```

Replace

The *replace* option specifies whether to overwrite existing files on your workstation, or to prompt you for your selection when you restore or retrieve files.

Attention: The *replace* option does not affect recovery of directory objects. Directory objects are always recovered, even when specifying **REPlace=no**. To prevent overwriting existing directories, use the FILESONly option.

You can use this option with the following commands:

- **restore**
- **retrieve**
- **restore backupset**
- **restore group**
- **restore was**

Note: Replace prompting does not occur during a scheduled operation. If you set the *replace* option to prompt, Tivoli Storage Manager skips files without prompting during a scheduled operation.


Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client user options file (dsm.opt). You can set this option on the **Restore** category, **Action for files that already exist** section of the Preferences editor.

Syntax



Parameters

Prompt

You are prompted whether to overwrite a file that already exists on your workstation. If the existing file is read-only, you are prompted whether to overwrite it. This is the default.

All

All existing files are overwritten, including read-only files. If access to a file is denied, you are prompted to skip or overwrite the file. No action is taken on the file until there is a response to the prompt.

Yes

Any existing files are overwritten, *except* read-only files. If a file is read-only, you are prompted to overwrite the file or skip it. No action is taken on the file until there is a response to the prompt. If access to a file is denied, the file is skipped.

No Existing files are not overwritten. No prompts will be displayed.

Examples

Options file:

replace all

Command line:

-replace=no

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This 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 UNIX and Linux client with *passwordaccess=generate* specified.

A client can use more than the default number of sessions when connecting to a server that is Version 3.7 or higher. For example, *resourceutilization=10* permits up to eight sessions with the server. Multiple sessions can be used for querying the server and sending file data.

Multiple query sessions are used when you specify multiple file specifications with a backup or archive command. For example, if you enter:

```
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 204 for more information.

Regulating restore sessions

When you request a restore, the default is to use a maximum of one session, based on how many tapes the requested data is stored on, how many tape drives are available, and 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.

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.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the *dsm.sys* file *within* a server stanza. You can set this option on the **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 7
```

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 client system options file after your client node successfully contacts the server.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

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

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the 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 460 for more information.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the *dsm.sys* file *within* a server stanza.

Syntax



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*.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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 *dsmsched.log* file is created in the same directory as the *dsmerror.log* file. See “Errorlogname” on page 243 for more information on placement of the *dsmsched.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

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the *dsm.sys* file. You can set this option on the **Scheduler** category, **Schedule Log** button of the Preferences editor.

Attention: Set the *DSM_LOG* environment variable to name a directory where the log is to be placed. The directory specified must have permissions which allow write access from the account under which the client is run. The root directory is *not* a valid value for *DSM_LOG*.

Syntax

►►—SCHEDLOGName— *filespec* —————►►

Parameters

filespec

Specifies the path and file name where you want to store schedule log information when processing scheduled work. If any part of the path you specify does not exist, Tivoli Storage Manager attempts to create it.

If you specify a file name only, the file is stored in your current directory. The default is the current working directory with a file name of *dsmsched.log*. The *dsmsched.log* file *cannot* be a symbolic link.

Examples

Options file:

```
schedlogname /home/mydir/schedlog.jan
```

Command line:

```
-schedlogname=/home/mydir/schedlog.jan
```

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

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*.


Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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.

Pruned entries are copied to the *dsmsched.pru* file that is stored in the same directory as the schedule 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 *dsm.sys* file or on the **schedule** command; the client can then be contacted at either an address or a port of your choice (useful for client machines 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.


Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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.

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.


Supported Clients

This option is valid for all clients. The server cannot define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the *dsm.sys* file within a server stanza for the scheduler. You can set this option on the **Scheduler** 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** **image**
- **query** **group**
- **query** **nas**
- **query** **node**
- **query** **options**

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

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

```
scrolllines 25
```

Command line:

```
-scroll=25
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This 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** **filesystem**
- **query** **archive**
- **query** **backup**
- **query** **backupset**
- **query** **filesystem**
- **query** **image**
- **query** **group**
- **query** **nas**
- **query** **node**
- **query** **options**

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

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

```
scrollprompt yes
```

Command line:

```
-scrollp=yes
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the

beginning of the interactive session. This 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      server_a
COMMMethod      TCPip
TCPpPort        1500
TCPSeveraddress server_hostname2.domain.company.com
PASSWORDAccess  prompt
GRoups          tsm
USERS           sullivan mushock tallan
INCLExcl       /adm/tsm/backup.excl

SErvername      server_b
COMMMethod      SHAREdmem
shmpoort        1520
PASSWORDAccess  generate
MAILprog        /usr/bin/xsend root
GRoups          system tsm
INCLExcl       /adm/tsm/archive.excl
```

In your client user options file (`dsm.opt`), the *servername* option specifies which server, of those named in your `dsm.sys` file, to contact for backup-archive services. When specified in a client user options file (`dsm.opt`) or on the command line, the *servername* option overrides the default server specified in your client system options file.

Notes:

1. You cannot use the *servername* option to override the server that is specified for migration in your client system options file.
2. If the Tivoli Storage Manager server name changes or Tivoli Storage Manager clients are directed to a different Tivoli Storage Manager server, all clients will need to have a new password initialized for the new server name.

Supported Clients

This option is for all UNIX clients.

Options 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_a

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 CAD-managed schedulers, the command line client, backup-archive client Java GUI, and Web client GUI 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.
3. The Tivoli Storage Manager client API does not support this option.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

Place this option in the *dsm.sys* file *within* a server stanza. You can set this option on the **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 364). Except for CAD-managed schedulers, the command line client, backup-archive client Java GUI, and Web client GUI 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.

Shmport

Authorized User

The *shmport* option specifies the TCP/IP port address on which the Tivoli Storage Manager server listens to establish a Shared Memory connection. To use Shared Memory, TCP/IP must be installed on your workstation.

Note: The value specified for the *shmport* option in the *dsm.sys* file must match the value specified for *shmport* in the server options file.

Supported Clients

This option is valid for AIX, HP-UX, Linux, and Solaris clients *only*.

Options File

Place this option in the client system options file (*dsm.sys*).

Syntax

▶—SHMPort *port_address*—————▶

Parameters

port_address

Specifies the shared memory port address that the server is listening on to establish a Shared Memory connection. The range of values is 1 through 32767; the default is 1.

Examples

Options file:

shmport 1520

Command line:

Does not apply.

Showmembers

Use the *showmembers* option with the following commands to display all members of a group:

- **query group**
- **query was**
- **restore group**
- **restore was**

The *showmembers* option is not valid with the *inactive* option. If you want to display members of a group that are not currently active, use the *pitdate* and *pittime* options.

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

▶▶—SHOWMembers—————▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
restore group /virtfs/* -pick -showmembers
```

Skipacl

The *skipacl* option allows you to back up or restore ACL data. When set to *yes* (default is *no*), the Tivoli Storage Manager client 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.


Supported Clients

This option is valid for all AIX clients.

Options 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.

Supported Clients

This option is valid for all AIX clients.

Options File

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

Syntax

►—SKIPACLupdatecheck

No
Yes

►►

Parameters

No If you specify *No*, the Tivoli Storage Manager client performs checksum and size comparisons of the ACL data, before and after backup and during incremental processing. This is the default.

Yes If you specify *Yes*, the Tivoli Storage Manager client does not perform checksum and size comparisons of the ACL data.

Examples

Options file:
skipaclup yes

Snapshotcachesize

Authorized User

Use the *snapshotcachesize* option to specify an appropriate snapshot size so that all old data blocks can be stored. A snapshot size of 100 percent will ensure a valid snapshot.

For online image backups, use the *snapshotcachesize* option with the **backup image** command, the *include.image* option, or in your `dsm.sys` file.

Supported Clients

This option is valid for Windows 32-bit platformsLinux x86/x86_64 clients *only*. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the `dsm.sys` file. You can set this option on the **Image-Snapshot** category of the Preferences editor.

Syntax

▶—SNAPSHOTCACHESize— *snapshotcachesize* —▶

Parameters

snapshotcachesize

Specifies an appropriate snapshot size so that all old data blocks can be stored during a snapshot image backup. The value is the percent of the used space that is expected to change on the volume on which the snapshot is taken. On a volume with high filesystem write activity, this value can be increased to remove the condition where the snapshot cache fills up during the operation. The snapshot cache location is specified with the *snapshotcachelocation* option. The range of values is one through 100 percent; the default is 100 percent. The minimum cache file size is 100 megabytes. If 100 megabytes of free space is not available on the *snapshotcachelocation* volume, then the cache size will be reduced to fit on the volume.

Examples

Options file:

```
snapshotcachesize 10
```

Command line:

```
-snapshotcachesize=10
```

Snapshotroot

Use the *snapshotroot* option with the **incremental**, **selective**, or **archive** commands in conjunction with a third-party application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server.

The *snapshotroot* option can be used to back up NFS mounted file systems. Both the back up specification (source) and the *snapshotroot* value can be a NFS mounted file specification. For example, the *snapshotroot* option can be used to back up a NFS file system hosted on a network-attached storage (NAS) that supports snapshot.

In the following example, /philo/test945 is NFS mounted from a NAS filer and /philo/test945.snapshot/backupsnap represents the snapshot created at the NAS filer.

```
dsmc i /philo/test945 -snapshotroot=philo/test945/.snapshot/backupsnap
```

You can also specify a directory with the *snapshotroot* option when backing up each fileset as a separate file space.

The *snapshotroot* option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot.

For example, consider an application that takes a snapshot of the /usr file system and mounts it as /snapshot/day1. If you back up this data using the following command:

```
dsmc incremental /snapshot/day1
```

a unique file space called /snapshot/day1 is created on the server. However, you might want to associate the snapshot data with the data already processed for the /usr file system. Using the *snapshotroot* option, you can associate the data with the file space corresponding to the /usr file system on the Tivoli Storage Manager server:

```
dsmc incremental /usr -snapshotroot=/snapshot/day1
```

On a subsequent day, you can back up a snapshot that was written to an alternate location, but managed under the same file space on the server:

```
dsmc incremental /usr -snapshotroot=/snapshot/day2
```

You can perform incremental backups, selective backups, or archives of a single directory, directory structure, or single file using the *snapshotroot* option. In all instances, the *snapshotroot* option must identify the root of the logical volume that was created by the snapshot. For example:

```
dsmc incremental /usr/dir1/* -subdir=yes
  -snapshotroot=/snapshot/day1
dsmc selective /usr/dir1/sub1/file.txt
  -snapshotroot=/snapshot/day1
dsmc archive /usr/dir1/sub1/*.txt
  -snapshotroot=/snapshot/day1
```

If you want to include or exclude specific file specifications, the include and exclude statements should contain the name of the file system that was the *source* of the snapshot (the /usr file system), and *not* the name of the target of the snapshot (/snapshot/day1). This allows you to preserve a set of include and

exclude statements regardless of the name of the logical volume to which the snapshot is written. Examples of include and exclude statements are:

```
include /usr/dir1/*.txt 1yrmgmtclass
exclude /usr/mydocs/*.txt
```

The following include-exclude statements are not valid because they contain the name of the snapshot:

```
include /snapshot/day1/dir1/*.txt 1yrmgmtclass
exclude /snapshot/day1/mydocs/*.txt
```

You must use the *snapshotroot* option in conjunction with a single file specification for a incremental, selective, or archive operation. You cannot specify multiple file specifications or no file specifications. For example, these commands are valid:

```
dsmc incremental /usr -snapshotroot=/snapshot/day1
dsmc incremental /usr/dir1/* -snapshotroot=/snapshot/day1
```

The following command is invalid because it contains two file specifications:

```
dsmc incremental /usr/dir1/* /home/dir2/*
-snapshotroot=/snapshot/day1
```

The following command is invalid because it contains no file specification:

```
dsmc incremental -snapshotroot=/snapshot/day1
```

Processing SAN File System FlashCopy images

SAN File System provides snapshot support called FlashCopy. FlashCopy images can only be created using the SAN File System administrative interface. FlashCopy images are created on a per fileset (container) basis and are accessible in the .flashcopy directory of each fileset's root directory. For example, there are two filesets, root fileset (the default one) and another fileset called tivoli1. An administrator creates one FlashCopy image for each fileset, called RootImage-1 and tivoli1Image-1 respectively. A SAN File System is mounted in the /tank directory on AIX. In this case, the FlashCopy images are accessible in the following directories:

```
/tank/TIVOLI/.flashcopy/RootImage-1/
/tank/TIVOLI/tivoli1/.flashcopy/tivoli1Image-1/
```

You can specify these directories with the *snapshotroot* option when backing up each fileset as a separate file space as follows:

1. Define a virtual mount point for the fileset, using the *virtualmountpoint* option in your dsm.sys file:

```
virtualmountpoint /tank/TIVOLI
virtualmountpoint /tank/TIVOLI/tivoli1
```

2. Define virtual mount points for the FlashCopy image locations:

```
virtualmountpoint /tank/TIVOLI/.flashcopy/RootImage-1
virtualmountpoint /tank/TIVOLI/tivoli1/.flashcopy
/tivoli1Image-1
```

3. Exclude the .flashcopy directory itself, so the data is not backed up twice:

```
exclude.dir /.../.flashcopy
```

4. On the command line, enter the following backup command:

```
dsmc incremental /tank/TIVOLI -snapshotroot=/tank/TIVOLI
/.flashcopy/RootImage-1
dsmc incremental /tank/TIVOLI/tivoli1 -snapshotroot=/tank
/TIVOLI/tivoli1/
.flashcopy/tivoli1Image-1
```


Notes:

1. Ensure that the `snapshotroot` references a snapshot of the correct volume. Ensure that `snapshotroot` refers to the root of the snapshot. If these rules are not followed, unintended results such as files expiring incorrectly can result.
2. You cannot use the `snapshotroot` option in conjunction with the `filelist` option.
3. You can use the `snapshotroot` option in conjunction with the `preschedulecmd` and `postschedulecmd` options, or in a automated script that you execute with the Tivoli Storage Manager client scheduler.

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

▶▶—SNAPSHOTRoot =- *snapshot_volume_name*————▶▶

Parameters

snapshot_volume_name

Specifies the root of the logical volume created by the third-party snapshot application.

Examples

Command line:

```
dsmc incremental /usr -snapshotroot=/snapshot/day1
```

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 machine, when performing scheduled operations. The *srvprepostscheddisabled* option can be used in conjunction with the *schedcmddisabled* and *srvprepostsnapdisabled* options to disable the execution of any unwanted Operating System command by a Tivoli Storage Manager administrator on a Tivoli Storage Manager client node.

Supported Clients

This option is valid for all Tivoli Storage Manager clients that use the Tivoli Storage Manager client scheduler. The server cannot define this option.

Options File

Place this option in the `dsm.sys` file within a server stanza for the scheduler. You can set this option on the **Scheduler** tab of the Preferences editor.

Syntax



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 machine, 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 machine, 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* and *srvprepostsnapdisabled* options.

Examples

Options file:

```
srvprepostscheddisabled yes
```

Command line:

```
Does not apply.
```

Srvprepostsnapdisabled

Authorized User

The *srvprepostsnapdisabled* option specifies whether to prevent the pre-snapshot and post-snapshot commands specified by the Tivoli Storage Manager administrator from executing on the client machine, when performing scheduled image snapshot backup operations. The *srvprepostsnapdisabled* option can be used in conjunction with the *schedcmddisabled* and *srvprepostscheddisabled* options to disable the execution of any unwanted Operating System command by a Tivoli Storage Manager administrator on a Tivoli Storage Manager client node.

Supported Clients

This option is valid for Linux clients that support the image snapshot backup command. The server cannot define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the `dsm.sys` file within a server stanza for the scheduler. You can set this option on the **Image-Snapshot** category of the Preferences editor.

Syntax



Parameters

- No* Specifies that Tivoli Storage Manager allows pre-snapshot and post-snapshot commands defined by the Tivoli Storage Manager administrator to execute on the client machine, when performing scheduled image snapshot backup operations. If a pre-snapshot or a post-snapshot command is defined by both the client and the Tivoli Storage 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 allows pre-snapshot and post-snapshot commands defined by the Tivoli Storage Manager administrator to execute on the client machine, when performing scheduled image snapshot backup operations. If a pre-snapshot or a post-snapshot command is defined by both the client and the Tivoli Storage 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* and *srvprepostscheddisabled* options.

Examples

Options file:
`srvprepostsnapdisabled 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.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the 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 318.

If a subdirectory is a mounted file system, it will not be processed even if you specify *subdir=yes*.

Examples

Options file:

```
subdir no
```

Command line:

To restore the structure:

```
/path2/dir1  
/path2/dir1/file1  
/path2/dir1/dir2  
/path2/dir1/dir2/file1
```

enter any of the following commands:

```
dsmc rest "/path/dir1/*" /path2/ -su=yes  
dsmc rest "/path/dir1/file*" /path2/ -su=yes  
dsmc rest "/path/dir1/file1*" /path2/ -su=yes
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This 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 “No query restore” on page 124 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.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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*.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the *dsm.sys* file *within* a server stanza. You can set this option on the **Communication** category, **Admin Port** field of the Preferences editor. See the *IBM Tivoli Storage Manger 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 *tcpport* 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.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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.

Tpcadaddress

Authorized User

The *tpccadaddress* option specifies a TCP/IP address for dsmscad. 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 206).

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza.

Syntax

▶▶—TCPCADAddress— *cad_address*—————▶▶

Parameters

cad_address

Specifies a TCP/IP Internet domain name or a dotted address.

Examples

Options file:

```
tcpcada dsmclnt.sanjose.ibm.com
```

Command line:

```
-tcpcadaddress=128.33.10.249
```

or

```
-tcpcadaddress=khoyt.mycompany.mydomain.com
```

This option is valid only on the initial command line of the dsmscad 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 *shedmode* option.

If *sessioninitiation* is set to *serveronly*, the value for the *tcpclientaddress* client option should be the same as the value for the *HLAddress* server option.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the *dsm.sys* file *within* a server stanza. 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 dot address.

Examples

Options file:

```
tcpclienta dsmclnt.sanjose.ibm.com
```

Command line:

```
-tcpclientaddress=128.33.10.249  
or  
-tcpclientaddress=khoyt.mycompany.mydomain.com
```

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.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the *dsm.sys* file *within* a server stanza. 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.

Supported Clients

This option is valid for all UNIX and Linux clients, except UNIX System Services client for z/OS.

Options File

Place this option in the client system options file (*dsm.sys*) *within* a server stanza. You can set this option on the **Communication** category of the Preferences editor.

Syntax



Parameters

- No* Specifies that the server does not allow successive small packets to be sent immediately over the network. Setting this option to *no* can degrade performance.
- Yes* Specifies that the server or client allows successive small packets to be sent immediately over the network. The default is *yes*.

Examples

Options file:
tcpnodelay yes

Command line:
Does not apply.

Tcpport

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.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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 server address from your administrator.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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 dot address.

Examples

Options file:

tcps dsmchost.endicott.ibm.com

Command line:

Does not apply.

|

Tcpwindowsize

Authorized User

Use the *tcpwindowsize* option to specify, in kilobytes, the size you want to use for the TCP/IP sliding window for your client node. The sending host cannot send more data until it receives an acknowledgment and a TCP receive window update. Each TCP packet contains the advertised TCP receive window on the connection. A larger window allows the sender to continue sending data and can improve communication performance.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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

▶—TCPWindowsize— *window_size* —▶

Parameters

window_size

Specifies the size, in kilobytes, to use for your client node TCP/IP sliding window. The range of values is 0 through 2048. A value of 0 allows Tivoli Storage Manager to use the operating system default TCP window size. Values from 1 to 2048 indicate that the window size is in the range of 1KB to 2MB. 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 2048, the TCP window size defaults to 2048.

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:

```
tcpwindowsize 63
```

Command line:

```
-tcpw=63
```

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

Timeformat

The *timeformat* option specifies the format in which you want to display system time.

The AIX, Solaris, and HP-UX clients support locales other than English that describe every user interface that varies with location or language. See Table 18 on page 47 for supported locales.

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 asr**
- **query backup**
- **query filespace**
- **query image**
- **query nas**
- **restore**
- **restore image**
- **restore nas**
- **retrieve**
- **set event**

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

Place this option in the client user 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.

0 Use the locale-defined time format.

For AIX, HP-UX, and Solaris: This is the default if the locale-specified format consists of digits, separator characters, and, if applicable, the AM or PM string.

1 23:00:00

For AIX, HP-UX, and Solaris: This is the default if the locale-specified format does not consist of digits, separator characters, and, if applicable, the AM or PM string.

2 23,00,00

3 23.00.00

4 12:00:00 A/P

For AIX, HP-UX, and Solaris: To set a particular time format, edit the source file for your locale and modify the `t_fmt` line to support your needs. Whatever time format you select applies both to output and to input.

`"%H:%M:%S"`

Displays time in the form `hh:mm:ss` with `hh` ranging from 0 through 23.

`"%H,%M,%S"`

Displays time in the form `hh,mm,ss` with `hh` ranging from 0 through 23.

`"%I,%M,0p"`

Displays time in the form `hh,mm,ssA/P` with `hh` ranging from 1 through 12 and `A/P` is the local abbreviation for ante-meridian (AM in English) or post-meridian (PM in English).

Examples

Options file:

`timeformat 4`

Command line:

`-time=3`

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This will be the value from the `dsm.opt` file unless overridden by the initial command line or by an option forced by the server.

Toc

Use the *toc* option with the **backup nas** command or the *include.fs.nas* option to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup. You should consider the following when deciding whether you want to save TOC information:

- If you save TOC information, you can use the **query toc** server command to determine the contents of a file system backup in conjunction with the **restore node** server command to restore individual files or directory trees.
- You can also use the Tivoli Storage Manager Web client to examine the entire file system tree and select files and directories to restore.
- Creation of a TOC requires that you define the TOCDESTINATION attribute in the backup copy group for the management class to which this backup image is bound. Note that TOC creation requires additional processing, network resources, storage pool space, and possibly a mount point during the backup operation.
- If you do not save TOC information, you can still restore individual files or directory trees using the **restore node** server command, provided that you know the fully qualified name of each file or directory and the image in which that object was backed up.


Supported Clients

This option is valid for AIX and Solaris clients *only*. The Tivoli Storage Manager client API does not support this option.

Options File

Place the *include.fs.nas* statement containing the *toc* value in the dsm.sys file.

Syntax



Parameters

Yes Specifies that Tivoli Storage Manager saves TOC information during a NAS file system image backup. However, the backup will fail if an error occurs during creation of the TOC.

No Specifies that Tivoli Storage Manager does not save TOC information during a NAS file system image backup.

Preferred

Specifies that Tivoli Storage Manager saves TOC information during a NAS file system image backup. The backup does not fail if an error occurs during creation of the TOC. This is the default.

Note: If the *mode* option is set to *differential* and you set the *toc* option to *preferred* or *yes*, but the last full image does not have a TOC, Tivoli Storage Manager performs a full image backup and creates a TOC.

Examples

Options file:

```
include.fs.nas netappsj/vol/vol0 homemgmtclass toc=yes
```

Command line:

```
backup nas -nasnodename=netappsj /vol/vol0 -toc=yes
```

Todate

Use the *todate* option with the *totime* option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation. 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**
- **restore group**
- **retrieve**
- **restore was**

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—TODate =- *date*—————▶▶

Parameters

date

Specifies an ending date. Enter the date in the format you selected with the *dateformat* option.

When you include *dateformat* with a command, it must precede the *fromdate*, *pitdate*, and *todate* options.

Examples

Command line:

```
dsmc restore "/home/user1/*" -todate=12/11/2003
```

Totime

Use the *totime* option with the *todate* option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation. 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**
- **restore group**
- **restore was**
- **retrieve**

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—Totime =— *time* —————►►

Parameters

time

Specifies an ending time. If you do not specify a time, the time defaults to 23:59:59. Specify the time in the format you selected with the *timeformat* option.

When you include the *timeformat* option in a command, it must precede the *fromtime*, *pittime*, and *totime* options.

Examples

Command line:

```
dsmc restore "/home/user1/*" -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.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options 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.

Supported Clients

This option is valid for AIX and Solaris clients *only*. The Tivoli Storage Manager client API does not support this option.

Syntax



Parameters

- any*
Specifies all nodes registered at the server. This is the default.
- nas*
Specifies all NAS nodes registered at the server.
- server*
Specifies client nodes that are other Tivoli Storage Manager servers.
- client*
Specifies client nodes that are backup-archive clients.

Examples

Command line:
query node -type=nas

Users

Authorized User

The *users* option authorizes specific users on your workstation to request services from a server. You can use this option more than once to specify a large number of user IDs. If you do not specify group names with the *groups* option, or user IDs with the *users* option, all users can request Tivoli Storage Manager services. If you use the *groups* option, the *users* option, or both, only users included in one of the specified groups, or included in the list of users, can request Tivoli Storage Manager services.

Define your root user name only with the *users* option to exclude all other users from accessing the server.

See “How Tivoli Storage Manager handles long user and group names” on page 86 for information about the character limits for group names.

Supported Clients

This option is valid for all UNIX clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the *dsm.sys* file *within* a server stanza.

Syntax



Parameters

username

Names a user that you want to authorize to request Tivoli Storage Manager services.

Examples

Options file:

```
users carol larry davecd kathyba michelle  
users amyb tkaspar kbsmith egray srjames
```

Command line:

Does not apply.

V2archive

Use the *v2archive* option with the **archive** command to archive only files to the server. Tivoli Storage Manager will not process directories that exist in the path of the source file specification.

This option differs from the *filesonly* option in that the *filesonly* option archives the directories that exist in the path of the source file specification.

The *v2archive* and *dirsonly* options are mutually exclusive and an error message is displayed if you use both options in the same **archive** command.

If you use this option, you might want to consider the following:

- You might experience performance problems when retrieving large amounts of data archived with this option.
- You might want to use this option only if you are concerned about expiration performance on a server that already contains extremely large amounts of archived data.
- If there are multiple files with the same name for the *v2archive* option, the files are archived multiple times, with their directory structure. The *v2archive* option archives only the files.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶—V2archive—▶

Parameters

There are no parameters for this option.

Examples

This command:

```
dsmc archive "/home/relx/dir1/*" -v2archive -su=y.
```

Archives these files:

```
/home/relx/dir1/file1  
/home/relx/dir1/file2  
/home/relx/dir1/file3  
/home/relx/dir1/dir2/file4  
/home/relx/dir1/dir2/file5
```

Note: Tivoli Storage Manager does not archive `/home/relx/dir1` and `/home/relx/dir1/dir2`.

Verbose

The *verbose* option specifies that you want to display detailed processing information on your screen. This is the default. When you run the **incremental**, **selective**, or **archive** commands, information is displayed about each file that is backed up. Use the *quiet* option if you do not want to display this information.

The following behavior applies when using the *verbose* and *quiet* options:

- If the server specifies either the *quiet* or *verbose* option in the server client option set, the server settings override the client values, even if **force** is set to *no* on the server.
- If you specify *quiet* in your *dsm.opt* file, and you specify *-verbose* on the command line, *-verbose* prevails.
- If you specify both *-quiet* and *-verbose* on the same command, the last option encountered during options processing prevails. If you specify *-quiet -verbose*, *-verbose* prevails. If you specify *-verbose -quiet*, *-quiet* prevails.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the 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.

Verifyimage

Use the *verifyimage* option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log.

Supported Clients

This option is valid for AIX, HP-UX, all Linux clients, and Solaris *only*. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶—VERIFYImage—————▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore image /usr -verifyimage
```

Virtualfsname

Use the *virtualfsname* option with the **backup group** command to specify the name of the virtual file space for the group on which you want to perform the operation.

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

▶▶—VIRTUALFSname =- *fsname*—————▶▶

Parameters

fsname

Specifies the name of the container for the group on which you want to perform the operation.

Examples

Command line:

```
backup group -filelist=/home/dir1/filelist1 -groupname=group1  
-virtualfsname=virtfs -mode=full
```

Virtualmountpoint

Authorized User

The *virtualmountpoint* option defines a virtual mount point for a file system if you want to consider files for backup that begin with a specific directory within that file system. Using the *virtualmountpoint* option to identify a directory within a file system provides a direct path to the files you want to back up, saving processing time. It is more efficient to define a virtual mount point within a file system than it is to define that file system using the *domain* option, and then to use the *exclude* option in your include-exclude options list to exclude the files that you do not want to back up.

Use the *virtualmountpoint* option to define virtual mount points for multiple file systems, for local and remote file systems, and to define more than one virtual mount point within the same file system. Virtual mount points cannot be used in a file system handled by automounter.

Note: If the directory that you want to specify as a virtual mount point is a symbolic link, set the *followsymbolic* option to *Yes*. If that option is set to *no* (the default), you are not permitted to use a symbolic link as a virtual mount point. Also, if you back up a file system, then add a *virtualmountpoint*, and then do another incremental on the file system, the files and directories in the *virtualmountpoint* directory will be expired, because they are logically contained within the *virtualmountpoint* directory and not the file system.

After you define a virtual mount point, you can specify the path and directory name with the *domain* option in either the default client options file or on the **incremental** command to include it for incremental backup services. When you perform a backup or archive using the *virtualmountpoint* option, the **query filespace** command will list the virtual mount point in its response along with other file systems. Generally, directories that you define as virtual mount points are treated as actual file systems and require that the *virtualmountpoint* option is specified in the *dsm.sys* file to restore or retrieve the data.

Note: When you specify a *virtualmountpoint* option, the path that it specifies is added to the default backup domain (domain ALL-LOCAL). The *virtualmountpoint* path is always considered a local "mount point" regardless of the real file system type it points to.

Supported Clients

This option is valid for all UNIX clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the *dsm.sys* file.

Syntax



Parameters

directory

Specifies the path and directory name for the directory you want to use as the virtual mount point for a file system. You cannot use wildcard characters in either the path or directory names.

Define only one virtual mount point with each *virtualmountpoint* option that you include in your client system options file. Use the *virtualmountpoint* option as many times as necessary to define all of the virtual mount points that you want to use.

Examples

Options file:

```
virtualmountpoint /afs/xyzcorp.com/home/ellen  
virtualmountpoint /afs/xyzcorp.com/home/ellen/test/data
```

Command line:

Does not apply.

Virtualnodename

The *virtualnodename* option specifies the node name of your workstation when you want to restore or retrieve files to a different workstation.

When you use the *virtualnodename* option in your client user options file, or with a command:

- You must specify the name you specified with the *nodename* option in your client system options file (dsm.sys). This name should be different from the name returned by the **hostname** command on your workstation.
- Tivoli Storage Manager prompts for the password assigned to the node you specify, if a password is required. If you enter the correct password, you have access to all backups and archives that originated from the specified node.

When connecting to a server, the client must identify itself to the server. This login identification is determined in the following ways:

- If the *nodename* and *virtualnodename* options are not specified, or a virtual node name is not specified on the command line, the default login ID is the name returned by the **hostname** command.
- If the *nodename* option is specified, the name specified with the *nodename* option overrides the name returned by the **hostname** command.
- If the *virtualnodename* option is specified, or a virtual node name is specified on a command line, it cannot be the same name as the name returned by the **hostname** command.

When the virtual node name is accepted by the server, a password is required (assuming authentication is on), even if the *passwordaccess* option is *generate*. Once a connection to the server is established, then access is permitted to any file backed up using this login ID.

| For the Java GUI, the *virtualnodename* option will only work if you are root user.
| If you need to use the *virtualnodename* option with a non-root user account,
| please use the command line client (or native GUI, for platforms that still have
| one).

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

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

Syntax

▶▶—VIRTUALNodename— *nodename* —————▶▶

Parameters

nodename

Specifies a 1- to 64-character name that identifies the node for which you want to request Tivoli Storage Manager services. There is no default.

Examples

Options file:

```
virtualnodename cougar
```

Command line:

```
-virtualn=banshee
```

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

Wasexphome

To back up the WebSphere Application Server-Express, use the *wasexphome* option to specify the fully qualified installation path of the WebSphere Application Server-Express. Ensure that the *washome* option is *not* set.

Supported Clients

This option is valid for AIX, Solaris, and Linux x86/x86_64 clients.

Options File

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

Syntax

▶▶—WASExpHome— —*pathvalue*—————▶▶

Parameters

pathvalue

Specifies the fully qualified path of the home directory of the WebSphere Application Server - Express. If the path contains spaces, enclose the path in double quotes.

Examples

Options file:

```
wasexphome /usr/mywas/Express
```

Washome

Use the *washome* option in your client user options file (dsm.opt) to specify an override base install path for the Application Server. You can use this option if your entire WebSphere installation is corrupted and must be reinstalled, or you are installing WAS on a new machine. See “WAS instance restore procedures” on page 481 for more information.

If you do not specify a value for this option, Tivoli Storage Manager uses the default installation path. If there are multiple installations of the Application Server on the same machine, use a different options file for each installation with the proper path to the installation directory.

Supported Clients

This option is valid for AIX, Solaris, and Linux x86/x86_64 clients.

Options File

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

Syntax

►►—WASHOME— *pathvalue*—————►►

Parameters

pathvalue

Specifies the fully qualified path of the home directory of the WebSphere Application Server installation. This value is the path of the directory where the configuration information and properties reside. If the path contains spaces, enclose the path in double quotes.

Examples

Options file:

```
washome /usr/mywas/appserver
```

Wasndhome

Use the *wasndhome* option in your client user options file (dsm.opt) to specify an override base install path for the Network Deployment Manager. You can use this option if your entire WebSphere installation is corrupted and must be reinstalled, or you are installing WAS on a new machine. See “WAS instance restore procedures” on page 481 for more information.

If you do not specify a value for this option, Tivoli Storage Manager uses the default installation path. If there are multiple installations of the Network Deployment Manager on the same machine, use a different options file for each installation with the proper path to the installation directory.

Supported Clients

This option is valid for AIX, Solaris, and Linux x86/x86_64 clients clients.

Options File

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

Syntax

▶▶—WASNDHome— —*pathvalue*—————▶▶

Parameters

pathvalue

Specifies the fully qualified path of the home directory of the Network Deployment Manager installation. This value is the path of the directory where the configuration information and properties reside. If the path contains spaces, enclose the path in double quotes.

Examples

Options file:

```
wasndhome /usr/mywas/DeploymentManager
```

Wasnode

Use the *wasnode* option with the **set waspassword** command to specify the WAS node name when performing the operation on the WAS Network Deployment Manager or Application Server.

Supported Clients

This option is valid for AIX, Solaris, and Linux x86/x86_64 clients.

Syntax

►►—WASNode =— *nodename*—————►►

Parameters

nodename

Specifies the WAS node name when performing the operation on the WAS Network Deployment Manager or Application Server.

Examples

Command line:

```
backup was -wasnode=ednode -wastype=app
```

Wastype

Use the *wastype* option with the **backup was**, **query was**, or **restore was**, or **set waspassword** commands to perform the operation on the WebSphere Application Server (WAS) Network Deployment Manager (contains setup, application files, and configuration information), the Application Server, or both.

Supported Clients

This option is valid for AIX, Solaris, and Linux x86/x86_64 clients clients.

Syntax



Parameters

ND

Specifies that you want to perform the operation on the Network Deployment Manager (ND) associated with the node name that you specify. This is the default for the **backup was** and **restore was** commands.

APP

Specifies that you want to perform the operation on the Application Server (APP) associated with the node name that you specify.

ANY

Specifies that you want to query all backups of Network Deployment Manager and Application Server associated with the node name that you specify, including instances of ND and APP. This parameter is valid for the **query was** command only, and is the default.

LOCAL

Specifies that you want to query all the of the Application Servers, Network Deployment Manager, and their instances on your local machine. This parameter displays the instance name, hostname, soap port information, installed path, the type of the WAS (ND or APP), and whether security is enabled. This parameter is valid for the **query was** command only.

Examples

Command line:

```
dsmc query was -wastype=local
```

Wasuser

If WAS security is enabled, use the *wasuser* option with the **set waspassword** command to set the WAS user name for each installation of WAS on your machine.

Supported Clients

This option is valid for AIX, Solaris, and Linux x86/x86_64 clients clients.

Syntax

▶▶—wasuser =— *username*—————▶▶

Parameters

username

Specifies the WAS user name when performing the operation on the WAS Network Deployment Manager or Application Server.

Examples

Command line:

```
dsmc set waspassword -wasnode=wasnode -wastype=app -wasuser=ed
```

Webports

The *webports* option enables the use of the Web client outside a firewall by specifying the TCP/IP port number used by the Tivoli Storage Manager Client Acceptor daemon and Web Client Agent service for communications with the Web client.

Values for both the Client Acceptor daemon and the Web Client Agent service are required.

If you do not specify this option, the default value, zero (0), is used for both ports. This causes TCP/IP to randomly assign a free port number for the Client Acceptor daemon and the Web Client Agent service.

Note: The Tivoli Storage Manager client API does not support this option.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

Place this option in the *dsm.sys* file *within* a server stanza. You can set this option on the **Web Client** category, **WEB Ports** fields of the Preferences editor.

Syntax

▶—WEBPorts— *cadport*— *agentport*—▶

Parameters

cadport

Specifies the *required* Tivoli Storage Manager Client Acceptor daemon port number. The range of values is 1000 through 32767. If a value is not specified, the default, zero (0), causes TCP/IP to randomly assign a free port number.

agentport

Specifies the *required* Tivoli Storage Manager Web client agent service port number. The range of values is 1000 through 32767. If a value is not specified, the default, zero (0), causes TCP/IP to randomly assign a free port number.

Examples

Options file:

```
webports 2123 2124
```

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 64 shows a list of tasks related to entering commands.

Table 64. Entering commands

Task	Page
Starting and ending a client command session	395
Entering client commands	397
Remembering previous commands	399
Using wildcard characters	400

Table 65 provides an alphabetical list of the commands, a brief description, and where to locate more information.

Table 65. Commands

Command	Description	Page
archive	Archives files from a workstation to Tivoli Storage Manager storage.	402
backup group	Creates and backs up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.	404
backup image	Creates an image backup of one or more file systems or logical volumes that you specify.	406
backup nas	Creates an image backup of one or more file systems belonging to a Network Attached Storage (NAS) file server.	410
backup was	Backs up the WebSphere Application Server (WAS) Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server to the Tivoli Storage Manager server.	413
cancel process	Displays a list of current NAS (if NDMP support is enabled) image backup and restore processes for which the administrative user has authority.	415
cancel restore	Displays a list of restartable restore sessions from which you can select one to cancel.	416
delete access	Deletes authorization rules for files or images that are stored on the server.	417
delete archive	Deletes archived files from Tivoli Storage Manager server storage.	418
delete backup	Deletes active and inactive backup files from Tivoli Storage Manager server storage.	420
delete filespace	Deletes file spaces in Tivoli Storage Manager server storage.	423

Table 65. Commands (continued)

Command	Description	Page
delete group	Deletes a group backup on the Tivoli Storage Manager server.	425
expire	Inactivates backup objects that you specify in the file specification or with the <i>filelist</i> option.	427
help	Displays a Table of Contents of help topics for the command line client..	429
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.	430
loop	Starts an interactive command session.	435
macro	Executes commands within a macro file that you specify.	437
monitor process	Displays a list of current NAS image backup and restore processes from which you can select one to cancel	438
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.	439
query access	Displays a list of current authorization rules.	441
query archive	Displays a list of archived files.	442
query backup	Displays a list of backup versions.	444
query backupset	Queries a backup set from a local file, tape device, or the Tivoli Storage Manager server.	447
query filespace	Displays a list of file spaces in Tivoli Storage Manager storage. You can also specify a single file space name to query.	449
query group	Displays information about group backups and their members.	451
query image	Displays information about image backups.	453
query inlexcl	Displays a list of include-exclude statements in the order in which they are processed during backup and archive operations.	455
query mgmtclass	Displays information about available management classes.	456
query node	Displays all the nodes for which an administrative user ID has authority to perform operations.	457
query options	Displays all or part of your options and their current settings.	458
query restore	Displays a list of your restartable restore sessions in the server database.	459
query schedule	Displays information about scheduled events for your node.	460
query session	Displays information about your session, including the current node name, when the session was established, server information, and server connection information.	461

Table 65. Commands (continued)

Command	Description	Page
query systeminfo	Gathers Tivoli Storage Manager system information and outputs this information to a file or the console.	462
query was	Displays backups of the WebSphere Application Server (WAS) Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server that match the node name and type of the WAS group backup that you specify.	464
restart restore	Displays a list of restartable restore sessions from which you can one to restart.	465
restore	Restores copies of backup versions of your files from a Tivoli Storage Manager server.	466
restore backupset	Restores a backup set from the Tivoli Storage Manager server or a local file. You can also restore a backup from a tape device.	470
restore group	Restores specific members or all members of a group backup.	474
restore image	Restores a file system or raw volume image backup.	476
restore nas	Restores the image of a file system belonging to a Network Attached Storage (NAS) file server.	479
restore was	Restores the WebSphere Application Server (WAS) Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server from the Tivoli Storage Manager server.	481
retrieve	Retrieves copies of archived files from the Tivoli Storage Manager server.	484
schedule	Starts the client scheduler on the workstation.	486
selective	Backs up selected files.	488
set access	Authorizes another user to access your backup versions, archived copies, or image backupsbackup or archive data.	491
set event	Allows you to specify the circumstances for when archived data is deleted.	493
set password	Changes the Tivoli Storage Manager password for your workstation.	495
set waspassword	When WAS security is enabled, this command allows you to set your WebSphere node name, user name, and password for the WAS Network Deployment Manager or the Application Server.	496

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 shell 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 169 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**. Tivoli Storage Manager processes the command and returns to the shell command prompt. 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.

Notes:

1. In loop mode, following a restore operation directly from tape, the mount point is not released in case additional restore requests are made to that volume. If you request a backup operation in the same session and that mount point is the only one available, the backup operation will stop with the following message:

```
Waiting for mount of offline media
```

In this case, the mount point is not released until one of the following conditions is met:

- The device class MOUNTRETENTION limit is satisfied.
 - The client idletimeout period is satisfied.
 - The dsmc loop session is closed after the restore operation completes, allowing you to start a subsequent loop mode session to perform the backup operation.
2. In interactive mode, you cannot enter a file specification that contains national language characters. If a command contains national characters, process the command in batch mode by preceding the command with the executable program name, **dsmc**.

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 194.
- **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 194.

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*:

- Do not use national language characters within the Tivoli Storage Manager command line client interactive mode. If a file specification contains national characters, process the command in batch mode, preceding it with the executable program name **dsmc**.
- 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 */home/me*, then the *destinationfilespec* would be */home/me/mydir* in the following command:

```
dsmc restore "/fs/dir1/*" mydir/
```

- The only command that accepts a simple file space name is the **incremental** command. The following example is valid:

```
dsmc i /fs
```

The following example is *not* valid:

```
dsmc sel /fs
```

- When entering the *sourcefilespec*, if the directory name ends with */*, then */** is implied.

When entering a *destinationfilespec*, if the name ends with */*, then it is considered a directory, otherwise it is considered a file.

The following example illustrates these two rules. Even though `mydir` and `yourdir` are directories, the command will fail because `/*` is implied after `mydir`, and `yourdir` is considered a file:

```
restore /home/mydir/ /away/yourdir
```

The following example illustrates the second rule. Even though `mydir` and `yourdir` are directories, the command will fail because `mydir` and `yourdir` are considered files:

```
restore /home/mydir /away/yourdir
```

- 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 `/home` directory, enter:

```
set access backup /home/* * *
set access backup /home/*/* * *
```

Do not use wildcards for the directory path name, for example:

```
/home/j*asler/file1.c
```

- The maximum number of characters for a file name is 256. The maximum number of characters for a path name is 1024 characters.
- 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 250 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 250 for more information.
2. You can also use the *removeoperandlimit* option to specify that Tivoli Storage Manager removes the 20-operand limit for UNIX-family platforms. 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 324

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 234.

Pressing the Up arrow key displays the previous command in memory. Pressing the Down arrow key displays the next command in memory. Table 66 lists other functions you can perform when you recall commands.

Note: The following keys or key combinations are not guaranteed to work in all terminal applications. Depending on your application's specifics, they may not work at all or map to wrong commands.

Table 66. 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 beginning of the command.	Home
Move to the end of the command.	End
Move to the left.	Left arrow
Move to the right.	Right arrow
Move five spaces to the left.	Tab left
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.	Delete
Delete a character to the left of the cursor.	Backspace
Insert a character.	Toggle the Insert key
Erase to the end of the line.	Ctrl-delete or Ctrl-D
Finish or execute the command.	Enter
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 67 shows examples of each wildcard.

Table 67. Wildcard characters

Pattern	Matches	Does not match
<i>Asterisk (*)</i>		
ab*	ab, abb, abxxx	a, b, aa, bb
ab*rs	abrs, abtrs, abrsrs	ars, aabrs, abrss

Table 67. Wildcard characters (continued)

Pattern	Matches	Does not match
ab*ef*rs	abefrs, abefghrs	abefr, abers
abcd.*	abcd.c, abcd.txt	abcd, abcdc, abcdtxt
<i>Question Mark (?)</i>		
ab?	abc	ab, abab, abzzz
ab?rs	abrs	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 double quotes. For example:

```
dsmc selective "/home/me/*.c"
```

Because the shell expands unquoted wildcards, it is easy to exceed the 20 operand limit in batch mode. It is more efficient to let the client process wildcard file specifications because many fewer server interactions are needed to complete the task.

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 xiii 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 “File system and ACL support” on page 84 for supported file systems and ACL support.

Removing operand limits

You can use the *removeoperandlimit* option to specify that Tivoli Storage Manager removes the 20-operand limit for UNIX-family platforms. 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 324


Associating a local snapshot with a server file space

Use the *snapshotroot* option with the **archive** command in conjunction with a third-party application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. The *snapshotroot* option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot. See “Snapshotroot” on page 351 for more information.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax



Parameters

options

Table 68. Archive command: Related options

Option	Where to use	Page
<i>archmc</i>	Command line only.	195
<i>archsymblinkasfile</i>	Client user options file (dsm.opt) or command line.	196
<i>changingretries</i>	Client system options file (dsm.sys) or command line.	201
<i>compressalways</i>	Client user options file (dsm.opt) or command line.	210
<i>compression</i>	dsm.sys file <i>within</i> a server stanza or command line.	211
<i>deletefiles</i>	Command line only.	218
<i>description</i>	Command line only.	219
<i>dirsonly</i>	Command line only.	223

Table 68. Archive command: Related options (continued)

<i>filelist</i>	Command line only.	250
<i>filesonly</i>	Command line only.	253
<i>preservelastaccessdate</i>	Client user options file (dsm.opt) or command line.	316
<i>removeoperandlimit</i>	Command line only.	324
<i>snapshotroot</i>	Command line only.	351
<i>subdir</i>	Client user options file (dsm.opt) or command line.	356
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	358
<i>v2archive</i>	Command line only.	378

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 operations” on page 86 for the maximum file size for archive processing.

Examples

- Task** Archive a single file named budget in the /home/proj1 directory.
Command: archive /home/proj1/budget
- Task** Archive all files in the /home/proj1 directory with a file extension of .txt.
Command: archive "/home/proj1/*.txt"
- Task** Archive all files in the directory tree headed by the /home directory.
Command: archive -subdir=yes "/home/*"
- Task** Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, archive the /usr/dir1/sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name /usr.
Command: dsmc archive /usr/dir1/sub1/ -subdir=yes -snapshotroot=/snapshot/day1

Backup Group

Authorized User

Use the **backup group** command to create and back up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.

A *group backup* allows you to create a consistent point-in-time backup of a group of files that is managed as a single logical entity:

- All objects in the group are assigned to the same management class. See “Include options” on page 269 for more information about using the *include* option to bind a group to a management class.
- Existing *exclude* statements for any files in the group are ignored.
- All objects in the group are exported together.
- All objects in the group are expired together as specified in the management class. No objects in a group are expired until all other objects in the group are expired, even when another group they belong to gets expired.
- If you are performing full and differential group backups to a sequential device, during a restore the data will be in no more than two locations. To optimize restore time, perform periodic full backups to back up the data to one location on the sequential media.
- During a full group backup, all objects in the filelist are sent to the server. During a differential group backup, only data that has changed since the last full backup is sent to the server. Objects in the filelist that have not changed since the last full backup are assigned as members of the differential group backup. This data is not resent to the server, reducing backup time.

The **backup group** command requires the following options:

filelist Specifies a list of files to add to a new group. See “Filelist” on page 250 for more information.

groupname

Specifies the fully qualified name of the group containing a list of files. See “Groupname” on page 259 for more information.

virtualfsname

Specifies the name of the virtual file space for the group on which you want to perform the operation. See “Virtualfsname” on page 381 for more information.

mode

Specifies whether you want to back up all of the files in the filelist or only files that have changed since the last full backup. See “Mode” on page 291 for more information.

Notes:

1. If any file in the group backup fails, the entire group backup will fail.
2. Use the **query group** command to query members of a group backup on the Tivoli Storage Manager server. See “Query Group” on page 451 for more information.
3. Use the **restore group** command to restore specific members or all members of a group backup on the Tivoli Storage Manager server. See “Restore Group” on page 474 for more information.

4. Use the **delete group** command to delete a specific group backup from the Tivoli Storage Manager server. See “Delete Group” on page 425 for more information.
5. Use the **query filespace** command to display virtual file space names for your node that are stored on the Tivoli Storage Manager server. See “Query Filespace” on page 449 for more information.
6. A group backup can be added to a backup set. See “Restoring data from a backupset” on page 130 for more information about backup sets.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

►► Backup GGroup options ◀◀

Parameters

options

Table 69. Backup Group command: Related options

Option	Where to use	Page
<i>filelist</i>	Command line only.	250
<i>groupname</i>	Command line only.	259
<i>mode</i>	Command line only.	291
<i>virtualfsname</i>	Command line only.	381

Examples

Task Perform a full backup of all the files in the /home/dir1/filelist1 file to the virtual file space name /virtfs containing the group leader /home/group1 file.

Command:

```
backup group -filelist=/home/dir1/filelist1 -groupname=group1
-virtualfsname=/virtfs -mode=full
```

Backup Image

The **backup image** command creates an image backup of one or more volumes on your system.

Notes:

1. For the Linux clients, image backup is only supported on partitions with id 0x83 or logical volumes created with the Linux Logical Volume Manager. Backing up other partitions, such as extended partitions that contain mounted file systems or database data, can produce inconsistent backup data if the data changes during the image backup operation.
2. The usage of GPFS snapshot, and image backup is not supported for GPFS file systems on Linux x86/x86_64, Linux on POWER and Linux zSeries.
3. Image backup of the Sun QFS and IBM TotalStorage SAN File Systems is not supported.
4. The Tivoli Storage Manager API must be installed to use the **backup image** command.

The Tivoli Storage Manager client must support the raw device type on the specific platform to perform an image backup of a raw device. You can only perform an image backup on local devices. Clustered devices or file systems as well as devices or file systems shared between two or more systems are not supported. If you want to perform an image backup for a file system mounted on a raw device, the raw device must be supported. See “Volume device type support for an image backup” on page 103 for specific information about supported devices for the **backup image** command.

Use the *include.image* option to include a file system or logical volume for image backup, or to specify volume-specific options for image backup.

The **backup image** command uses the *compression* option value specified in the *dsm.sys*. You can also specify the *compression* option with the **backup image** command.

Static, dynamic, and snapshot image backup

The traditional image backup prevents write access to the volume by other system applications during the operation. Use the *imagetype=dynamic* option to back up the volume *as is* without remounting it read-only. Corruption of the backup can occur if applications write to the volume while the backup is in progress. In this case, run **fsck** after a restore.

For Linux x86/x86_64 clients *only*: Tivoli Storage Manager performs a snapshot image backup of file systems residing on a logical volume created by the Linux Logical Volume Manager during which the volume is available to other system applications. Snapshot image backup requires a Version 5.1 Tivoli Storage Manager server.

You can use the *imagetype* option with the **backup image** command or the *include.image* option to specify whether to perform a static, dynamic, or snapshot image backup. See “Imagetype” on page 265 for more information.

1. A snapshot of the logical volume is taken.
2. Old data blocks containing changes made to the volume during the backup are saved into the cache files.
3. The cache files are backed up to the server and the changes are applied to the volume image when you restore the volume.

The Linux Logical Volume Manager allows the creation of a snapshot of a logical volume while the logical volume itself is still online. The snapshot is created inside the same volume group as the source logical volume. You must ensure that the volume group provides enough free disk space to create the snapshot. The snapshot contains the old data blocks while the modified data is stored in the source logical volume. Use the *snapshotcachesize* option with the **backup image** command, in the *dsm.opt* file, or with the *include.image* option to specify an appropriate snapshot size so that all old data blocks can be stored while the image backup occurs. A snapshot size of 100 percent will ensure a valid snapshot. See “Snapshotcachesize” on page 350 for more information.

Utilizing image backup to perform file system incremental backup

There are two methods of utilizing image backups to perform efficient incremental backups of your file system. These backup methods allow you to perform point-in-time restore of your file systems and improve backup and restore performance. You can perform the backup only on formatted volumes; not on raw logical volumes. You can use one of the following methods to perform image backups of volumes with mounted file systems.

Method 1 Using image backup with file system incremental:

1. Perform a full incremental backup of the file system, for example:

```
dsmc incremental /myfilesystem
```

2. Perform an image backup of the same file system, for example:

```
dsmc backup image /myfilesystem
```

3. Periodically, perform incremental backups, for example:

```
dsmc incremental /myfilesystem
```

You must follow these steps in the order shown to ensure that the server records additions and deletions accurately.

4. The following command restores the file system to its exact state as of the last incremental backup:

```
dsmc restore image /myfilesystem -incremental -deletefiles
```

During the restore, the client does the following:

- Restores the most recent image on the server.
- Deletes all of the files restored in the previous step which are inactive on the server. These are files which existed at the time of the image backup, but were subsequently deleted and recorded by a later incremental backup.
- Restores new and changed files from the incremental backups.

If you do not follow the steps exactly, two things can occur:

- After the original image is restored, all files backed up with the **incremental** command are restored individually.
- If you perform a **backup image** before performing an **incremental**, files deleted from the original image are *not* deleted from the final restored file system.

Method 2 Using image backup with image incremental mode:

1. Perform an image backup of the same file system, for example:

```
dsmc backup image /myfilesystem
```

2. Perform an incremental image backup of the file system, for example:

```
dsmc backup image /myfilesystem -mode=incremental
```

This sends only those files that were added or changed since the last image backup to the server. For more information, see "Mode" on page 291.

3. Periodically, perform full image backups, for example:

```
dsmc backup image /myfilesystem
```

4. Restore the image as follows:

```
dsmc restore image /myfilesystem -incremental
```

On restore, Tivoli Storage Manager ignores the *deletefiles* option when the image+image incremental technique of backing up has been used. The restore will include files that were deleted after the last full image backup plus the latest versions of files added or changed after the last image backup.

Note: You should perform full image backups periodically in the following cases:

- When a file system changes substantially (more than 40%), as indicated in step 3 of methods 1 and 2.
- Once each month.
- As appropriate for your environment.

This will improve restore time because fewer changes are applied from incrementals.

The following restrictions apply when using method 2:


- The file system can have no previous full incremental backups produced by the **incremental** command.
- Incremental-by-date image backup does not inactivate files on the server; therefore, when files are restored, none can be deleted.
- If this is the first image backup for the file system, a full image backup is performed.
- Using *mode=incremental* backs up only files with a changed date, not files with changed permissions.
- If file systems are running at or near capacity, an out-of-space condition could result during the restore.

To help you decide which method is appropriate for your environment, see "Comparing methods 1 and 2" on page 106.

Supported Clients

This command is valid for AIX, HP-UX, all Linux clients, and Solaris.

Syntax



Parameters

options

Table 70. Backup Image command: Related options

Option	Where to use	Page
<i>compressalways</i>	Client user options file (dsm.opt) or command line.	210
<i>compression</i>	Client system options file (dsm.sys) or command line.	211
<i>imagetype</i>	Use with the backup image command or the <i>include.image</i> option in your client system options file (dsm.sys).	265
<i>mode</i>	Command line only.	291
<i>snapshotcachesize</i>	Use with the backup image command, the <i>include.image</i> option, or in the dsm.opt file.	350

filespec

Specifies the name of one or more logical volumes. If you want to back up more than one file system, separate their names with spaces. Do not use pattern matching characters. If you do not specify a volume name, the logical volumes specified with the *domain.image* option will be processed. If you do not use the *domain.image* option to specify file systems to process, an error message is displayed and no image backup occurs.

Specify the file space over which the logical volume is mounted or the logical volume name. If there is a file system configured in the system for a given volume, you cannot back up the volume with the device name. For example, if /dev/lv01 is mounted on /home you can issue backup image /home but backup image /dev/lv01 will fail with an error: ANS1063E Invalid path specified.

For Sun systems: Specify either a file system name or a raw device name (block device type).

Examples

Task Back up the /home/test file space over which the logical volume is mounted and perform an image incremental backup that backs up only new and changed files after the last full image backup.

Command: dsmc backup image /home/test -mode=incremental

Task Perform a static image backup of the logical volume mounted at the /home directory.

Command: dsmc backup image /home -imagetype=static

Task Perform a snapshot image backup of the /home directory.

Command: dsmc backup image /home -imagetype=snapshot

Task Back up the /dev/lv01 raw logical volume.

Command: dsmc backup image /dev/lv01

Backup NAS

The **backup nas** command creates an image backup of one or more file systems belonging to a Network Attached Storage (NAS) file server. The NAS file server performs the outboard data movement. A server process starts in order to perform the backup.

Use the *nasnodename* option to specify the node name for the NAS file server. When using an interactive command line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID. The NAS node name identifies the NAS file server to the Tivoli Storage Manager server; the NAS node name must be registered at the server. Place the *nasnodename* option in your client system options file (dsm.sys). The value in the client system options file is the default, but can be overridden on the command line. See “Nasnodename” on page 294 for more information.

Use the *toc* option with the **backup nas** command or the *include.fs.nas* option to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup. See “Toc” on page 371 for more information. If you save TOC information, you can use the **query toc** server command to determine the contents of a file system backup in conjunction with the **restore node** server command to restore individual files or directory trees. You can also use the Tivoli Storage Manager Web client to examine the entire file system tree and select files and directories to restore. Creation of a TOC requires that you define the TOCDESTINATION attribute in the backup copy group for the management class to which this backup image is bound. Note that TOC creation requires additional processing, network resources, storage pool space, and possibly a mount point during the backup operation. If you do not save TOC information, you can still restore individual files or directory trees using the **restore node** server command, provided that you know the fully qualified name of each file or directory and the image in which that object was backed up. The *toc* option is only supported for images backed up by Version 5.2 or later client and server.

Specifying **MODE =differential** on the **backup node** server command or the **backup nas** command where no full image exists, shows that a full backup was started. Using the **query process** server command shows that a full backup is in process.

Use the *mode* option to specify whether to perform a full or differential NAS image backup. A full image backup backs up the entire file system. The default is a differential NAS image backup on files that change after the last full image backup. If an eligible full image backup does not exist, a full image backup is performed. If a full image exists, whether it is restorable, or expired and being maintained because of dependent differential images, specifying **MODE =differential** sends a differential image backup. If a full image is sent during a *differential* backup, it is reflected as a full image using the **query nasbackup** server command. The **query nasbackup** server command also displays NAS images that are restorable and will display full image or differential image as the object type. See “Mode” on page 291 for more information.

Use the *monitor* option to specify whether you want to monitor a NAS file system image backup and display processing information on your screen. See “Monitor” on page 293.

Use the **monitor process** command to display a list of all processes for which an administrative user ID has authority. The authorized administrative user ID should

have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web.

Use the **cancel process** command to stop NAS backup processing. For more information, see “Cancel Process” on page 415.

Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.

- Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.
- NAS file system designations on the command line require brace delimiters {} around the file system names, such as: {/vol/vol0}.

Supported Clients

This command is valid for AIX and Solaris clients only.

Syntax



Parameters

options

Table 71. Backup NAS command: Related options

Option	Where to use	Page
<i>mode</i>	Command line only.	291
<i>monitor</i>	Command line only.	293
<i>nasnodename</i>	Client system options file (dsm.sys) or command line.	294
<i>toc</i>	Command line or with the <i>include.fs.nas</i> option in your dsm.sys file.	371

filespec

Specifies the name of one or more file systems on the NAS file server. If you do not specify this parameter, Tivoli Storage Manager processes all of the file systems defined by the *domain.nas* option. For more information about this option, see “Domain.nas” on page 232.

If you do not specify the *filespec* or the *domain.nas* option, the default *all-nas* value is used for *domain.nas* and all file systems on the NAS file server are backed up.

Examples

Task Perform the NAS image backup of the entire file system.

Command: backup nas -mode=full -nasnodename=nas1 /vol/vol0 /vol/vol2

Task Perform the NAS image backup of the entire file server.

Command: backup nas -nasnodename=nas1

Task Perform the NAS image backup of the entire file system and save Table of Contents (TOC) information for the file system backup.

Command: backup nas -mode=full -nasnodename=netappsj /vol/vol0
-toc=yes

Backup WAS

Root User

The **backup was** command specifies whether to back up the WebSphere Application Server (WAS) Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server (also contains setup, application files, and configuration information) to the Tivoli Storage Manager server. You can back up both the Network Deployment Manager and the Application Server using separate sessions.

Notes:

1. If WAS security is enabled, user name and password validation for Data Protection for WebSphere Application Server is required. To avoid backup failure, you *must* use the **set waspassword** command to set the user name and password for each installation of WAS on your machine. You only need to perform this task once, unless you change your WAS user name or password. See “Set Waspassword” on page 496 for more information.

To determine if WAS security is enabled, enter the following command:

```
dsmc query was -wast=local
```

Tivoli Storage Manager displays the WAS security status under the **Sec** heading.

2. Multiple backup sessions of the same node are not supported.
3. Use the *mode* option to specify whether to perform a full (the default) or differential backup. See “Mode” on page 291 for more information.
4. Use the *wastype* option to specify whether to back up the Network Deployment Manager (ND) or Application Server (APP) associated with the node name of the instance of WAS that you want to back up. The default is ND. See “Wastype” on page 390 for more information.
5. Use the *include* option in your dsm.sys file to assign a management class to a WAS group backup. For example:
 - For the Network Deployment Manager: include /WAS_ND_NDNODE mgmtclass
 - For the Application Server: include /WAS_APPNODE mgmtclass
6. WAS backups can also be added to a backup set. See “Restoring data from a backupset” on page 130 for more information about WAS backups.

Supported Clients

This command is valid for AIX, Solaris, and Linux x86/x86_64 clients.

Syntax

```
▶▶ Backup WAS [ options ] —nodename▶▶
```

Parameters

options

Table 72. Backup WAS command: Related options

Option	Where to use	Page
<i>mode</i>	Command line only.	291

Table 72. Backup WAS command: Related options (continued)

<i>wastype</i>	Command line only.	390
----------------	--------------------	-----

nodename

Specifies the node name of the instance of WAS to back up. This is a required parameter.

Examples

Task Back up the Network Deployment Manager associated with the node name *wasnode*.

Command: backup was wasnode

Task Back up the Application Server associated with the node name and instance *ednode_instance1*.

Command: backup was ednode_instance1 -wastype=app

Task Perform a differential backup of the Network Deployment Manager associated with the node name and instance *ednode_instance2*.

Command: backup was ednode_instance2 -wastype=nd
-mode=differential

Cancel Process

The **cancel process** command displays a list of current NAS (if NDMP support is enabled) image backup and restore processes for which the administrative user has authority. From the list, the administrative user can select one process to cancel. Client owner privilege is sufficient authority to cancel the selected NAS image backup or restore processes.

When using an interactive command line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID.

Supported Clients

This command is valid for AIX and Solaris clients only.

Syntax

▶▶—Cancel Process—————▶▶

Parameters

There are no parameters for this command.

Examples

Task Cancel current NAS image backup or restore processes.

Command: cancel process

Cancel Restore

The **cancel restore** command displays a list of your restartable restore sessions in the server database. You can only cancel one restartable restore session at a time. Run the **cancel restore** command again to cancel additional restores. To restart restartable restore sessions, use the **restart restore** command.

Use the **cancel restore** command 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.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

▶▶—Cancel Restore—————▶▶

Parameters

There are no parameters for this command.

Examples

Task Cancel a restore operation.

Command: cancel restore

Delete Access

The **delete access** command deletes authorization rules for files or images that are stored on the server. When you delete an authorization rule, you revoke user access to any files or images specified by that rule.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

►►—Delete Access—◄◄

Parameters

There are no parameters for this command.

Examples

Task Display a list of current authorization rules and select the rules you want to delete.

Command: delete access

See the following screen example:

Index	Type	Node	Owner	Path
1	Backup	NODE1	USER1	home/dev/proja/list/
2	Archive	NODE3	LUIE	home/fin/budg/depta/
3	Backup	NODE4	USER2	home/plan/exp/deptc/
4	Archive	NODE5	USER2S	home/mfg/invn/parta/

Enter Index of rule(s) to delete, or quit to cancel:

To delete the authorization rules that let **luie** and **user2s** access your files or images, type: **2 4** or **(2,4)** and press Enter.

Delete Archive

The **delete archive** command deletes archived files from Tivoli Storage Manager server storage. Your administrator must give you authority to delete archived files.

Attention: When you delete archived files, *you cannot retrieve them*. Verify that the files are obsolete *before* you delete them.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

```
▶▶ Delete ARchive [ options ] [ filespec ]
```

Parameters

options

Table 73. Delete Archive command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	214
<i>description</i>	Command line only.	219
<i>filelist</i>	Command line only.	250
<i>noprompt</i>	Command line only.	300
<i>numberformat</i>	Client user options file (dsm.opt) or command line.	301
<i>pick</i>	Command line only.	308
<i>subdir</i>	Client user options file (dsm.opt) or command line.	356
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	358
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	369

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.

Command: delete archive /user/home/proj1/budget

Task Delete all files archived from the /user/home/proj1 directory with a file extension of .txt.

Command: del arch "/user/home/proj1/*.txt"

Task Delete files archived from the /user/project directory using the *pick* option to display a list of archive copies that match the file specification. From the list, you can select the versions to process.

Command: delete archive "/user/project/*" -pick

Task Delete selected files from the group of files archived with the description "Monthly Budgets 1999".

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.

Supported Clients

This command is valid for all UNIX and Linux clients, however *objtype=image* is not supported on z/OS.

Syntax



Parameters

options

Table 74. Delete Backup command: Related options

Option	Where to use	Page
<i>description</i>	Command line only.	219
<i>filelist</i>	Command line only.	250
<i>fromdate</i>	Command line, and in GUI find function.	250
<i>frontime</i>	Command line, and in GUI find function.	250
<i>noprompt</i>	Command line only.	300
<i>pick</i>	Command line only.	308
<i>pitdate</i>	Command line, and in GUI find function.	250
<i>pittime</i>	Command line, and in GUI find function.	250
<i>subdir</i>	Client user options file (dsm.opt) or command line.	356
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	358
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	369
<i>todate</i>	Command line, and in GUI find function.	250
<i>totime</i>	Command line, and in GUI find function.	250

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.

IMAGE

Specifies that you want to delete an image backup.

Examples

Task Delete all active and inactive file objects named budget in directory /data/plan/proj1.

Commands:

```
delete backup /data/plan/proj1/budget.jan -deltype=inactive
delete backup /data/plan/proj1/budget.jan -deltype=active
```

Note: To delete active and inactive file objects, you must first delete the inactive objects, followed by the active objects.

Task Delete all inactive files with names ending in .txt backed up from the /data/plan/proj1 directory and its subdirectories.

Command: delete backup "/data/plan/proj1/*.txt" -deltype=inactive -subdir=yes

Task Delete selected active files backed up from the /home/marymb/project directory. Use the *-pick* option to display a list of backup copies that match the file specification. From the list, you can select which versions to delete.

Command: delete backup "/home/marymb/project/*" -pick

| **Task** Delete all active and inactive versions of files and subdirectories in the
| /home/storman/myproject directory. Then delete all active and inactive
| versions of the /user/myproject directory.

| **Command:** delete backup "/home/storman/myproject*" -deltype=all

| **Note:** The backup versions of directory object /home/storman/myproject
| are not deleted.

Delete Filespace

Authorized User

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 or images 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.

Deleting NAS file spaces

You can use the **delete filesystem** command to interactively delete NAS file spaces from server storage.

Use the *nasnodename* option to identify the NAS file server. When using an interactive command line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID. Place the *nasnodename* option in your client system options file (dsm.sys). The value in the client system options file is the default, but this value can be overridden on the command line. If the *nasnodename* option is not specified in the client system options file, you must specify this option on the command line when processing NAS file systems. See “Nasnodename” on page 294 for more information.

Use the *class* option to specify the class of the file space to delete. To display a list of file spaces belonging to a NAS node so that you can choose one to delete, use the *-class=nas* option. See “Class” on page 202 for more information.

To delete NAS file spaces using the Web client, see Chapter 4, “Backing up your data,” on page 79.

Deleting WebSphere Application Server (WAS) file spaces

Use the **delete filesystem** command to delete a WAS file space on the Tivoli Storage Manager server.

Use the **delete group** command to delete WAS group backups on the Tivoli Storage Manager server. See “Delete Group” on page 425 for more information.

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 or images*. Verify that the files or images are obsolete *before* you delete them.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

Parameters

options

Table 75. Delete Filespace command: Related options

Option	Where to use	Page
<i>class</i>	Command line only.	202
<i>detail</i>	Command line only.	221
<i>nasnodename</i>	Client system options file (dsm.sys) or command line.	294
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	339
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	340

Examples

Task Delete a file space.

Command: delete filesystem

Task Delete NAS file spaces from the **dagordon** NAS file server stored on the server.

Command: delete filesystem -nasnodename=dagordon -class=nas

Task Delete WAS file spaces stored on the server.

Command: delete filesystem

Delete Group

Authorized User

Use the **delete group** command to delete a group backup on the Tivoli Storage Manager server. You can also delete WAS group backups using this command.

After deleting a group, the group leader (virtualfsname) remains on the Tivoli Storage Manager server. It contains no members (file or directories) but is reported in a subsequent **query filespace** command. It will have no files listed if the *showmembers* option is added. Deleting a group does not remove the file space that it resides in because there might be other groups in it. Use **delete filespace** if you want to remove the file space and all the data it contains.

Notes:

1. Use the *inactive* option to display both active and inactive group backup versions. By default, Tivoli Storage Manager only displays active versions. See “Inactive” on page 267 for more information.
2. Use the *pick* option to select a specific group to delete from the Tivoli Storage Manager server. See “Pick” on page 308 for more information.
3. Use the *noprompt* option if you want to suppress the confirmation prompt that normally appears before you delete a group backup version. By default, Tivoli Storage Manager prompts you for confirmation before deleting the group backup. Using this option can speed up the delete procedure. However, it also increases the danger of accidentally deleting a group backup version that you want to save. Use this option with caution. See “Noprompt” on page 300 for more information.
4. Use the **query filespace** command to display virtual file space names for your node that are stored on the Tivoli Storage Manager server. See “Query Filespace” on page 449 for more information.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

```
→ Delete GRoup- filespec [ options ] →
```

Parameters

filespec

Specifies the virtual file space name and the group name that you want to delete from the server storage.

options

Table 76. Delete Group command: Related options

Option	Where to use	Page
<i>inactive</i>	Command line only.	267
<i>noprompt</i>	Command line only.	300
<i>pick</i>	Command line only.	308
<i>pitdate</i>	Command line only.	309

Table 76. Delete Group command: Related options (continued)

<i>pittime</i>	Command line only.	310
----------------	--------------------	-----

Examples

Task Delete the current active version of the /vfs/group1 group.

Command:

```
delete group /vfs/group1
```

Task Delete a backup version of the /vfs/group1 group from a list of active and inactive versions.

Command:

```
delete group /vfs/group1 -inactive -pick
```

Task Delete the current active APP/ND WAS group.

Command: delete group /WAS_APPNODE/WASGROUP

Task Delete a backup version of the WAS ND/APP group from a list of active and inactive versions.

Command: delete group /WAS_ND_NDNODE/WASGROUP -inactive -pick

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.

Note: Because the **expire** command changes the server's picture of the client file system without actually changing the client file system, the **expire** command is not allowed on files located on a file system monitored by the Tivoli Storage Manager journaldaemon.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

```
►►—EXPIre— [ options ]— filespec—►►
```

Parameters

options

Table 77. Expire command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	214
<i>filelist</i>	Command line only.	250
<i>noprompt</i>	Command line only.	300
<i>numberformat</i>	Client user options file (dsm.opt) or command line.	301
<i>pick</i>	Command line only.	308
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	369

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: `expire "/home/letter1.txt"`

Task Inactivate all files in the /admin/mydir directory.

Command: `expire /admin/mydir/*`

Task Inactivate all files named in the /home/avi/filelist.txt file.

Command: `expire -filelist=/home/avi/filelist.txt`

Help

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.

Supported Clients

This command is valid for all UNIX and Linux clients.

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 “File system and ACL support” on page 84 for supported file systems 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 227 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. This management class attribute is ignored during a journal-based backup.

Mode Permits you to back up only files that changed since the last backup (*modified*) or back up the files whether they changed or not (*absolute*).

Serialization

Permits or denies backup of files or directories according to the following values:

- **static**: In order to be backed up, data must not be modified during backup or archive.
- **shared static**: If data in the file or directory changes during each of the allowed attempts to back up or archive it, it is not backed up or archived. The value of the *changingretries* option determines how many attempts are made. The default is 4.
- **dynamic**: The object is backed up or archived on the first attempt whether or not data changes during the process.
- **shared dynamic**: The object is backed up or archived on the last attempt, even if data changes during the process.

For more information on management classes, see Chapter 8, “Understanding storage management policies,” on page 157.

Using the *include* option in an include-exclude list, you can override the default management class for a file or group of files.

You can perform either a *full incremental* backup or an *incremental by date* backup. The default is a full incremental backup.

If you are journaling a file system and the journal is valid, the full incremental backup performs a journal-based backup. More than one journal-based backup session can be started, but only one journal-based backup session can proceed. All other journal-based backup sessions that need access to the same filespace must wait until the current journal-based backup session has completed before the next session can proceed. See “Journal-based backup” on page 431 for more information. You can perform a full incremental backup without the journal by using the *nojournal* option. See “Nojournal” on page 299 for more information.

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 488.

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

Please see “Removeoperandlimit” on page 324 for information about operand limits.

Journal-based backup

Journal-based backup is supported on the AIX Backup-Archive client, on JFS, JFS2, and VxFS filesystems. A backup for a particular file system is journal-based when the Tivoli Storage Manager journal daemon is installed and configured to journal the file system, and a valid journal has been established.

If the *journal daemon* is installed and running, then by default the **incremental** command will perform a journal-based backup on file systems which are being monitored by the journal engine daemon. The following conditions must be met in order to successfully perform a journal-based backup:

- The journal daemon must be set up to monitor the file system that contains the files and directories being backed up.
- A full incremental backup must have been run successfully at least once on the file system being backed up.
- The file space image of the file system at the server cannot have been modified by an administrative command since the last full incremental backup.
- The storage management policy for the files being backed up cannot have been updated since the last full incremental backup.

The *journal daemon* records changes to an object or its attributes in a journal database. During a journal-based backup, the client obtains a list of files that are eligible for backup from the journal database. Journal-based backup can increase backup performance because the client does not scan the local file system or contact the server to determine which files to process. Journal-based backup also reduces network traffic between the client and server

Tivoli Storage Manager filters the list based on the current include-exclude list and processes, expires, and updates the resulting files according to policy constraints, such as serialization. However, the client ignores the server frequency attribute during a journal-based backup. The reason for this is because a journal-based backup eliminates the backup version query to the server; therefore, the client does not know how many days have transpired since the last backup of the file.

The journal daemon does not record changes in UNIX special files.

The journal daemon excludes specific system files from having changes recorded in the journal. Because changes to these files are not journaled, Tivoli Storage Manager does not back up these files. See the journal daemon configuration file `tsmjbbd.ini` located in the Tivoli Storage Manager installation directory for specific system files that are excluded.

Notes:

1. Journal-based backup is supported on a Version 5.2 or higher Tivoli Storage Manager server.
2. When using anti-virus software, there are limitations to journal-based backup.
3. A journal-based backup might not fall back to the traditional incremental backup if the policy domain of your node is changed on the server. This depends on when the policy set within the domain was last updated and the date of the last incremental backup. In this case, you must force a full traditional incremental backup to rebind the files to the new domain. Use the *nojournal* option with the **incremental** command to specify that you want to perform a traditional full incremental backup, instead of the default journal-based backup.

Considerations: Under the following conditions, the journal database is considered invalid and the client reverts to the traditional full incremental backup:

- A journaled file space name has changed.
- The client node name has changed.
- The client contacts a different server to do the backup.
- Policy changes have occurred (new policy set activation).
- The journal is corrupt (out of space conditions, disk error).
- The journal is not running.

Journal-based backup differs from the traditional full incremental backup in the following ways:

- Tivoli Storage Manager does not enforce non-default copy frequencies (other than 0).
- Changes to UNIX special files are not detected.

You can use the *nojournal* option with the **incremental** command to perform a traditional full incremental backup instead of the default journal-based backup. See “Nojournal” on page 299 for more information.

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.


Associating a local snapshot with a server file space

Use the *snapshotroot* option with the **incremental** command in conjunction with a third-party application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. The *snapshotroot* option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot. See “Snapshotroot” on page 351 for more information.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax



Parameters

options

Table 78. Incremental command: Related options

Option	Where to use	Page
<i>changingretries</i>	dsm.sys file or command line.	201
<i>compressalways</i>	Client user options file (dsm.opt) or command line.	210
<i>compression</i>	dsm.sys file <i>within</i> a server stanza or command line.	211
<i>dironly</i>	Command line only.	223
<i>domain</i>	dsm.sys file or the client user options file (dsm.opt) or command line only.	227
<i>filelist</i>	Command line only.	250
<i>filesonly</i>	Command line only.	253
<i>incrbydate</i>	Command line only.	274
<i>memoryefficientbackup</i>	Client user options file (dsm.opt), client system options file (dsm.sys), server, or command line.	289
<i>nojournal</i>	Command line only.	299
<i>preservelastaccessdate</i>	Client user options file (dsm.opt) or command line.	316
<i>removeoperandlimit</i>	Command line only.	324

Table 78. Incremental command: Related options (continued)

<i>snapshotroot</i>	Command line only.	351
<i>subdir</i>	Client user options file (dsm.opt) or command line.	356
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	358

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 operations” on page 86 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 backup for the /home, /usr, and /proj file systems.
Command: Incremental /home /usr /proj
- Task** Run an incremental backup for the /proj/test directory.
Command: Incremental /proj/test/
- Task** Run an incremental-by-date backup for the /home file system.
Command: Incremental -incrbydate /home
- Task** Run an incremental backup of the abc file in the /fs/dir1 directory.
Command: Incremental -subdir=yes /fs/dir1/abc
- Task** Run an incremental backup of the directory object /fs/dir1, but not any of the files in the /fs/dir1 directory.
Command: Incremental /fs/dir1
- Task** Run an incremental backup of the directory object /fs/dir1, all of the files in the /fs/dir1 directory, and all files and subdirectories under /fs/dir1.
Command: Incremental -subdir=yes /fs/dir1/
- Task** Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, run an incremental backup of all files and directories under the local snapshot and manage them on the Tivoli Storage Manager server under the file space name /usr .
Command: dsmc inc /usr -snapshotroot=/snapshot/day1

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 169 for options that you cannot use in interactive mode.

Notes:

1. In loop mode, following a restore operation directly from tape, the mount point is not released in case additional restore requests are made to that volume. If you request a backup operation in the same session and that mount point is the only one available, the backup operation will stop with the following message:

```
Waiting for mount of offline media
```

In this case, the mount point is not released until one of the following conditions is met:

- The device class MOUNTRETENTION limit is satisfied.
 - The client idletimeout period is satisfied.
 - The dsmc loop session is closed after the restore operation completes, allowing you to start a subsequent loop mode session to perform the backup operation.
2. In interactive mode, you cannot enter a file specification that contains national language characters. If a command contains national characters, process the command in batch mode by preceding the command with the executable program name, **dsmc**.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

▶▶—LOOP—▶▶

Parameters

There are no parameters for this command.

Examples

Task Start an interactive command line session.

Command: `dsmc`

At the **tsm>** prompt, enter a command.

- 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.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

►►—Macro— *macroname* —————►►

Parameters

macroname

Specifies the fully qualified name of the file containing the commands.

Examples

The following is an example of how to use the **macro** command.

Task Selectively back up files in the following directories:

/devel/project/proja
/devel/project/projb
/devel/project/projc

Command: macro backabc.mac

where backabc.mac contains the following statements:

Selective /devel/project/proja/
Selective /devel/project/projb/
Selective /devel/project/projc/

Monitor Process

The **monitor process** command displays a list of current NAS (if NDMP support is enabled) image backup and restore processes for which the administrative user has authority. The administrative user can then select one process to monitor. Client owner privilege is sufficient authority to monitor the selected NAS image backup or restore processes.

When using an interactive command line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID.

Supported Clients

This command is valid for AIX and Solaris clients only.

Syntax

►►—MONitor Process—◄◄

Parameters

There are no parameters for this command.

Examples

Task Monitor current NAS image backup or restore processes.

Command: monitor process

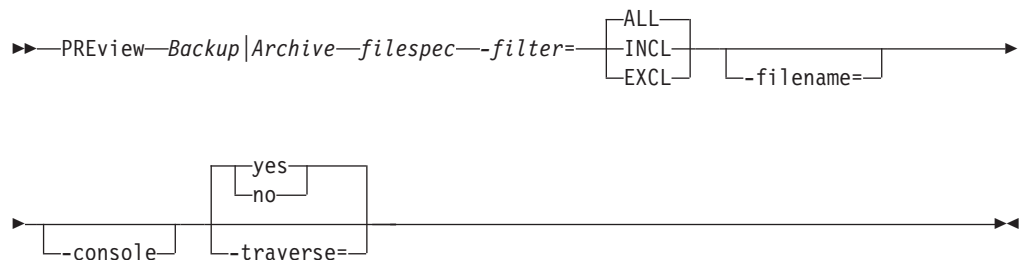
Preview

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.

Supported Clients

This command is valid for all UNIX and Linux clients.

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.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

►►—Query Access—◄◄

Parameters

There are no parameters for this command.

Examples

Task Display a list of users who have access to your files.

Command: query access

Query Archive

The **query archive** command displays a list of your archived files and the following information about each file:

- File size
- Archive date
- File specification
- Expiration date
- Archive description

If you use the *detail* option with the **query archive** command, the client displays the following additional information:

- Last modification date
- Last access date

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

```
►► Query Archive [ options ] [ filespec "filespec" ]
```

Parameters

options

Table 79. Query Archive command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	214
<i>description</i>	Command line only.	219
<i>detail</i>	Command line only.	221
<i>dironly</i>	Command line only.	223
<i>filelist</i>	Command line only.	250
<i>filesonly</i>	Command line only.	253
<i>fromdate</i>	Command line only.	255
<i>fromnode</i>	Command line only.	256
<i>fromowner</i>	Command line only.	257
<i>fromtime</i>	Command line only.	258
<i>numberformat</i>	Client user options file (dsm.opt) or command line.	301
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	339
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	340
<i>subdir</i>	Client user options file (dsm.opt) or command line.	356
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	369
<i>to date</i>	Command line only.	373
<i>totime</i>	Command line only.	374

filespec

Specifies the path and file name that you want to query. Use wildcard

characters to specify a group of files or all the files in a directory. If you use wildcard characters, enclose the file specification in double quotation marks. Specify an asterisk (*) to query all archived files in the current directory.

Examples

- Task** Display a list of all your archived files in the current working directory.
Command: `q archive "*"`
- Task** Display a list of all your archived files in the /devel directory and all of its subdirectories.
Command: `query archive "/devel/*" -subdir=yes`
- Task** Display a list of all your archived files in the current directory. Use the *dateformat* and *timeformat* options to reformat the dates and times.
Command: `q ar -date=5 -time=1 "*"`
- Task** Display a list of all your archived files in the current directory. Use the *detail* option to display the last modification date and the last access date of each file.
Command: `q ar -detail "*"`
- Task** Display a list of archived files in the /home/proj directory whose first four characters of the file name begin with proj.
Command: `q ar "/home/proj/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
- Last access date

Querying NAS file system images

You can use the **query backup** command to display information about file system images backed up for a NAS file server.

Use the *nasnodename* option to identify the NAS file server to query. When using an interactive command line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID. Place the *nasnodename* option in your client system options file (dsm.sys). The value in the client system options file is the default, but this value can be overridden on the command line. See “Nasnodename” on page 294 for more information.

Use the *class* option to specify the class of the file space to query. To display a list of images belonging to a NAS node, use the *-class=nas* option. See “Class” on page 202 for more information.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

```
►► Query Backup [options] [filespec] -BACKUPSetname=backupsetname►►
►► -LOcation=server►►
```

Parameters

options

Table 80. Query Backup command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	200
<i>class</i>	Command line only.	202
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	214

Table 80. Query Backup command: Related options (continued)

<i>detail</i>	Command line only.	221
<i>dirsonly</i>	Command line only.	223
<i>filelist</i>	Command line only.	250
<i>filesonly</i>	Command line only.	253
<i>fromdate</i>	Command line only.	255
<i>fromnode</i>	Command line only.	256
<i>fromowner</i>	Command line only.	257
<i>frontime</i>	Command line only.	258
<i>inactive</i>	Command line only.	267
<i>location</i>	Command line only.	283
<i>nasnodename</i>	Client system options file (dsm.sys) or command line.	294
<i>numberformat</i>	Client user options file (dsm.opt) or command line.	301
<i>pitdate</i>	Command line only.	309
<i>pittime</i>	Command line only.	310
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	339
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	340
<i>subdir</i>	Client user options file (dsm.opt) or command line.	356
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	369
<i>todate</i>	Command line only.	373
<i>totime</i>	Command line only.	374

filespec

Specifies the path and file name that you want to query. Use wildcard characters to specify a group of files or all the files in a directory. If you use wildcard characters, enclose the file specification in double quotation marks. Specify an asterisk (*) to display information about backup versions for all of your files in the current directory. Do not use wild cards when you query NAS file system images with *-class=nas* option.

Examples

- Task** Display a list of all active and inactive backup versions of your files in the current directory.
- Command:** `query backup -inactive "*"`
- Task** Display a list of all your backups in the current directory. Use the *detail* option to display the last modification date and the last access date of each file.
- Command:** `q backup -detail "*"`
- Task** Display a list of files that were backed up from the /home/proj directory with file names that begin with proj.
- Command:** `q b "/home/proj/proj*"`
- Task** Display a list of active and inactive backup file versions in the /home file system.
- Command:** `q b -ina -su=yes /home/`

Task Query file system images from the **nas2** NAS file server.
Command: `query backup -nasnodename=nas2 -class=nas`

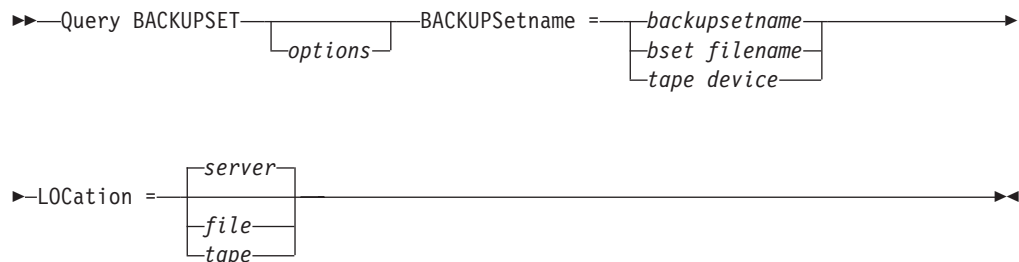
Query Backupset

The **query backupset** command queries a backup set from a local file, tape device, or the Tivoli Storage Manager server. See “Location” on page 283 for information on specifying locally-resident backup sets. This command displays the backup set name, generation date, retention, and description.

Supported Clients

This command is valid for all UNIX and Linux clients. Tape support is only available on AIX, Solaris, and HP-UX.

Syntax



Parameters

options

Table 81. Query Backupset command: Related options

Option	Where to use	Page
<i>description</i>	Command line only.	219
<i>location</i>	Command line only.	283
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	339
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	340

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*.

tape device

Specifies the name of the tape device where the tape containing the start of the backup set is located. The *location* option should be set to *tape*.

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.

tape Specifies that Tivoli Storage Manager searches for the backup set on a local tape device.

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="/home/budget/weekly_budget_data.ost" -loc=file`

Task Query the backup set on the `/dev/rmt0` tape device.

Command: `dsmc query backupset -backupsetname=/dev/rmt0 -loc=tape`

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. The file space name is the same as the file system name.

Querying NAS file spaces


Use the *nasnodename* option to identify the NAS file server to query. When using an interactive command line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID. Place the *nasnodename* option in your client system options file (dsm.sys). The value in the client system options file is the default, but this value can be overridden on the command line. If the *nasnodename* option is not specified in the client system options file, it must be specified on the command line when processing NAS file systems. See “Nasnodename” on page 294 for more information.

Use the *class* option to specify the class of the object to query. To display a list of file spaces belonging to a NAS node, use the *-class=nas* option. See “Class” on page 202 for more information.

Supported Clients

This command is valid for all UNIX and Linux clients.

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 82. Query Filespace command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	200
<i>class</i>	Command line only.	202
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	214
<i>detail</i>	Command line only.	221
<i>fromnode</i>	Command line only.	256
<i>fromowner</i>	Command line only.	257

Table 82. Query Filespace command: Related options (continued)

<i>location</i>	Command line only.	283
<i>nasnodename</i>	Client system options file (dsm.sys) or command line.	294
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	339
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	340
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	369

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 /home file space.

Command: query filesystem /home

Task Display file space names that include the pattern smith.

Command: query filesystem "*smith*"

Task Query a file space from the **nas2** NAS file server.

Command: query filesystem -nasnodename=nas2 -class=nas

Query Group

Authorized User

Use the **query group** command to display information about a group backup and its members.

Notes:

1. Use the *pick* option to display a list of groups from which you can select one group to query.
2. Use the *showmembers* option to display and select individual group members that you want to query. The *showmembers* option is not valid with the *inactive* option. If you want to display members of a group that are not currently active, use the *pitdate* and *pittime* options to specify the backup date and time of the member you want to query.
3. Use the **query filesystem** command to display virtual file space names for your node that are stored on the Tivoli Storage Manager server. See “Query Filespace” on page 449 for more information.
4. If you perform a full and differential group backup, a query of this group using the *-inactive* option displays two active backups of the same name, one of type FULL and one of type DIFF. These backups inactivate any previous full and differential backups:

Size	Backup Date	Mgmt Class	A/I	Type	File
433 B	10/09/2002 14:40:07	NOARCH	A	FULL	VFS/GROUP1
433 B	10/10/2002 07:58:43	NOARCH	A	DIFF	VFS/GROUP1
433 B	10/09/2002 14:39:58	NOARCH	I	FULL	VFS/GROUP1
433 B	10/09/2002 14:39:53	NOARCH	I	DIFF	VFS/GROUP1

If you query a group backup without the *-inactive* option, the query displays only the latest group backup, whether it is type FULL or type DIFF:

Size	Backup Date	Mgmt Class	A/I	Type	File
433 B	10/10/2002 07:58:43	NOARCH	A	DIFF	VFS/GROUP1

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

►►—Query GRoup— *filespec* [*options*]

Parameters

options

Table 83. Query Group command: Related options

Option	Where to use	Page
<i>fromnode</i>	Command line only.	256
<i>fromowner</i>	Command line only.	257

Table 83. Query Group command: Related options (continued)

<i>inactive</i>	Command line only.	267
<i>pitdate</i>	Command line only.	309
<i>pittime</i>	Command line only.	310
<i>showmembers</i>	Command line only.	347

filespec

Specifies the virtual file space name and the group name on the server that you want to query.

Examples

Task Display all the groups in the /vfs file space.

Command:

```
query group /vfs/*
```

Task Display active and inactive versions of the /vfs/group1 filespec.

Command:

```
query group /vfs/group1 -inactive
```

Task Display the /vfs/group1 filespec. Use the *showmembers* option to display a list of group members from which you can select one or more to query.

Command:

```
query backup /vfs/group1 -showmembers
```

Query Image

The **query image** command displays the following information about file system images that are stored on the Tivoli Storage Manager server, or inside a backupset when a backupset is specified.

- Image Size - This is the volume size which was backed up.
- Stored Size - This is the actual image size stored on the server. The stored image on the Tivoli Storage Manager server is the same size as the volume capacity.
- File system type
- Backup date and time
- Management class assigned to image backup
- Whether the image backup is an active or inactive copy
- The image name

Note: The Tivoli Storage Manager API must be installed to use the **query image** command.

Supported Clients

This command is valid for AIX, HP-UX, all Linux clients, and Solaris.

Syntax

```

>> query Image [options] [logicalvolumename | filespaceName]
>> -BACKUPSetName=backupsetname -LOCation=server
  
```

Parameters

options

Table 84. Query Image command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	200
<i>dateformat</i>	Client user option file (dsm.opt) or command line.	214
<i>fromnode</i>	Command line only.	256
<i>fromowner</i>	Command line only	257
<i>inactive</i>	Command line only.	267
<i>location</i>	Command line only.	283
<i>numberformat</i>	Client user option file (dsm.opt) or command line.	301
<i>pitdate</i>	Command line only.	309
<i>pittime</i>	Command line only.	310
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	339
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	340
<i>timeformat</i>	Client user option file (dsm.opt) or command line.	369

logicalvolumename

The name of a logical volume you want to query. You must specify the exact name of the image. You cannot use wildcards. The default is all active images (unless restricted by one or more options).

filespace

Specifies the file system name that you want to query.

Omitting *logicalvolumename* and *filespace* causes all images to be displayed.

Examples

Task Display all backed up images.

Command: q image

Task Display all backed up images owned by **kutras** at node avalon.

Command: query image -fromnode=avalon -fromowner=kutras

Task Display active and inactive version of the /usr image.

Command: q i /usr -inactive

Task Display all images that are contained within the backupset name weekly_backup_data.3214.

Command: query image -backupsetname=weekly_backup_data.3214

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 wish to use in your include-exclude list before you actually insert them in your options file. See the *test pattern* explanation below.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

►►—Query INCLexcl—┐
└—*test pattern*—┘

Parameters

test pattern

Use for testing the validity of patterns you wish 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 Mgmtclass

The **query mgmtclass** command displays information about the management classes available in your active policy set.

Your administrator defines management classes that contain attributes controlling whether a file is eligible for backup or archive services. Management classes also determine how backups and archives are managed on the server.

Your active policy set contains a default management class; it can contain any number of additional management classes. You can assign specific management classes to files using *include* options that are located in the client user options file (dsm.opt). If you do not assign a management class to a file, Tivoli Storage Manager uses the default management class.

When you archive files, you can override the assigned management class by using the *archmc* option.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

```
▶▶ Query Mgmtclass [ options ] ▶▶
```

Parameters

options

Table 85. Query Mgmtclass command: Related options

Option	Where to use	Page
<i>detail</i>	Command line only.	221
<i>frommode</i>	Command line only.	256

Examples

Task Display default and available management classes.

Command: query mgmtclass

Query Node

The **query node** command displays all the nodes for which an administrative user ID has authority to perform operations. The authorized administrative user ID should have at least client owner authority over the client workstation node they are using either from command line or from the web.

When using an interactive command line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID.

Use the *type* option to specify the type of node to filter for. Valid values are *nas*, *client*, *server*, and *any*. The default is *any*. See “Type” on page 376 for more information.

Supported Clients

This command is valid for AIX and Solaris clients only.

Syntax

►► Query Node options ◀◀

Parameters

options

Table 86. Query Node command: Related options

Option	Where to use	Page
<i>type</i>	Command line only.	376
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	339
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	340

Examples

Task Display all NAS nodes.

Command: query node -type=nas

Query Options

Use the **query options** command to display all or part of your options and their current settings relevant to the command line client.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

►► Query Options options pattern ◄◄

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 87. Query Options command: Related options

Option	Where to use	Page
<i>scrolllines</i>	Client user options file (dsm.opt) or command line.	339
<i>scrollprompt</i>	Client user options file (dsm.opt) or command line.	340

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.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

►►—Query Restore—◄◄

Parameters

There are no parameters for this command.

Examples

Task Display your restartable restore session in the server database.

Command: query restore

Query Schedule

The **query schedule** command displays the events scheduled for your node. Your administrator can set up schedules to perform automatic backups and archives for you. To plan your work, use this command to determine when the next scheduled events occur.

The enhanced **query schedule** command on a Tivoli Storage Manager Version 5.3 and above client reports new parameters. The **query schedule** command prior to Tivoli Storage Manager Version 5.3 is referred to as classic. Refer to the Administrator's Guide and Administrator's Reference for your operating system for more information on the classic and enhanced commands.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

►►—Query Schedule—◄◄

Parameters

There are no parameters for this command.

Examples

Task Display your scheduled events.

Command: query schedule

Query Session

The **query session** command displays information about your session, including the current node name, when the session was established, server information, and server connection information.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

►►—Query SEssion—◄◄

Parameters

There are no parameters for this command.

Examples

Task Display your session information.

Command: query session

A sample **query session** display follows:

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
```

Task

A sample **query session** display with LAN-free enabled follows:

TSM Server Connection Information

```
Server Name.....: SERVER1
Server Type.....: Windows
Server Version.....: Ver. 5, Rel. 3, Lev. 0.0
Last Access Date.....: 07/02/2004 12:45:02
Delete Backup Files.....: "Yes"
Delete Archive Files.....: "Yes"
```

```
Node Name.....: GOETHE
User Name.....:
```

```
Storage Agent Name.....: AVI_STA
Storage Agent Type.....: Windows
Storage Agent Version...: Ver. 5, Rel. 2, Lev. 3.
```

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 (includes ULIMIT information for UNIX).
- POLICY - Policy set dump.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually dsmsched.log).
- CLUSTER - AIX cluster information.
- ENCRYPT - Available encryption methods.

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 dsminfo.txt file. See “Filename” on page 252 for more information.
2. Use the *console* option if you want to output the information to the console. See “Console” on page 213 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.

Supported Clients

This command is valid for all UNIX and Linux clients.

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 88. Query Systeminfo command: Related options

Option	Where to use	Page
<i>console</i>	Command line only.	213
<i>filename</i>	Command line only.	252

Examples

Task Gather and store the contents of the dsm.opt file and the Tivoli Storage Manager error log file in the tsminfo.txt file.

Command: query systeminfo dsmoptfile errorlog
-filename=tsminfo.txt

Query WAS

Authorized User

Use the **query was** command to display backups of the WebSphere Application Server (WAS) Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server that match the node name and type of the WAS group backup that you specify.

Use the *wastype* option to specify whether to query the Network Deployment Manager (ND), Application Server (APP), or both (ANY), that are associated with the node name of the instance of WAS that you want to query. The default is ANY. You can also set the *wastype* option to LOCAL to query all the instances of the Application server and Network Deployment Manager on your local machine. See “Wastype” on page 390 for more information.

Supported Clients

This command is valid for AIX, Solaris, and Linux x86/x86_64 clients clients.

Syntax

```
►► Query WAS [ options ] -nodename◄◄
```

Parameters

options

Table 89. Query WAS command: Related options

Option	Where to use	Page
<i>fromnode</i>	Command line only.	256
<i>fromowner</i>	Command line only	257
<i>inactive</i>	Command line only.	267
<i>pitdate</i>	Command line only.	309
<i>pittime</i>	Command line only.	310
<i>showmembers</i>	Command line only.	347
<i>wastype</i>	Command line only.	390

nodename

Specifies the node name of the group that you want to query. If you do not specify a node name, Tivoli Storage Manager queries all WAS backups.

Examples

Task Query all WAS backups (Network Deployment Manager and Application Server) on the Tivoli Storage Manager server.

Command: query was -ina -wastype=any

Task Query the group backups of the Network Deployment Manager associated with the node name *wasnode*. Use the *showmembers* option to display a list of group members from which you can select one or more to query.

Command: query was wasnode -showmembers

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.

Supported Clients

This command is valid for all UNIX and Linux clients.

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 318.

See “File system and ACL support” on page 84 for supported file systems.

Note: On UNIX and Linux systems when a symbolic link is created its modification time is set to the current system time and cannot be changed. So, when restoring a symbolic link its modification date and time is set to the date and time of the restore, not to the date and time the link had when it was backed up. As a result, Tivoli Storage Manager will back up the symbolic link during the next incremental backup because its modification time changed since the last backup.

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.

Supported Clients

This command is valid for all UNIX and Linux clients.

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 90. Restore command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	200
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	214
<i>dironly</i>	Command line only.	223
<i>filelist</i>	Command line only.	250
<i>filesonly</i>	Command line only.	253
<i>followsymbolic</i>	Client user options file (dsm.opt) or command line.	254
<i>fromdate</i>	Command line only.	255
<i>fromnode</i>	Command line only.	256
<i>fromowner</i>	Command line only.	257
<i>fromtime</i>	Command line only.	258
<i>ifnewer</i>	Command line only.	263
<i>inactive</i>	Command line only.	267
<i>latest</i>	Command line only.	281
<i>location</i>	Command line only.	283

Table 90. Restore command: Related options (continued)

<i>numberformat</i>	Client user options file (dsm.opt) or command line.	301
<i>pick</i>	Command line only.	308
<i>pitdate</i>	Command line only.	309
<i>pittime</i>	Command line only.	310
<i>preservepath</i>	Command line only.	318
<i>replace</i>	Client user options file (dsm.opt) or command line.	325
<i>subdir</i>	Client user options file (dsm.opt) or command line.	356
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	358
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	369
<i>todate</i>	Command line only.	373
<i>totime</i>	Command line only.	374

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.

Note: If you do not specify a destination, Tivoli Storage Manager determines whether the original file system can be reached. If the original file system cannot be reached, Tivoli Storage Manager will not restore the file. In this case, you can specify a different destination and retry the command.

Examples

Task Restore a single file named budget in the /Users/user1/Documents directory.

Command: restore /home/devel/projecta/budget

Task Restore a single file named budget which resides in the current directory.

Command: restore file budget

Task Restore all files with a file extension of .c from the /home/devel/projecta directory.

Command: restore "/home/devel/projecta/*.c"

Task Restore files in the /user/project directory. Use the *pick* and *inactive* options to select active and inactive backup versions.

Command: restore "/user/project/*" -pick -inactive

Task Restore all files from the /home/devel/projecta directory that end with the character .c to the /home/newdevel/projectn/projecta directory. If the projectn or the projectn/projecta directory does not exist, it is created.

Command: restore "/home/devel/projecta/*.c" /home/newdevel/projectn/

Task Restore all files in the /home/mydir directory to their state as of 1:00 PM on August 17, 2002.

Command: restore -pitd=8/17/2002 -pitt=13:00:00 /home/mydir/

Task Restore all objects in the /home/myid/ directory. Since this restore is fully wildcarded, if the restore process is interrupted, a restartable restore session is created.

Command: res /home/myid/*

Task Restore files specified in the filelist to a different location.

Command: res -filelist=/home/avi/restorelist.txt
/home/NewRestoreLocation/

Restore Backupset

The **restore backupset** command restores a backup set from the server, a local file, or a local tape device.

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 283 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=<backupsetname>
  -location=<server|tape|file> {/fsname}/* /destfs/ -subdir=yes
```

You can restore a group from a backup set with the following considerations:

- 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 mybackupset /virtfs/* /home/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.

- If you want to restore a local backup set, the virtual file space is not a known real file system on the local system, so the *sourcefilespec* must be enclosed in {} braces, followed with a terminating directory delimiter. For example:

```
restore backupset mybackupset {/virtfs}/* /home/devel/projectb/
  -backupsetname=mybackupset -loc=file -subdir=yes
```

Considerations:

- You must be a root 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 *tape* or *file*, no attempt is made to contact the server.
- There is no support in the API for the backup set format. Therefore, backup set data that was backed up via the API cannot be restored or used.
- 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.

Restoring backup sets in a SAN environment

You can restore backup sets in a storage area network (SAN) in the following ways:


- If the backup set is on a SAN-attached storage device, specify the device using the *filename* parameter and use the *location=tape* option. Tivoli Storage Manager restores the backup set directly from the SAN-attached storage device, gaining high-speed restore performance.

- Note:** You must ensure that the correct tape is mounted in the SAN-attached tape drive prior to issuing the **restore** command. The backup-archive client will not initiate a SCSI autochanger to mount the tape automatically.
- If the backup set is not on local media or a SAN-attached storage device, you can specify the backup set using the *backupsetname* option. Use the *location=server* option to restore the backup set directly from the server via the LAN.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax



Parameters

options

Table 91. Restore Backupset command: Related options

Option	Where to use	Page
<i>dirsonly</i>	Command line only.	223
<i>filesonly</i>	Command line only.	253
<i>ifnewer</i>	Command line only.	263
<i>location</i>	Command line only.	283
<i>preservepath</i>	Command line only.	318
<i>quiet</i>	Client user options file (dsm.opt) or command line.	323
<i>replace</i>	Client user options file (dsm.opt) or command line.	325
<i>subdir</i>	Client user options file (dsm.opt) or command line.	356

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 via LAN. If the backup set contains files from several owners, the backup set itself will be owned by root, and non-root users will not see the backup set. In this case, non-root users can restore their files by obtaining the backup set name from the Tivoli Storage Manager administrator.

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*.

tape device

Specifies the name of the tape device where the tape containing the start of the backup set is located. The *location* option should be set to *tape*.

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.

tape Specifies that Tivoli Storage Manager searches for the backup set on a local tape device.

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 budget directory.

Command: `dsmc restore backupset -backupsetname="/home/budget/
weekly_budget_data.ost" -loc=file`

Task Restore a backup set from the `/dev/rmt0` device.

Command: `dsmc restore backupset "-backupsetname=/dev/rmt0"
-loc=tape`

Task Restore a single file named `budget.dev` from the `/dev/rmt0` device, to the original source path.

Command: `dsmc restore backupset /dev/rmt0 "/home/jones/budget.dev"
-loc=tape`

Task Restore all files in the budget directory that contain a file extension of `.txt` from the tape(s) on the `/dev/rmt0` device, to the original source path.

Command: `dsmc restore backupset "/home/budget/*.txt"
-backupsetname=/dev/rmt0 -loc=tape`

Task Restore the backup set `bset01.001` from the Tivoli Storage Manager server.

Command: `dsmc restore backupset -backupsetname=bset01.001
-loc=server`

Task Restore a group from the backup set `mybackupset` on the Tivoli Storage Manager server to the `/home/devel/projectb` directory.

Command: `dsmc restore backupset /virtfs/ /home/devel/projectb/
-backupsetname=mybackupset -loc=server -subdir=yes`

Task Restore a group from the local backup set `mybackupset.ost` to the `/home/devel/projectb/` directory.

Command: `dsmc restore backupset {/virtfs}/ /home/dev1/projectb/
-backupsetname=mybackupset.ost -loc=server -subdir=yes`

Task Restore the backup set contained in local file `"/home/jones/bset01.file"`.

Command: `dsmc restore backupset -backupsetname="/home/jones/
bset01.file" -loc=file`

Restore Group

Authorized User

Use the **restore group** command to restore specific members or all members of a group backup.

Notes:

1. Use the *pick* option to display a list of groups from which you can select one group to restore.
2. Use the *showmembers* option with the *pick* option to display and restore one or more members of a group. In this case, you first select the group from which you want to restore specific members, then you select one or more group members to restore.
3. You can restore a group from a backup set. See “Restore Backupset” on page 470 for more information.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

```
►► restore group [ options ] sourcefilespec [ destinationfilespec ] ◀◀
```

Parameters

options

Table 92. Restore Group command: Related options

Option	Where to use	Page
<i>followsymbolic</i>	Client user options file (dsm.opt) or command line.	254
<i>fromdate</i>	Command line only.	255
<i>fromnode</i>	Command line only.	256
<i>fromowner</i>	Command line only.	257
<i>fromtime</i>	Command line only.	258
<i>ifnewer</i>	Command line only.	263
<i>inactive</i>	Command line only.	267
<i>latest</i>	Command line only.	281
<i>pick</i>	Command line only.	308
<i>pitdate</i>	Command line only.	309
<i>pittime</i>	Command line only.	310
<i>replace</i>	Client user options file (dsm.opt) or command line.	325
<i>showmembers</i>	Command line only.	347
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	358
<i>to date</i>	Command line only.	373
<i>totime</i>	Command line only.	374

sourcefilespec

Specifies the virtual file space name and the group name on the server that you want to restore.

destinationfilespec

Specifies the path where you want to place the group or one or more group members. If you do not specify a destination, the client restores the files to their original location.

Examples

Task Restore all members in the `/virtfs/group1` group backup to their original location on the client machine.

Command:

```
restore group /virtfs/group1
```

Task Display all groups within the `/virtfs` virtual file space. Use the *showmembers* option to display a list of group members from which you can select one or more to restore.

Command:

```
restore group /virtfs/* -pick -showmembers
```

Task Display a list of groups within the `/virtfs` virtual file space from which you can select one or more groups to restore.

Command:

```
restore group /virtfs/* -pick
```

Restore Image

The **restore image** command restores a file system or raw volume image that was backed up using the **backup image** command. The restore obtains the backup image from a Tivoli Storage Manager server, or inside a backupset when a backupset is specified. This command can restore an active base image, or a point-in-time base image, with associated incremental updates.

Note: Using the *incremental* option with the **restore image** command to perform a dynamic image backup is not supported.

You can use the *verifyimage* option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log. See “Verifyimage” on page 380 for more information.

If bad sectors are present on the target volume, you can use the *imagetofile* option with the **restore image** command to specify that you want to restore the source image to a file. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume. See “Imagetofile” on page 264 for more information.

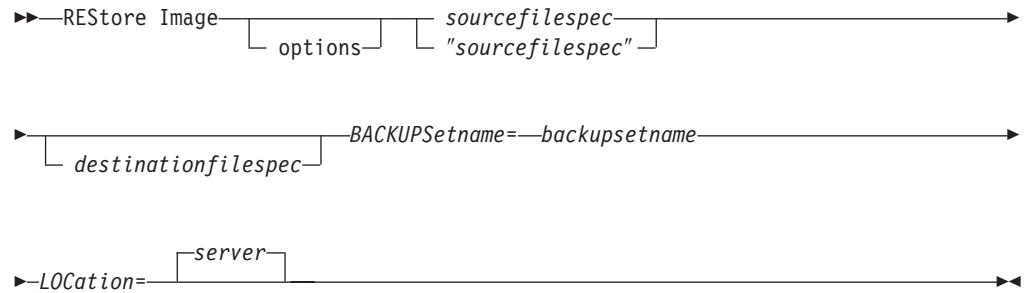
Considerations:

- The API must be installed to use the **restore image** command.
- Image restore of the Sun QFS file system is not supported.
- Image restore is not supported for GPFS file systems on Linux x86/x86_64, Linux on POWER and Linux zSeries.
- If you use the *pick* option, the following information is displayed for file system images backed up by the client:
 - Image Size
 - Stored Size - This is the actual image size stored on the server. The stored image on the Tivoli Storage Manager server is the same size as the volume capacity.
 - File system type
 - Backup date and time
 - Management class assigned to image backup
 - Whether the image backup is an active or inactive copy
 - The image name
- If for some reason a restored image is corrupted, you can use the **fsck** tool to attempt to repair the image.

Supported Clients

This command is valid for AIX, HP-UX, all Linux clients, and Solaris.

Syntax



Parameters

options

Table 93. Restore Image command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	200
<i>dateformat</i>	Client user option file (dsm.opt) or command line.	214
<i>deletefiles</i>	Command line only.	218
<i>fromnode</i>	Command line only.	256
<i>fromowner</i>	Command line only.	257
<i>imagetofile</i>	Command line only.	264
<i>inactive</i>	Command line only.	267
<i>incremental</i>	Command line only.	275
<i>noprompt</i>	Command line only.	300
<i>pick</i>	Command line only.	308
<i>pitdate</i>	Command line only.	309
<i>pittime</i>	Command line only.	310
<i>timeformat</i>	Client user option file (dsm.opt) or command line.	369
<i>verifyimage</i>	Command line only.	380

sourcefilespec

Specifies the name of a source image file system to be restored. Only a single source image can be specified; you cannot use wildcard characters.

destinationfilespec

Specifies the name of an existing mounted file system or the path and file name to which the source file system will be restored. The default is the original location of the file system.

The **restore image** command does not define or mount the destination file space. The destination volume must exist, must be large enough to hold the source, and if it contains a file system, must be mounted. If an image backup contains a file system, and you restore them to a different location, be aware of the following points:

- If the destination volume is smaller than the source volume, the operation will fail.

- If the destination volume is larger than the source, after the restore operation you will *lose* the difference between the sizes. The lost space can be recovered by increasing the size of the volume. This will also increase the size of the restored volume.

Examples

Task Restore the /home/test directory over which the logical volume is mounted, to its original location.

Command: `dsmc rest image /home/test`

Task Restore the /home/proj directory over which the logical volume is mounted, to its original location and apply the changes from the last incremental backup of the original image recorded at the server. The changes include deletion of files.

Command: `dsmc restore image /home/proj -incremental -deletefiles`

Task Restore the /usr file system to its original location. Use the *verifyimage* option to enable detection of bad sectors on the target volume.

Command: `dsmc restore image /usr -verifyimage`

Task If bad sectors present on the target volume, use the *imagetofile* option to restore the /usr file system to the /home/usr.img file to avoid data corruption.

Command: `dsmc restore image /usr /home/usr.img -imagetofile`

Restore NAS

The **restore nas** command restores the image of a file system belonging to a Network Attached Storage (NAS) file server. The NAS file server performs the outboard data movement. A server process performs the restore.

If you used the *toc* option with the **backup nas** command or the *include.fs.nas* option to save Table of Contents (TOC) information for each file system backup, you can use the **query toc** server command to determine the contents of a file system backup in conjunction with the **restore node** server command to restore individual files or directory trees. You can also use the Web client to examine the entire file system tree and select files and directories to restore. If you do not save TOC information, you can still restore individual files or directory trees using the **restore node** server command, provided that you know the fully qualified name of each file or directory and the image in which that object was backed up.

Use the *nasnodename* option to specify the node name for the NAS file server. When using an interactive command line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID. The NAS node name identifies the NAS file server to the Tivoli Storage Manager server. You must register the NAS node name at the server. Place the *nasnodename* option in your client system options file (*dsm.sys*). The value in the client system options file is the default, but this value can be overridden on the command line. See “Nasnodename” on page 294 for more information.

You can use the *pick* option to display a list of NAS images owned by the NAS node you specify. From this list you can select one or more images to restore. If you select multiple images to restore using the *pick* option, do not use the *monitor* option or you will serialize the restores. To start multiple restore processes simultaneously when restoring multiple images, do not specify *monitor=yes*.

Use the *monitor* option to specify whether you want to monitor a NAS file system image restore and display processing information on your screen. See “Monitor” on page 293.

Use the **monitor process** command to display a list of current restore processes for all NAS nodes for which your administrative user ID has authority. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web.

Use the **cancel process** command to stop NAS restore processing. For more information, see “Cancel Process” on page 415.

Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.

Supported Clients

This command is valid for AIX and Solaris clients only.

Syntax

```
►►—REStore NAS options sourcefilespec destinationfilespec►►
```

Parameters

options

Table 94. Restore NAS command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client user option file (dsm.opt) or command line.	214
<i>inactive</i>	Command line only.	267
<i>mode</i>	Command line only.	291
<i>monitor</i>	Command line only.	293
<i>nasnodename</i>	Client system options file (dsm.sys) or command line.	294
<i>numberformat</i>	Client user option file (dsm.opt) or command line.	301
<i>pick</i>	Command line only.	308
<i>pitdate</i>	Command line only.	309
<i>pittime</i>	Command line only.	310
<i>timeformat</i>	Client user option file (dsm.opt) or command line.	369

sourcefilespec

Specifies the name of the NAS file system image you want to restore. This parameter is required unless you use the *pick* option to display a list of NAS images from which to choose. You cannot use wildcard characters when specifying the *sourcefilespec*.

destinationfilespec

Specifies the name of an existing mounted file system on the NAS device over which you want to restore the image. This parameter is optional. The default is the original location of the file system on the NAS device.

Examples

Task Restore the NAS file system image /vol/vol1 to the /vol/vol2 file system on the NAS file server called **nas1**.

Command: restore nas -nasnodename=nas1 /vol/vol1 /vol/vol2

Task Restore inactive NAS images.

Command: restore nas -nasnodename=nas2 -pick -inactive

Restore WAS

Root User

The **restore was** command specifies whether to restore the WebSphere Application Server (WAS) Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server from the Tivoli Storage Manager server. The Application Server must be stopped for the restore to proceed.

You can also restore the Network Deployment Manager, the Application Server, and their instances simultaneously in separate sessions.

Use the *wastype* option to specify whether to restore the Network Deployment Manager (ND) or Application Server (APP) associated with the node name of the instance of WAS that you want to restore. The default is ND. See “Wastype” on page 390 for more information.

WAS instance restore procedures

You can use the following procedure if your entire WebSphere installation is corrupted and it must be reinstalled, or you are installing WAS on a new machine:

1. Install WAS on a machine with the same name and IP address as the machine on which the back up was performed. If you backed up a Network Deployment Manager on machine *ebarton*, the restore must also be performed on machine *ebarton*. Install the same type of WAS as the back up (either ND or APP). When installing, choose the same node and cell names as the backup. For example, if you backed up an ND with a cell name of EdCell and node name of EdNode then the new server must also use these names. Install to the same location in which the backup was performed. For example, if the backup was performed on an APP server installed at `/home/WebSphere/App`, the new server must also be installed at `/home/WebSphere/App`. After installation ensure that the server is stopped.

Notes:

- a. You can use the *washome* option in your client user options file (`dsm.opt`) to specify an override base install path for the Application Server. See “Washome” on page 387 for more information.
 - b. You can use the *wasndhome* option in your client user options file (`dsm.opt`) to specify an override base install path for the Network Deployment manager. See “Wasndhome” on page 388 for more information.
2. Restore the WAS node.
 3. Start the server.
 4. If the server is an ND, you can now attach any remote nodes. The ND and remote nodes will then synchronize.

Use the following procedure to restore to the same existing instance of a server. This procedure assumes that there is a preexisting WAS installation that is running and has a backup stored on the Tivoli Storage Manager server:

1. Stop the WAS server to be restored.
2. Restore the correct WAS node (same *wastype* and node name as the server that was stopped) to the same location.
3. Start the server.

On an ND, the server will synchronize the restored data with the remote nodes since they are already connected.

Notes:

1. For proper operation, the was node must be restored to same location and under same name.
2. WARNING! Restoring data other than at the group level can corrupt your WAS installation. It is strongly recommended that you restore data at the Network Deployment Manager node or Application Server node level only.

Supported Clients

This command is valid for AIX, Solaris, and Linux x86/x86_64 clients clients.

Syntax

```
➤ REStore WAS [ options ] -nodename ➤
```

Parameters

options

Table 95. Restore WAS command: Related options

Option	Where to use	Page
<i>fromdate</i>	Command line only.	255
<i>fromnode</i>	Command line only.	256
<i>fromowner</i>	Command line only.	257
<i>fromtime</i>	Command line only.	258
<i>ifnewer</i>	Command line only.	263
<i>inactive</i>	Command line only.	267
<i>latest</i>	Command line only.	281
<i>pick</i>	Command line only.	308
<i>pitdate</i>	Command line only.	309
<i>pittime</i>	Command line only.	310
<i>preservepath</i>	Command line only.	318
<i>replace</i>	Client user options file (dsm.opt) or command line.	325
<i>showmembers</i>	Command line only.	347
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	358
<i>todate</i>	Command line only.	373
<i>totime</i>	Command line only.	374
<i>wastype</i>	Command line only.	390

nodename

Specifies the node name of the instance of WAS that you want to restore.

Examples

Task Restore the Network Deployment Manager associated with the node name *wasnode*. Use the *pick* option to restore a specific group backup. Use the *showmembers* option to display a list of group members from which you can select one or more to restore.

Command: restore was wasnode -showmembers -pick

Task Restore the Network Deployment Manager wasnode to the original location.

Command: restore was wasnode -wastype=nd

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 194.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

```
➤➤ RETRIEVE [options] [sourcefilespec] [destinationfilespec]➤➤
```

The diagram shows the syntax of the RETRIEVE command. It starts with two right-pointing arrows (➤➤) followed by the word RETRIEVE. Then, there are three optional arguments in square brackets: [options], [sourcefilespec], and [destinationfilespec]. The sourcefilespec argument is shown with an example value "sourcefilespec" in quotes. The diagram ends with two right-pointing arrows (➤➤).

Parameters

options

Table 96. Retrieve command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client user options file (dsm.opt) or command line.	214
<i>description</i>	Command line only.	219
<i>dirsonly</i>	Command line only.	223
<i>filelist</i>	Command line only.	250
<i>filesonly</i>	Command line only	253
<i>followsymbolic</i>	Client user options file (dsm.opt) or command line.	254
<i>fromdate</i>	Command line only	255
<i>fromnode</i>	Command line only.	256
<i>fromowner</i>	Command line only	257
<i>fromtime</i>	Command line only	258
<i>ifnewer</i>	Command line only	263
<i>pick</i>	Command line only.	308
<i>preservepath</i>	Command line only.	318
<i>replace</i>	Client user options file (dsm.opt) or command line.	325
<i>subdir</i>	Client user options file (dsm.opt) or command line.	356
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	358
<i>timeformat</i>	Client user options file (dsm.opt) or command line.	369
<i>todate</i>	Command line only.	373

Table 96. Retrieve command: Related options (continued)

<i>totime</i>	Command line only.	374
---------------	--------------------	-----

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 operations” on page 86 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.

Note: If you do not specify a destination, Tivoli Storage Manager determines whether the original file system can be reached. If the original file system cannot be reached, Tivoli Storage Manager will not restore the file. This failure can also occur if you remove the *virtualmountpoint* option from the dsm.sys file. In this case, you can specify a different destination or restore the original *virtualmountpoint* option to the dsm.sys file, restart the client, and retry the command.

Examples

Task Retrieve a single file named budget.

Command: retrieve /home/devel/projecta/budget

Task Retrieve all files with an extension of .c from the /home/devel/projecta directory.

Command: retrieve "/home/devel/projecta/*.c"

Task Retrieve all files in the /home directory.

Command: retrieve /home/

Task Retrieve all files with a file extension of .c from the /home/devel/projecta directory to the /home/newdevel/projectn/projecta directory. If the /projectn or the /projectn/projecta directory does not exist, it is created.

Command: retrieve "/home/devel/projecta/*.c" /home/newdevel/projectn/

Task Retrieve files in the /user/project directory. Use the *pick* option.

Command: ret "/user/project/*" -pick

Task Retrieve all files archived from the /proj directory with the description "2002 survey results."

Command: retrieve "/proj/*" -desc="2002 survey results"

Task Retrieve archived file /home/devel/budget with description "my budget" to the /dev/rmt1 tape drive.

Command:

```
mkfifo fifo
dd if=fifo of=/dev/rmt1&
dsmc retrieve -replace=yes -description="mybudget"
/home/devel/budget fifo
```

Schedule

Authorized User

The **schedule** command starts the client scheduler on your workstation. The client scheduler must be running before scheduled work can start.

Notes:

1. The **schedule** command cannot be used if the *managedservices* option is set to *schedule*.
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 *schedmode* option is set to polling, the client scheduler contacts the server for scheduled events at the hourly interval you specified with the *queryschedperiod* option in your client user options file (dsm.opt). If your administrator sets the *queryschedperiod* option for all nodes, that setting overrides the client setting.

If you are using TCP/IP communications, the server can prompt your workstation when it is time to run a scheduled event. To do so, set the *schedmode* option to *prompted* in the client user options file (dsm.opt) or on the **schedule** command.

You can use the *sessioninitiation* option with the **schedule** command to control whether the server or client initiates sessions through a firewall. See “Sessioninitiation” on page 344 for more information.

After you start the client scheduler, it continues to run and to start scheduled events until you press Ctrl+C, stop the scheduler process with the UNIX **kill** command, start the workstation again, or turn off the workstation to end it.

Note: You *cannot* enter this command in interactive mode.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

►►—Schedule—┐
└ options ┘

Parameters

options

Table 97. Schedule command: Related options

Option	Where to use	Page
<i>maxcmdretries</i>	Client system options file (dsm.sys) or command line.	288
<i>password</i>	client user options file (dsm.opt)	304
<i>queryschedperiod</i>	Client system options file (dsm.sys) or command line.	322
<i>retryperiod</i>	Client system options file (dsm.sys) or command line.	329
<i>schedlogname</i>	Client system options file (dsm.sys) or command line.	333
<i>schedmode</i>	Client system options file (dsm.sys) or command line.	336

Table 97. Schedule command: Related options (continued)

<i>sessioninitiation</i>	Client system options file (dsm.sys) or command line.	344
<i>tcpclientport</i>	Client system options file (dsm.sys) or command line.	364

Examples

Task For AIX: Start the scheduler at system bootup time by entering this command in the /etc/inittab file. Ensure the *passwordaccess* option is set to *generate*.

Command: tsm::once:/usr/bin/dsmc sched > /dev/null 2>&1 #TSM

Task Interactively start the scheduler and keep it running in the background.

Command: nohup dsmc sched 2> /dev/null &

When you run the **schedule** command, all messages regarding scheduled work are sent to the dsmsched.log file or to the file you specify with the *schedlogname* option in your client system options file (dsm.sys). If you do not specify a directory path with the file name in the *schedlogname* option, the dsmsched.log will reside in the current working directory. file will reside in the default directory:

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 “File system and ACL support” on page 84 for supported file systems.

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 "/dir1/dir2/bogus.txt"
```

still backs up *dir1* and *dir2* 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 UNIX-family platforms. 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 324


Associating a local snapshot with a server file space

Use the *snapshotroot* option with the **selective** command in conjunction with a third-party application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. The *snapshotroot* option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot. See “Snapshotroot” on page 351 for more information.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax



Parameters

options

Table 98. Selective command: Related options

Option	Where to use	Page
<i>changingretries</i>	Client system options file (dsm.sys) or command line.	201
<i>compressalways</i>	Client user options file (dsm.opt) or command line.	210
<i>compression</i>	Client user options file (dsm.opt) or command line.	211
<i>dirsonly</i>	Command line only.	223
<i>filelist</i>	Command line only.	250
<i>filesonly</i>	Command line only.	253
<i>preservelastaccessdate</i>	Client user options file (dsm.opt) or command line.	316
<i>removeoperandlimit</i>	Command line only.	324
<i>snapshotroot</i>	Command line only.	351
<i>subdir</i>	Client user options file (dsm.opt) or command line.	356
<i>tapeprompt</i>	Client user options file (dsm.opt) or command line.	358

filespec

Specifies the path and file name that you want to back up. Use wildcard characters to select a group of files or all the files in a directory. When backing up a file system, specify the file system with a trailing slash; for example: /home/.

Examples

Task Back up the proj file in the /home/devel directory.

Command: selective /home/devel/proj

Task Back up all files in the /home/devel directory whose file names begin with proj.

Command: selective "/home/devel/proj*"

Task Back up all files in the /home/devel directory whose file names begin with proj. Back up the single file named budget in the /user/home directory.

Command: selective "/home/devel/proj*" /user/home/budget

Task Back up the /home file system.

Command: selective /home/ -subdir=yes

Task Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, run a selective backup of the

/usr/dir1/sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name /usr.

Command: dsmc sel "/usr/dir1/sub1/*" -subdir=yes
-snapshotroot=/snapshot/day1

Set Access

The **set access** command gives users at other nodes access to your backup versions, archived copies, or backup images. You can give another user access to a specific file or image, multiple files or images, or all files in a directory. When you give access to another user, that user can restore or retrieve your objects. Specify in the command whether you are giving access to archives or backups.

Note: You cannot give access to both archives and backups using a single command.

Supported Clients

This command is valid for all UNIX and Linux clients.

Syntax

```
▶▶—SET Access— [ Archive Backup ] [ filespec image-fs ] node [ user ]▶▶
```

Parameters

Archive

Permits access to archived files or images.

Backup

Permits access to backup versions of files or images.

filespec

Specifies the path, file, image, or directory to which you are giving access to another node or user. Use wildcard characters to specify a group of files or images, or all files in a directory; all objects in a directory branch; or all objects in a file system. Use a single asterisk "*" for the file spec to give access to all files or images owned by you and backed up on the server. When the command **set access backup "*" node** is entered, no check is made with the server; it is assumed you have at least one object backed up.

If you give access to a branch of the current working directory, you only need to specify the branch. If you give access to objects that are not in a branch of the current working directory, you must specify the complete path. The file spec to which you gave access must have at least one backup version or archive copy object (file or directory) on the server.

To specify all files in a named directory, enter `/home/mine/proj1/*` on the command line.

To give access to all objects *below* a certain level, use an asterisk, directory delimiter, and an asterisk at the end of your file spec. For example, to give access to all objects below `home/test`, use file spec `home/test/*/*`.

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: `/home/sub1/subsub1`.
- The `/home` directory contains the `h1.txt` and `h2.txt` files.
- The `/home/sub1` directory contains file `s1.htm`.
- The `/home/sub1/sub2` directory contains the `ss1.cpp` file.

To allow access to all files in the `/home/sub1/sub2` directory, enter:

```
set access backup /home/sub1/sub2/* * *
```

To allow access to only those files in the `/home` directory, enter:

```
set access backup /home/* * *
```

To allow access to all files in all directories *in* and *below* the `/home` directory, enter:

```
set access backup /home/* * *  
set access backup /home/*/* * *
```

image-fs

The name of the image file system to be shared. This can be specified as an asterisk (*) to allow access to all images owned by the user granting access.

node

Specifies the client node of the user to whom you are giving access. Use wildcards to give access to more than one node with similar node names. Use an asterisk (*) to give access to all nodes.

user

This is an optional parameter that restricts access to the named user at the specified node.

Examples

Task Give the user at `node_2` authority to restore the budget file from the `/home/user` directory.

Command: `set access backup /home/user/budget node_2`

Task Give `node_3` authority to retrieve all files in the `/home/devel/proja` directory whose file names end with `.c`.

Command: `set access archive "/home/devel/proja/*.c" node_3`

Task Give `node_3` the authority to retrieve all files in the `/home/devel/proja` directory.

Command: `set ac archive /home/devel/proja/ node_3`

Task Give all nodes whose names end with `bldgb` the authority to restore all backup versions from directories with a file space name of `project`.

Command: `set ac b "{project}/*" "*bldgb"`

Task Give user `serena` at `node_5` authority to restore all images of the file space mounted on directory `/home/devel/proja`.

Command: `set acc backup "home/devel/proja/*/*" node_5 serena`

Set Event

The **set event** command allows you to specify the circumstances for when archived data is deleted. You can use the **set event** command to:

- Prevent the deletion of data at the end of its assigned retention period (Deletion hold)
- Allow the expiration to take place, as defined by the archive copy group (Release a deletion hold)
- Start the expiration clock running when a particular event occurs (Notify the server that an event has occurred)

Objects affected can be specified with a standard Tivoli Storage Manager filespec (including wildcards), a list of files whose names are in the file specified using the *filelist* option, or a group of archived files with the description specified with the *description* option.

Note: When only a <filespec> is used, *all* archived copies of files or folders matching the filespec are affected. If you only want to affect certain versions of a file, use the *-pick* option and select from the displayed list.

Interaction with down-level servers:

If the **set event** command is issued when the client is connected to a server that does not support event-based policy (previous to Tivoli Storage Manager 5.2.2), the command will be rejected with an error message indicating that the current server does not support event-based policy.

Supported Clients

This command is valid for all UNIX and Linux clients.

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 214)
- *Numberformat* (See “Numberformat” on page 301)
- *Noprompt* (See “Noprompt” on page 300)
- *Subdir* (See “Subdir” on page 356)
- *Timeformat* (See “Timeformat” on page 369)

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

Authorized User

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

Supported Clients

This command is valid for all UNIX and Linux clients.

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

Set Waspassword

Root User

If WAS security is enabled, user name and password validation for Data Protection for WebSphere Application Server is required. If you do not set the WAS password for the security, the backup will failover to an offline backup.

Recommendation: Set the WAS security password to perform consistent backups.

Use the **set waspassword** command to set the user name and password for each installation of WAS on your machine. You only need to perform this task once, unless you change your WAS user name or password. You can only perform this task on the Tivoli Storage Manager command line.

To determine if WAS security is enabled, enter the following command:

```
dsmc query was -wast=local
```

Tivoli Storage Manager displays the WAS security status under the **Sec** heading.

Supported Clients

This command is valid for AIX, Solaris, and Linux x86/x86_64 clients clients.

Syntax

►►—SET WASPassword—┐
 └ WASNode— WASType— WASUser—┘

Parameters

WASNode

Specifies the node name on which each installation of WAS is installed. This parameter is required. If you do not specify a value for this parameter, Tivoli Storage Manager prompts you. See “Wasnode” on page 389 for more information.

WASType

Specifies the WAS Network Deployment Manager (ND) or Application Server (APP). This parameter is required. If you do not specify a value for this parameter, ND is the default value. See “Wastype” on page 390 for more information.

WASUser

Specifies WAS user name. This parameter is required. If you do not specify a value for this parameter, Tivoli Storage Manager prompts you. See “Wasuser” on page 391 for more information.

Examples

Task: Set the WebSphere user name and password for each installation of WAS on your machine.

Command: `dsmc set waspassword -wasnode=wasnode -wastype=app -wasuser=ed`

Tivoli Storage Manager prompts for the WebSphere password, as follows:

Please enter the WebSphere password:

If you do not specify the *wasnode* or *wasuser* options, Tivoli Storage Manager prompts for this information as follows:

Please enter WebSphere node name:

Please enter the WebSphere user name:

Appendix A. Using the Tivoli Storage Manager central 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 503 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, imagebackup, imagerestore, command, and macro. The scheduled action that is most frequently used is incremental with the *objects* parameter left undefined. With this setting, the Tivoli Storage Manager client performs a domain incremental backup of all file systems defined by the client domain option. A schedule definition using the *command* action allows an operating system command or shell script to be executed. When automating tasks for *Tivoli Storage Manager for Data Protection* clients, you must use *command* action schedule definitions which invoke the command line utilities for the "Tivoli Storage Manager for" application.

The schedule *startup window* indicates the acceptable time period for a scheduled event to start. The startup window is defined by these schedule definition parameters: *startdate*, *starttime*, *durunits*, and *duration*. The *startdate* and *starttime* options define the beginning of the startup window for the very first scheduled event. The beginning of the startup windows for subsequent scheduled events 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 double quotes around each file specification that contains blanks, then single quotes around all of the specifications. Examples:

```
objects="/home/proj1/Some file.doc"
objects="/home/proj1/Some file.doc" "/home/Another file.txt"
/home/noblanks.txt'
objects="/home/My Directory With Blank Spaces/'
```

This will ensure that `/home/proj1/Some file.doc` is treated as a single file name, as opposed to two separate files (`/home/proj1/Some`, and `file.doc`).

You can also refer to the *objects* parameter information for the **define schedule** and **update schedule** commands in the appropriate *IBM Tivoli Storage Manager Administrator's Reference*

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>managedservices</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>queryschedperiod</i>	X		set queryschedperiod 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 *widget* command might exit with return code 3, if *widget* ran successfully. Therefore, it is possible that the *preschedulecmd* or *postschedulecmd* command might end with a nonzero return code and be successful. To prevent Tivoli Storage Manager from treating such commands as failed, you should wrap these commands in a script, and code the script so that it interprets the command return codes correctly. The script should exit with return code 0 if the command was successful; otherwise it should exit with a nonzero return code. The logic for a script running *widget* might look like this:

```
run 'widget'  
  if lastcc == 3  
    exit 0  
  else  
    exit 1
```

See the following references for more information:

- “*Postschedulecmd/Postnschedulecmd*” on page 311
- “*Preschedulecmd/Prenschedulecmd*” on page 314
- “Client return codes” on page 155

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 (CAD). The CAD 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 CAD to manage the scheduler is the preferred method. These methods are compared as follows:

CAD-managed Services

- Defined using the *manageservices schedule* option and started with CAD services (dsmcad).
- The CAD 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 CAD 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 CAD-managed and traditional methods of running the Tivoli Storage Manager client scheduler. When running a CAD-managed schedule, only the CAD 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 CAD to run as a background system task which starts automatically when your system is started. To configure the CAD to manage scheduled backups, you must set the *manageservices* option to manage the scheduler, or both the scheduler and web client. The method for setting up the CAD 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 sess`. Note that for testing purposes, you can always start the scheduler in the foreground by running `dsmc sched` from a command prompt (without a 'managedservices' stanza set).

To start the scheduler automatically at boot time, use either of the following methods:

CAD-managed

1. In your `dsm.sys` file, set the *managedservices* option to *schedule* or *schedule webclient*.
2. Add the following entry into the system startup file (`/etc/inittab` for most platforms):

```
tsm::once:/usr/bin/dsmcad > /dev/null 2>&1 # TSM Client Acceptor Daemon
```
3. In your `dsm.sys` file, set the *passwordaccess* option to *generate*.
4. Run a command like `dsmc query sess` to store the node password.

Traditional

1. In your client options file (`dsm.opt`), look for an entry called `MANAGEDSERVICES`. If it exists, and `SCHEDULE` is one of the parameters, delete the word `SCHEDULE`.
2. Add the following entry into the system startup file (`/etc/inittab` for most platforms):

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

For z/OS UNIX:

- a. Create a shell script called `/tivoli/tsm/client/ba/bin/rundsmc` which contains the following entries:

```
cd /usr/lpp/Tivoli/tsm/client/ba/bin
sleep 60
./dsmc schedule
```

This prevents the creation of two jobs with the same name and enables automatic shutdown. You might need to customize the time for your system.

- b. Add the following entries in the `/etc/rc` file to set environment variables to retrieve the *servername* and *nodename* options from `dsm.sys` and to start the client scheduler, as follows:

```
# Set environment variables to retrieve the servername and
# nodename options from dsm.sys.
export DSM_DIR=/tivoli/tsm/client/ba/bin
export DSM_CONFIG=/tivoli/tsm/client/ba/bin/dsm.opt
# Start the TSM Client scheduler and redirect outputs to
# schedule.out instead of the /etc/log file.
_BPX_JOBNAME='ADSMCLNT' /tivoli/tsm/client/ba/bin/rundsmc
1>/tivoli/tsm/client/ba/bin/schedule.out 2>&1 &
```

Note: Enter the `_BPX_JOBNAME` entry on a single line in the `/etc/rc` file.

3. In your `dsm.sys` file, set the *passwordaccess* option to *generate*.
4. Run a command like `dsmc query sess` to store the node password.
5. To start the client scheduler on your client node and connect to the server schedule, enter the following command:

```
dsmc schedule
```


If the current directory is not in your PATH environment variable, enter the following command:

```
./dsmc schedule
```

When you start the client scheduler, it runs continuously until you close the window, end the process, or log off your system.

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

```
nohup dsmc schedule 2> /dev/null &
```

If you are using the CAD to manage the scheduler, you must set the recovery properties for the "TSM Client Acceptor" service, but leave the settings for the "TSM Client Scheduler" to take no action for first, second, and subsequent failures.

Managing multiple schedule requirements on one machine

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 machine. 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 *tcplientport* 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, via Tivoli Storage Manager client option file or Tivoli Storage Manager server override options.

The disadvantages of using multiple schedule processes:

- A unique file space for each node name on the Tivoli Storage Manager server is created.
- When restoring the data, you must use the same node name associated with the backup.

Multiple schedule processes can run on UNIX and Linux platforms with either the CAD 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.

You must start each `dsmc sched` command or instance with the `-servername` option to reference its unique stanza name in `dsm.sys`. For `dsmcad`, it is necessary to define the environment variable `DSM_CONFIG` for each instance of `dsmcad` to reference its unique option file.

The following is an example configuration of two schedule processes managed by the CAD in the `dsm.sys` file. Note that you must use full paths for the log file names to avoid the files being written in the root directory):

```

servername tsm1_sched1
nodename      aixsvt01_sched1
tcpserv       firebat
tcpclientport 1507
passwordaccess generate
domain        /svt1
schedmode     prompted
schedlogname  /tsm/dsmsched1.log
errorlogname  /tsm/dsmerror1.log
managedservices schedule

servername tsm1_sched2
nodename      aixsvt01_sched2
tcpserv       firebat
tcpclientport 1508
passwordaccess generate
domain        /svt1
schedmode     prompted
schedlogname  /tsm/dsmsched2.log
errorlogname  /tsm/dsmerror2.log
managedservices schedule

```

Contents of `/test/dsm.opt1`:

```
servername tsm1_sched1
```

Contents of `/test/dsm.opt2`:

```
servername tsm1_sched2
```

Open two shell command windows:

- In shell command window 1, enter:


```
export DSM_CONFIG=/test/dsm.opt1
sudo dsmcad
```
- In shell command window 2, enter:


```
export DSM_CONFIG=/test/dsm.opt2
sudo dsmcad
```

Note: You should enter these commands into a shell script if you intend to have the `dsmcad` processes started directly from `/etc/inittab` so that the proper `DSM_CONFIG` variable can be set prior to launching `dsmcad`.

Restarting the scheduler process on a remote machine

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 machine.

You can create a shell script to search for and kill running Tivoli Storage Manager scheduler or CAD processes, and restart the processes. You can use software products, such as Symark Power Broker, to allow Tivoli Storage Manager administrators limited access to UNIX servers to manage the scheduler processes and copy off the Tivoli Storage Manager schedule log file. The following shell script is an example of how to recycle the Tivoli Storage Manager scheduler process:

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

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

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

# Kill the scheduler
kill -9 $PID

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

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

Using the scheduler on clustered systems

Using Tivoli Storage Manager client in cluster environment requires additional configuration steps and can be easily done. There are some concepts to keep in mind when configuring Tivoli Storage Manager client in cluster environment:

- The Tivoli Storage Manager client is not a cluster aware application and processes data from shared volumes in the same way as from local volumes.
- If the Tivoli Storage Manager client backs up the same shared volume from two different cluster nodes using same Tivoli Storage Manager node, two different file spaces are created on the Tivoli Storage Manager server containing the same data. To avoid this, you must set the *clusternode* option to *yes* in the Tivoli Storage Manager client options file (dsm.opt). This setting changes the file space naming convention to CLUSTER_NAME\VOLUME_NAME and allows the Tivoli Storage Manager client to successfully manage backups of shared volumes taken from different cluster nodes. At the same time, this change in file space naming will result in mixing data from physically different local volumes under one file space on the Tivoli Storage Manager server.

It is recommended that you separate the backup of shared and local volumes. If the cluster has several shared disk resources which can migrate from one cluster node to another, it is recommended to run a separate instance of the Tivoli Storage Manager client for each disk resource. This allows Tivoli Storage Manager client instances to migrate together with disk resources, in case of failover or failback. See cluster software documentation for information about how to create a program resource. In case of a failover or failback, the disk resources must be up and running before the Tivoli Storage Manager client starts.

The Tivoli Storage Manager client will prompt for the password if it is not readable from a password file. The Tivoli Storage Manager node password must be available for all Tivoli Storage Manager client cluster nodes where the Tivoli Storage Manager client resource can migrate to. For example, if a cluster has three nodes A, B and C and the Tivoli Storage Manager client resource can execute on nodes A and C, passwords must be readable on nodes A and C but not on B. See “Changing your password” on page 75 for more information about how to generate Tivoli Storage Manager client passwords.

For information on how to configure a Tivoli Storage Manager server to manage a cluster configured client, see Appendix B, “Configuring the backup-archive client in an HACMP takeover environment,” on page 511.

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 machines. 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 * * begind=today-1 begint=00:00:00 endd=today-1 endt=23:59:59
```

You can limit query results to exception cases with:

```
query event * * begind=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 99 summarizes the meaning of the event status codes which are likely to exist for a scheduled event that has already taken place:

Table 99. 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.

Table 99. Scheduled client event status codes (continued)

Status	Meaning
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. See “Restarting the scheduler process on a remote machine” on page 507 for tips on how to remotely recycle the hanging scheduler process. 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 begind=08/21/2003 begint=08:25:00
query actlog begind=08/21/2003 begint=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 machine 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:

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 182.

Appendix B. Configuring the backup-archive client in an HACMP takeover environment

High Availability Cluster Multi Processing (HACMP) allows scheduled Tivoli Storage Manager client operations to continue processing during a failover situation.

Tivoli Storage Manager in an HACMP environment can back up anything that Tivoli Storage Manager can normally back up. However, be careful when backing up non-clustered resources due to the "after failover" effects.

Local resources should never be backed up or archived from clustered Tivoli Storage Manager nodes. Local Tivoli Storage Manager nodes should be used for local resources.

For example, a scheduled incremental backup of a clustered volume is running on **machine-a**. A situation causes the client acceptor daemon (CAD) to failover to **machine-b**. **machine-b** then reconnects to the server. If the reconnection occurs within the start window for that event, the scheduled command is restarted. This scheduled incremental backup will reexamine files sent to the server before the failover. The backup will then "catch up" to where it terminated before the failover situation.

If a failover occurs during a user initiated client session, the Tivoli Storage Manager CAD starts on the node that is handling the takeover. This allows it to process scheduled events and provide Web client access. You can install Tivoli Storage Manager locally on each node of an HACMP environment. You can also install and configure the Tivoli Storage Manager Scheduler Service for each cluster node to manage all local disks and each cluster group containing physical disk resources.

Note: If you use the *httpport* option to specify specific ports for Web client access, this will help you remember which port to point to. Without it, the CAD will find the first available port, starting with 1581, and will likely change ports when failover or fallback occurs. Choosing a port value, such as 1585 can save you the inconvenience of trying to determine to which port the CAD might have changed.

The *clusternode* option determines if you want the Tivoli Storage Manager client to back up cluster resources and participate in cluster failover for high availability. See "Clusternode" on page 203 for more information.

The following software is required:

- HACMP 5.1 and 5.2
- AIX 5L (5.1 and 5.2)

The HACMP Cluster Information Daemon must also be running.

Note: See the README for the latest fix levels required for HACMP and AIX.

Installing the backup-archive client

Install the Tivoli Storage Manager Backup-Archive client software on a local disk on each node in the cluster you want to participate in an HACMP takeover. The following client configuration files must be stored locally:

- The client executables and related files should reside in the same location on each node in the cluster.
- The API executable and configuration files should reside in the default API installation directory (/usr/tivoli/tsm/client/api/bin)
- The system options file (dsm.sys) should reside in the default client installation directory (/usr/tivoli/tsm/client/ba/bin)

The following client configuration files must be stored externally in a shared disk subsystem so they can be defined as a cluster resource and be available to the takeover node during a failover. Each resource group must have the following configuration:

- The client option file (dsm.opt), include-exclude file, and password file must be placed in a directory on the shared disk.
- The log files must be placed on the shared disk volumes to maintain a single continuous log file.

Configuring the backup-archive client to process local nodes

You can edit your dsm.opt file on each local node to process local disk drives using the following options:

clusternode

Do not specify this option when processing local drives. See “Clusternode” on page 203 for more information.

nodename

If no value is specified, Tivoli Storage Manager uses the local machine name. Do not use the same local nodename that is being used for cluster processing. See “Nodename” on page 297 for more information.

domain

Recommendation: Use DOMAIN statement(s) to restrict local nodes to local filespaces and cluster nodes to cluster spaces. This will help to ensure that local filespaces from one system are not available to cluster nodes after failover to another system. See “Domain” on page 227 for more information on using Domain statements.

You can also configure the Tivoli Storage Manager Backup-Archive Scheduler Service to back up the local cluster nodes.

Configuring Tivoli Storage Manager backup-archive client to process cluster disk resources

Ensure that Tivoli Storage Manager manages each cluster group that contains physical disk resources as a unique node. This ensures that Tivoli Storage Manager correctly manages all disk resources, regardless of which cluster node owns the resource at the time of back up.

Step 1: Register the client to a server

A Tivoli Storage Manager client in an HACMP cluster must be registered to a Tivoli Storage Manager server with an assigned node name. Consider the following conditions when registering your node name:

- If local volumes that are *not* defined as cluster resources will be backed up, separate node names (and separate client instances) must be used for both non-clustered and clustered volumes.
- The node name used to back up clustered volumes defaults to the cluster name, not the host name.

Recommendation: choose a node name related to the cluster resource group to be managed by that node. For example, the HACMP node name should be the same as the Tivoli Storage Manager node name.

- If multiple resource groups are defined in the HACMP environment to failover independently, then separate node names must be defined per resource group.
- **Recommendations:** If you don't specify the file systems in a backup or archive command line client or in a defined client schedule for HACMP managed clients, define the Tivoli Storage Manager cluster resources in the domain statement. If you do specify the file system in a backup or archive command line, define the Tivoli Storage Manager cluster resources in the include and exclude statements, in order to restrict clustered operations to cluster resources and non clustered operations to non clustered resources.

Step 2: Configure the client system options file

Each node in the HACMP cluster that runs the Tivoli Storage Manager client must have the following settings defined in each respective `dsm.sys` file:

- Separate server stanzas to back up non-clustered volumes
- Separate server stanzas for each cluster resource group to be backed up

The server stanzas defined to back up non-clustered volumes must have the following special characteristics:

- The value of the *tcpclientaddress* option must be the *service* IP address. This is the IP address used for primary traffic to and from the node.
- If the client will back up and restore non-clustered volumes without being connected to the HACMP cluster, the value of the *tcpclientaddress* option must be the *boot* IP address. This is the IP address used to start the machine (node) before it rejoins the HACMP cluster.

The server stanzas defined to back up clustered volumes must have the following special characteristics:

- *clusternode yes*
- The *nodename* value must be related to the resource group. If *nodename* is not specified, the cluster name is used.
- The *tcpclientaddress* option must refer to the service IP address of the HACMP node.
- The *passworddir* option must point to a directory on the shared volumes that are part of the cluster resource group.
- The *errorlogname* and *schedlogname* options must point to files on the shared volumes that are part of the cluster resource group.
- All *inlexcl* statements must point to files on the shared volumes that are part of the cluster resource group.

- Set the *managedservices* statement to indicate that the scheduler (or Web client) should be managed by the client acceptor daemon.

Other options can be set as desired.

Step 3: Configure the client user options file

The client user options file (*dsm.opt*) for the Tivoli Storage Manager client that will manage your clustered file spaces must reside on the shared volumes in the cluster resource group. Define the *DSM_CONFIG* environment variable to point to this *dsm.opt* file. Make sure the *dsm.opt* file contains the following settings:

- The value of the *servername* option must be the server stanza in the *dsm.sys* file which defines parameters for backing up clustered volumes. The *dsm.sys* file might reside on shared space.
- If the *dsm.sys* file resides on a local disk, each node on the cluster must have a matching stanza.
- Define clustered filespaces to be backed up with the *domain* option.
- Other options can be set as desired.

Defining the client as an HACMP application

The Tivoli Storage Manager client must be defined as an application to HACMP to participate in failover processing. See *High Availability Cluster Multi-Processing for AIX - Planning and Installation Guide, SC23-4861*, for detailed instructions on how to perform this procedure.

The Storage Manager client must be in a resource group with a *cascading* or *rotating* takeover relationship. The client does not support a concurrent access resource group. See *High Availability Cluster Multi-Processing for AIX - Planning and Installation Guide, SC23-4861*, for additional information regarding HACMP topology and strategy.

Note: For installation instructions for HACMP versions later than 5.1, go to: http://www.ibm.com/servers/eserver/pseries/library/hacmp_docs.html.

Creating an HACMP resource group to add a client

You must first create an HACMP resource group so you can add the client to it. The following is a summary of this procedure for HACMP 5.1, using IP address takeover (IPAT):

1. Start HACMP for AIX system management with the following command:

```
smit hacmp
```
2. Select **Cluster Configuration, Cluster Resources, Define Resource Groups, and Add a Resource Group**. The Add a Resource Group window is displayed.
3. On the Add a Resource Group window, enter the following field values:

Resource Group Name

Enter an ASCII text string that identifies the resource group. The resource group name can include alphabetic and numeric characters and underscores. Use no more than 31 characters.

Node Relationship

Select **Cascading**.

Participating Node Names/Default Node Priority

Select the node names that are participating in the resource group. Add the nodes in order of priority. The node owner of the resource group should be the first node listed.

4. Click **OK**.
5. Press F10 to exit `smit` and return to the command line. Press F3 to perform other configuration tasks.

The Storage Manager client must be in a resource group with a *cascading* or *rotating* takeover relationship. The client does not support a concurrent access resource group. See *HACMP for AIX 5.1 Planning Guide, SC23-4277*, for additional information regarding HACMP topology and strategy.

Adding the client to an HACMP resource group

The Tivoli Storage Manager client must be defined to a cluster resource group. See *HACMP for AIX 5.1 Installation Guide, SC23-4278*, for detailed instructions on how to perform this procedure. Following is a summary of how to define resources as part of a resource group:

1. Start HACMP for AIX system management with the following command:
`smit hacmp`
2. Select **Cluster Configuration, Cluster Resources, and Change/Show Resources/Attributes for a Resource Group**. Press Enter.
3. Select the desired resource group.
4. Press Enter. The **Configure a Resource Group** screen appears.
5. Enter values that define all the resources you want to add to this resource group.
6. Synchronize cluster resources after entering field values in Step 5. Do this by selecting **Cluster Configuration, Cluster Resources, and Synchronize Cluster Resources**.
7. Press F10 to exit `smit` and return to the command line. Press F3 to perform other configuration tasks.

The Tivoli Storage Manager client must be added to the resource group that contains the file systems to be backed up. These file systems must also be the same file systems specified by the *domain* option in the `dsm.opt` file defined for this client instance.

Both JFS2 and NFS file systems can be defined as cluster resources. NFS supports only 2 node clusters in a cascading takeover relationship.

Using journal-based backup and HACMP

You can set the proper environment for HACMP by configuring a journal service on each node in the cluster and then setting some of the following options. Refer to Appendix C, "Journal daemon configuration," on page 517 for more information.

- **JournalDir** - The journal directory for shared resource journals must be accessible by all nodes in the cluster which might own the shared resource. The default location for the journal directory is in the file system being journaled. This is usable for a cluster environment but might be inefficient for local resources. It might, therefore, be appropriate to use the `JournalDir` setting to define different journal directory locations for shared and local resources. Note that the directory specified can be an absolute path (that is, it starts with a

directory delimiter) or a relative path name. In the later case, the location of the journal directory is formed by using the name of the file system being journaled and appending the relative path name.

- ***PreserveDbOnExit*** - The ***PreserveDbOnExit*** setting can be used to allow journals for shared resources to remain valid when a resource moves to a different node in the cluster.
- ***DeferFsMonStart*** - The ***DeferFsMonStart*** should be set to 1 to cause journal-based backup to periodically (***DeferRetryInterval***) check for a file system coming online or going offline.
- ***DeferRetryInterval*** - This controls how often the journal checks for a change in file system status.
- ***logFsErrors*** - This is used to eliminate excessive "File System unavailable" messages from being written to the log.

The previously listed options can be set for all file systems being journaled. Then, if desired, you can override any given file system with a file system override stanza.

When the journal daemon is set up to monitor file systems as they go online and offline, this is done in a periodic loop. The following actions occur during resource failover and fallback conditions:

- nodeA and nodeB share a resource and nodeA has a failure and, therefore, a failover occurs
- nodeB will detect the shared resource and start monitoring
- When nodeA takes the resource back (fallback), if nodeB has an open file descriptor to the shared journal database, it will receive a SIGKILL interrupt. This interrupt cannot be caught or ignored.
- nodeB must be informed of the resource fallback before the event, by removing the shared file system names from the journal list of file systems to monitor contained in the options file - `tsmjbbd.ini`. This can be done manually or it can be automated by using HACMP start and stop scripts.

Two sample perl scripts are provided with the journal: `jbbadd.pl` and `jbbdelete.pl`. The invocation of these scripts can be added to `StartClusterTsmClient.sh` and `StopTsmClient.sh`. The arguments for these scripts are just the file system names of the shared resources.

Appendix C. Journal daemon configuration

Journal-Based Backup is enabled by installing and configuring the Tivoli Storage Manager journal daemon. The Tivoli Storage Manager journal daemon can be configured by editing the Journal daemon configuration sample file, `tsmjbbd.ini.smp`, and saving it as `tsmjbbd.ini`. Both files should be in the default install directory.

Notes:

1. For customers already using journal-based backup, the first incremental backup after installing Tivoli Storage Manager Version 5.4 will result in a non-journal based backup, even if you are using the `PreserveDbOnExit` option in your `tsmjbbd.ini` file. This is because Tivoli Storage Manager Version 5.4 now uses a new database for better reliability and efficiency and the new database must go through normal validation before it can support journal-based backups.
2. Network and removeable file systems are not supported.

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

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

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

The following is the syntax for stanza and stanza settings:

Syntax for stanzas:

[StanzaName]

Syntax for stanza settings:

stanzaSetting=value

Notes:

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

JournalSettings stanza

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

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

Syntax for *JournalSettings* stanza:

[JournalSettings]

Syntax for stanza settings:
JournalSettings=value

You can specify the following *JournalSettings* values:

| *ErrorLog*

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

| `ErrorLog=/logs/jbberror.log`

| *JournalDir*

| Directory where journal database files are stored and written.

| If the path given is an absolute (for example, it begins with a `dir`
| delimiter) pathname, this is the directory used. If the path given is a relative
| directory name, then this path is appended to each file system name and
| the resulting path name is used.

| The default is a directory named `.tSm_JoUrNaL` (used within each file
| system being journaled).

| The advantage of having the journal database on the file system being
| monitored is that the database stays with the file system. The disadvantage
| is that the updates to the database must be processed and discarded. That
| is, directing the database to a non-journaled file system is more efficient
| and is strongly recommended, unless this file system is shared in a cluster
| environment.

| This setting applies to all journaled file systems but may be overridden via
| an override stanza for each journal file system.

JournalExcludeList stanza

This list of exclude statements filters changes from being recorded in the journal database. Changes to objects which match statements in this stanza are ignored and are not recorded in the journal database.

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

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

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

The following pattern matching meta characters are supported:

- % Matches exactly one character.
- * Matches zero or more characters.

%EnvVar%
Expands environment variable.

The following is an exclude statement syntax example:

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

JournaledFileSystemSettings stanza

Settings under this stanza apply to each specified journaled file system unless they are overridden for individual file systems in an override stanza. See “Overriding stanzas” on page 521 for more information.

The following is the syntax for the `JournaledFileSystemSettings` stanza:

Syntax for `JournaledFileSystemSettings` stanza:

```
[JournaledFileSystemSettings]
```

Syntax for stanza settings:

```
JournaledFileSystemSetting=value
```

You can specify the following `JournaledFileSystemSettings` values:

JournaledFileSystems

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

```
JournaledFileSystems=/home /other
```

Note: The journal selects object names based strictly on a string match.

The implication for the user is that care must be taken when selecting file systems to journal. For example, suppose you have a file system `/jbb` and another file system called `/jbb/mnt1`. If you ask the journal to monitor just `/jbb`, then all the changes for `/jbb/mnt1` will also match this string and be entered in the database. When, however, you do a back up on the client, it parses the name based on file systems, realizes the journal is not monitoring this file system and then tells the journal to remove the `/jbb/mnt1` files from the database. The solution is to either monitor both or use the `JournalExcludeList`. The same is true for the virtual mount point options. It is highly recommended that you are consistent with this list. For example, if you specify `/home/student1` as a virtual mount point in your `dsm.sys` option file and you want to journal `/home`, then you should specify `JournaledFileSystems=/home /home/student1`. In this case, two separate databases are created.

JournalDbSize

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

```
JournalDbSize=0x10000000
```

NotifyBufferSize, DirNotifyBufferSize

Specify change notification buffer sizes for a journaled file system. A large

amount of change activity on a journaled file system might require this to be increased. The default is 0x00020000 (128 k) for files and 0x00010000 (64 k) for directories.

```
NotifyBufferSize=0x00200000
```

PreserveDbOnExit setting

This setting allows a journal to remain valid when a journaled file system goes offline and comes back online. This is useful for preserving the journal during system reboots, and resource movement.

This setting allows a journal-based backup to continue processing when the daemon is restarted (or the file system comes back online) without performing a full incremental backup.

Note: Any change activity which occurs while the journal daemon is not running (or the file system is offline) will not be recorded in the journal.

A value of 1 specifies that the journaled file system journal database is not deleted when the journal file system goes offline. The database will also be valid when the journal file system comes back online. This value should be used with caution because any file system change activity which occurs while the journaled file system is offline will not be reflected in the journal database. The default setting of 0 deletes the journaled file system journal database.

Note: The journal will only be preserved when a journaled file system comes offline normally or is brought offline when the resource is no longer available and you specify the `deferFsMonStart` setting. If a file system comes offline due to an error such as a notification buffer overrun, the journal is not preserved.

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

```
preserveDBOnExit=1
```

deferFSMonStart setting

This setting defers an attempt to begin monitoring a file system in the following cases:

- When the specified journaled file system is not valid or available
- The journal directory for the specified journaled file system cannot be accessed or created

Resources are checked at the interval you specify using the `deferRetryInterval` setting.

A value of 1 indicates that the setting is on. A value of 0 indicates that the setting is off. The default value is off (set to 0) .

deferRetryInterval setting

This setting specifies the value in seconds that a deferred file systems with the `deferRetryInterval` setting enabled are checked for availability and brought online. The default value is 5 seconds.

logFSerrors setting

A value of 1 indicates that all errors encountered accessing a journaled file system or journal directory should be logged. A value of zero indicates that logging of errors encountered while checking deferred file systems and journal directories will be suppressed. This is usually used in conjunction with the `deferFSMonStart` setting to eliminate excessive File System

| unavailable messages from being written to the logs when bringing a
| journaled file system on-line is deferred. The default value is 1 (log all
| errors).

Overriding stanzas

Any setting in the *JournaledFileSystemSettings* stanza, except for the buffer sizes, can be overridden for a particular journaled file system by creating an override stanza.

HookFileName

In order for the journal to begin monitoring a file system, it must know the name of an existing file in that file system. This setting specifies an existing file. Access to this file is then used as a test of whether or not this file system is online. (The system definition of mounted cannot be used because we allow the use of virtual mount points in the backup-archive client. This means that the Tivoli Storage Manager system can treat a directory as a (virtual) file system).

Therefore, if this file system can be mounted and unmounted, a *HookFileName* needs to be provided.

If a *HookFileName* is not entered, the journal daemon attempts to create a temporary file in the highest directory, use it to begin monitoring, and then delete it.

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

Syntax for JournaledFileSystemSettings stanza:

[JournaledFileSystemSettings.fs]

Syntax for stanza settings:

JournaledFileSystemSetting=override value

| For example, the override stanza name for /home would be:

| JournaledFileSystemSettings./home
| HookFileName=/home/doNotDeleteThisFile

Notices

This information was developed for products and services offered in the U.S.A. IBM 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 be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right 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 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 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 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 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 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:

List of trademarks

AIX	RACF
Application System/400	Redbooks
AS/400	RISC System/6000
Client Access/400	RS/6000
DB2	Scalable POWERparallel Systems
FlashCopy	SP2
i5/OS	S/390
IBM	System/390
IBMLink	Tivoli Enterprise Console
Lotus	TME
Lotus Approach	TotalStorage
Magstar	VisualAge
Operating System/400	WebSphere
OS/390	xSeries
OS/400	z/OS
PowerPC	zSeries

Apple is a registered trademark of Apple Computer, Inc.

EMC and Celerra are registered trademarks of EMC Corporation.

HP-UX is a registered trademark of Hewlett-Packard Company.

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

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

Microsoft, Windows, Windows XP, Windows Server 2003, Windows Vista and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

NetApp is a registered trademark of Network Appliance, Inc.

NetWare is a registered trademark of Novell, Inc.

Solaris is a registered trademark of Sun Microsystems, Inc.

Tivoli is a trademark of Tivoli Systems Inc. in the United States, or other countries, or both.

Unicode is a trademark of Unicode, Inc.

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

VxFS is a registered trademark of VERITAS Software Corporation.

Other company, product and service names might 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*
```

automounted file system (AutoFS). A file system managed by an automounter daemon. The automounter daemon monitors a specified directory path and automatically mounts the file system to access data.

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 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. A Web client process that is an HTTP daemon that serves the Web client Java applet to the Web browsers. The program that starts the client acceptor process is called `dsmcad`. On UNIX and Linux, the client acceptor is run as a daemon.

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 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 **backuppools** 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.

dsm.opt file. See *options file*. See also *client user options file*. Also called client user options file.

dsm.sys file. See *client system options file*.

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.

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

group backup. Back up of a group containing a list of files from one or more file space origins.

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 (10^9) 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.

GPFS node set. A set of AIX SP nodes that can mount a defined group of GPFS file systems.

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

image. A full file system or raw logical volume backup as a single object.

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. The include-exclude list contain entries from any or all of the following sources: the client options file (Windows), the client system options file (UNIX and Linux), separate include-exclude files, or the Tivoli Storage Manager server. 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 file systems, directories, or files. These file systems, 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. Hardlinked 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*.

J

journal-based backup. Journal-based backup is a method of backup for 32-bit Windows clients and AIX clients that exploits the file change notification mechanism to improve incremental backup performance by reducing the need to fully scan the file system. By default, the **incremental** command will automatically perform a journal-based backup on selected file systems which are being monitored by the journal engine service. The **incremental** command performs the traditional full incremental backup on any non-journaled file systems.

journal daemon. The journal is a Tivoli Storage Manager program that tracks change activity for files residing in file systems that you specify. Tivoli Storage Manager processes the list of changed files during a journal-based backup.

L

LAN. Local area network.

LAN-free data transfer. The movement of client data between the client and a storage device over a SAN, bypassing the LAN.

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.

logical volume backup. A back up of a file system or logical volume as a single object

Loopback Virtual File System (LOFS). A file system created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

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

NAS node. A type of node that is a NAS file server. The NAS node name uniquely identifies the NAS file server and its data to Tivoli Storage Manager. Through support of Network Data Management Protocol (NDMP), Tivoli Storage Manager can efficiently back up and restore NAS file systems to tape drives or libraries that are locally attached to the NAS file servers.

NDMP. Network Data Management Protocol.

Network Attached Storage (NAS) file server. A network attached storage (NAS) device is a specialized file-serving box whose operating system is streamlined and optimized for file-serving functions. Through support of Network Data Management Protocol (NDMP), Tivoli Storage Manager can efficiently back up and restore NAS file systems to tape drives or libraries that are locally attached to the NAS file servers.

Network Data Management Protocol. Open standard network protocol. Enables efficient back up and restore of Network Attached Storage (NAS) file systems to tape drives or libraries that are locally attached to the NAS file servers.

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.

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*.

options file. A file that contains processing options.

- **dsm.opt**

Identifies the Tivoli Storage Manager server to contact, specifies backup, archive, restore, and retrieve options. Also called the client user options file.

- **dsm.sys**

Contains stanzas describing Tivoli Storage Manager servers to contact for services. These stanzas also specify communication methods, backup and archive options, and scheduling options. Also called the client system options file.

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

raw logical volume. A portion of a physical volume which is comprised of unallocated blocks and has no Journaled File System (JFS) definition. A raw logical volume is read/write accessible only through low level I/O functions.

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*.

root user (UNIX and Linux). The authority level for a root user permits this user to do authorized tasks for Tivoli Storage Manager.

S

SAN. Storage area network.

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 *sclient-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 on 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.

sparse file. A file that is created with a length greater than the data it contains, leaving empty spaces for future addition of data.

special files. Special files define devices for the system or temporary files created by processes. There are three basic types of special files: FIFO (first-in, first-out), block, and character. FIFO files are also called pipes. Pipes are created by one process to temporarily allow communication with another process. These files cease to exist when the first process finishes. Block and character files define devices. Tivoli Storage Manager processes only device and named pipe special files. Socket special files are not processed.

snapshot image backup. During an snapshot image backup, the volume is available to other system applications during the operation.

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 agent. A program that enables Tivoli Storage Manager to back up and restore client data directly to and from SAN-attached storage.

storage area network (SAN). A high-speed communications network optimized for storage.

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.

TCA. Trusted Communications Agent

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, *dsmj*, *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 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.

virtual file space. A representation of a directory on a network-attached storage (NAS) file system as a path to that directory. A virtual file space is used to back up the directory as a file space in Tivoli Storage Manager server storage.

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 via 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

- ? 119
- * 119
- * ? 120

Numerics

- 128-bit AES encryption support 84

A

- absolute mode 162
- access
 - permissions, archive 147
- access control lists
 - backup 91
 - extended permissions, backup 91
 - permissions, backup 91
- accessibility
 - disability xvi
 - keyboard xvi
 - shortcut keys xvi
- ACL
 - file systems supported 84
- active backup versions
 - displaying 76, 121, 444
 - restoring 121
- administrative client
 - allowing secure sessions within a private network 360
- AIX
 - HACMP takeover environment 511
 - installing Tivoli Storage Manager 511
- AIX client
 - client components 3
 - communication methods 4
 - considerations when upgrading from ADSM AIX 3.1 16
 - disk space requirements 4
 - hardware requirements 4
 - installing 14
 - memory requirements 4
 - operating system requirements 4
 - uninstalling 17
- application program interface (API)
 - setting 50
- archive
 - a list of files 145, 250
 - archive symbolic link and data or symbolic link only 196
 - assign description to 219
 - assigning description on command line 145
 - associating local snapshot with server file space 144, 351, 402
 - authorize users on your workstation to request services
 - from server 377
 - binding management class to 165, 195
 - binding management classes to files 145
 - command 402
 - compressing files 211
 - configuring client node proxy support 145
 - copy mode 162
 - delete file spaces 95, 140

- archive (*continued*)
 - delete files after 218
 - delete individual archives from server file space 147, 418
 - display details 221
 - display the last modification date and last access date 442
 - enabling cluster drive management 203
 - encrypting data during 84
 - encryptiontype 239
 - exclude objects from 246
 - files only 253, 378
 - grace period retention 159
 - hard links 148
 - how managed 150, 157
 - improving speed using share memory 170
 - include files for 269
 - information, query 442
 - managing file growth during compression 210
 - maximum file size 86
 - more than one file specification 144
 - only files; not directories 145
 - overriding management class during 164
 - primary tasks 143
 - process directories only; not files 223
 - purpose 80
 - query user access 441
 - regulate resources server and client use during processing 327
 - remove 20-operand limit for Unix-family platforms512-operand limit for the Mac OS X platform 324
 - retrieving using command line 149
 - revoke access 417
 - running 143
 - shared data on multiple clients under a single node name 145, 197
 - specify end date and time to search for 373, 374
 - specifying number of attempts to archive open files 201
 - specifying whether to include subdirectories for 356
 - specifying whether to update last access date 316
 - specifying whether to wait for tape to mount if it is required 358
 - starting a Web client session 72
 - subdirectories 145
 - summary of options 173
 - suppress confirmation prompt before deleting 300
 - symbolic links 147
 - testing include-exclude list 64
 - using commands 144, 147
- archive copy group 158
 - attributes 159
- archmc option 195
- archsymbklnk as file option 196
- asnodename option 197
- Attention
 - changing HFS file paths to UNIX file paths 3
 - DSM_LOG environment variable 333, 487
 - DSM_LOG environment variable name 49
 - incremental-by-date backup dates 274
 - insufficient system resources on Linux 206
 - securing local backup sets on portable media 132
 - sessioninitiation option and serveronly 344

Attention (*continued*)

- Using the asterisk (*) and not the question mark (?) for multibyte code pages 119, 401
- when Tivoli Storage Manager cannot write to the log file 49

authentication

- Tivoli Storage Manager client 69

authorization

- summary of options 184

authorized user

- definition xii, 41
- tasks 41

authorizing

- groups to use Tivoli Storage Manager 260
- Tivoli Storage Manager users 377
- user to restore or retrieve your files 138

automating backup services

- displaying scheduled work 151, 153
- options for 151
- overview 151, 499
- process commands after back up 311
- process commands before back up 314
- starting client scheduler 52

automount option 199

automounted file systems

- backing up 228

automounter

- platform support for 228

B

backup

- access control lists 91
- access permissions 91
- advanced considerations 114
- authorize users on your workstation to request services
 - from server 377
- automounted file systems 228
- beginning with a specific directory 92
- binding management class to files 165
- comparison: incremental, incremental-by-date 90
- comparison: incremental, journal-based, incremental-by-date 90
- compressing files 211
- configuring support for client node proxy backups 97
- copy mode 162
- define virtual mount point for file systems 382
- delete file spaces 95, 140
- delete individual backups from server file space 95
- delete individual files or images from server file space 420
- display details 221
- displaying processing status 112
- enable LAN-free data movement 237
- enabling cluster drive management 203
- enabling communications for LAN-free data movement 276, 279
- encrypting data during 84
- encryptiontype 239
- estimating processing time 93
- exclude objects from 246
- excluding domains 93, 227
- excluding files from backup services 83
- excluding system objects 247
- extended permissions 91
- files only 253
- files, assigning management class 163

backup (*continued*)

- fuzzy 162
- GPFS, multi-node cluster environment 86, 228
- grace period retention 159
- group 404
 - delete 425
 - specify name of 259
 - specify virtual file space name 381
- group, display all members of 347
- hard links 117
- how Tivoli Storage Manager manages 83
- image 102
 - client domain 231
 - revoke access 417
 - specify type 265
 - using command line 108
 - using DSM_DIR to point to plug-in library 48
 - using the GUI 106
 - using with incremental backup 105, 407
 - using with incremental mode 407
- image, considerations 102
- image, volume device type support 103
- image: static, dynamic, snapshot 102
- improving speed using share memory 170
- inactivate a list of 427
- include files for 269
- include-exclude processing for 60
- incremental
 - associating local snapshot with server file space 351, 433
 - client domain 227
 - directories, processing overview 88
 - new and changed files with modification date later than last backup 274
 - process a list of files 250
 - using command line 93
- incremental-by-date 90
 - directories, processing overview 90
 - using command line 94
 - using with image backup 105
- incremental-by-date, using client command line 93
- incremental-by-date, using client Java GUI 92
- incremental, associating local snapshot with server file space 101
- incremental, specify memory-conserving algorithm for 289
- incremental, using client command line 93
- incremental, using client Java GUI 92
- journal-based 89, 431
- LAN-based image 406
- managing 157
- managing file growth during compression 210
- maximum file size 86
- multi-session, send files contiguously to the server 204
- NAS
 - using DSM_DIR to point to plug-in library 48
- NAS file systems 108
- NAS file systems with command line 110
- NAS file systems with web client GUI 109
- NAS image
 - client domain 232
- Network Attached Storage (NAS) file systems 410
- new or changed files 87
- NFS file systems 111
- open files 119
- overview 79, 87
- performing traditional full incremental 89, 299, 431

- backup (*continued*)
 - planning 79
 - pre-backup considerations 81
 - primary tasks 79
 - process directories only; not files 223
 - purpose 80
 - query information 444
 - query user access 441
 - regulate resources server and client use during processing 327
 - remove 20-operand limit for Unix-family platforms 512-operand limit for the Mac OS X platform 324
 - revoke access 417
 - SAN File System FlashCopy images 352
 - saving encryption key password 240
 - selective
 - associating local snapshot with server file space 351, 488
 - back up list of files 94
 - overview 91
 - using command line 94
 - selective, associating local snapshot with server file space 101
 - selective, using client command line 93
 - selective, using client Java GUI 92
 - shared data on multiple clients under a single node name 197
 - skip acl processing 348
 - skip acl update checking 349
 - snapshot image
 - specify percent value of the target volume to create 350
 - snapshot image backup 406
 - sparse files 118
 - specify end date and time to search for 373, 374
 - specifying automounted file systems 199
 - specifying number of attempts to backup open files 201
 - specifying one server session per file specification 204
 - specifying shared memory port for LAN-free data movement 278
 - specifying whether to include subdirectories for 356
 - specifying whether to update last access date 316
 - specifying whether to wait for tape to mount if it is required 358
 - starting a Web client session 72
 - subdirectories 94
 - summary of options 173
 - symbolic links 116
 - testing include-exclude list 64
 - using LAN-free data movement 81
 - using multiple sessions 112
 - WAS
 - Application Server 389, 390, 413
 - assigning management class 413
 - Network Deployment Manager 389, 390, 413
 - set WebSphere node name 389
 - set WebSphere node name, user name, password 496
 - set WebSphere user name 391
 - WebSphere Application Server (WAS)
 - Application Server 112
 - Network Deployment Manager 112
- backup copy group 158
 - attributes 159
- backup group command 404
- backup image command 406
 - supported devices 103

- backup nas command 410
- backup set
 - enabling GUI for local restore of 282
 - specify location for restore or query 283
- backup sets
 - restoring in a SAN environment 470
- backup was command 413
- backup with client node proxy
 - agent node 97
 - target node 97
- Backup with client node proxy
 - overview 97
- backup-archive client
 - installation 13
 - overview 1
- backupset
 - enabling GUI for local restore 130
 - restore 130
- backupsetname option 200
- batch mode 396
 - starting a session 71
- bottom-up processing
 - include-exclude list 65
 - include-exclude options file 65
- Bourne and Korn shell
 - pointing to client user options file 44
- Bourne and Korn shell variables, setting 49

C

- C shell variables
 - setting 50
- cancel process command 415
- cancel restore command 416
- central scheduling
 - summary of options 182
- changingretries option 201
- class option 202
- client
 - registering with server 57
 - setting password 57
 - specify size for the TCP/IP sliding window for client node 368
 - specifying client TCP/IP address other than the one for first server contact 363
 - specifying client TCP/IP port number other than the one for first server contact 364
- client acceptor daemon
 - manage scheduler, Web client, or both 286
- Client Acceptor daemon
 - configuring to manage scheduler 51
- client command options
 - overview 397
- client components
 - AIX client 3
 - HP-UX Itanium 2 client 5
 - HP-UX PA-RISC client 4
 - Linux for zSeries client 9
 - Linux on POWER 8
 - Linux x86/x86_64 client 6
 - Solaris client 10
 - z/OS UNIX System Services client 10
- client node proxy archive
 - overview 145
- client node proxy backup
 - scheduling 97, 98
- client node proxy support 97, 145

- client options
 - display current settings 46, 458
 - exclude
 - exclude.archive 60
 - exclude.attribute.symlink 60
 - exclude.backup 60
 - exclude.compression 60
 - exclude.dir 60
 - exclude.file 60
 - exclude.file.backup 60
 - exclude.fs 60
 - exclude.image 60
 - guitreeviewafterbackup 261
 - order of processing (precedence) 187
 - overriding using command line 187
 - overview 397
 - setting in a file 45
 - using with commands 187
- client options file
 - creating and modifying 43
 - required options for 43
- client options reference 194
- client scheduler
 - displaying scheduled work 151, 153
 - options for 151
 - overview 151, 499
 - starting 52, 486
 - starting automatically 75
- client system options
 - exclude
 - exclude.archive 60
 - exclude.attribute.symlink 60
 - exclude.backup 60
 - exclude.compression 60
 - exclude.dir 60
 - exclude.file 60
 - exclude.file.backup 60
 - exclude.fs 60
- client system options file
 - copying and modifying 43
 - example of 43
 - minimum required statements 43
 - setting options 45
 - specifying include-exclude options 59
- client user options file
 - customizing 45
 - overriding using commands 187
 - setting options 45
- client-server communication
 - establishing 43
 - identify your workstation to the server 297
 - specify maximum disk I/O buffer size client uses when reading or writing files 224
 - specify size for the TCP/IP sliding window for client node 368
 - specifying client TCP/IP address other than the one for first server contact 363
 - specifying client TCP/IP port number other than the one for first server contact 364
 - specifying default server 217
 - specifying method 206
 - specifying name of server to contact for services 342
 - specifying number of kilobytes client buffers before sending transaction to server 375
 - specifying reconnection attempts after failure 208
 - specifying reconnection interval after failure 209
- client-server communication (*continued*)
 - specifying size of internal TCP/IP communication buffer 361
 - specifying TCP/IP address for dsmcad 362
 - specifying TCP/IP address of Tivoli Storage Manager server 367
 - specifying TCP/IP port address on which to establish Shared Memory connection 346
 - specifying whether to send small transactions to server without buffering them first 365
- client-server communicator
 - specifying TCP/IP port address of Tivoli Storage Manager server 366
- closed registration
 - permissions 57
 - using 57
- cluster drives
 - enabling management of 203
- clusternode option 203
- collocatebyfilespec option 204
- command line
 - archiving files 147
 - assigning description to archive 145
 - display current settings for client options 46, 458
 - displaying processing status 112
 - enabling command history capabilities 234
 - ending a session 77
 - entering commands 397
 - general rules when entering options with commands 188
 - NAS file systems backup 110
 - options 188
 - overriding management class during archive 164
 - overview of parameters 398
 - performing image backup 108
 - performing large restore operations 123
 - performing point-in-time restore 127
 - recall commands 399
 - restrictions for NAS file systems 108
 - retrieving archived files 149
 - return codes for operations 155
 - rules for entering commands 401
 - specifying file specification 398
 - specifying options file during session 303
 - starting a session 71
 - using wildcard characters 400
- command parameters
 - overview 398
- command processing, summary of options 184
- command session
 - ending 395
 - starting 395
- commands
 - archive 402
 - backup group 404
 - backup image 406
 - backup nas 410
 - backup was 413
 - batch mode 396
 - cancel process 415
 - cancel restore 416
 - delete access 417
 - delete archive 418
 - delete backup 420
 - delete filespace 423
 - delete group 425
 - entering 397
 - entering on command line 397

- commands (*continued*)
 - expire 427
 - general rules when entering options with 188
 - help 429
 - incremental 430
 - interactive (loop) mode 396
 - loop 435
 - macro 437
 - maximum file specifications permitted 399
 - monitor process 438
 - overview of parameters 398
 - preview 439
 - query access 441
 - query archive 442
 - query backup 444
 - query backupset 447
 - query filespace 449
 - query group 451
 - query image 453
 - query inxcl 455
 - query mgmtclass 456
 - query node 457
 - query options 458
 - query restore 459
 - query schedule 460
 - query session 461
 - query systeminfo 462
 - query was 464
 - recall previous 399
 - restart restore 465
 - restore 466
 - restore backupset 470
 - restore group 474
 - restore image 476
 - restore NAS 479
 - restore was 481
 - retrieve 484
 - rules for entering 401
 - schedule 486
 - scheduled, enabling or disabling 154
 - selective 488
 - set access 491
 - set event 493
 - set password 495
 - set waspassword 496
 - specifying file specification 398
 - using 393
 - using in executables 155
 - using in shell scripts 155
 - using options with 187
 - using wildcard characters 400
- commmethod option 206
- commrestartduration option 208
- commrestartinterval option 209
- communication methods
 - Shared Memory
 - AIX client 4
 - HP-UX client 5
 - HP-UX Itanium 2 client 6
 - Solaris client 12
 - summary 169
 - TCP/IP
 - AIX client 4
 - HP-UX client 5
 - HP-UX Itanium 2 client 6
 - Linux for zSeries client 10
 - Linux on POWER client 9

- communication methods (*continued*)
 - TCP/IP (*continued*)
 - Linux x86/x86_64 client 7
 - Solaris client 12
 - z/OS UNIX System Services client 10
- communications
 - establishing through firewall 54
- compressalways option 210
- compression
 - disabling processing 248, 272
 - enabling processing 248, 272
 - include-exclude statements 248, 272
- compression and encryption processing
 - back up 272
 - exclude from backup 272
 - exclude options 272
- compression option 211
- compression processing
 - exclude from backup 248, 272
 - exclude options 248, 272
 - excluding files from 246
 - include files for 269
- configuring 56
 - optional tasks 41
 - required tasks 41
 - the client scheduler 51
 - the Web client 50
- console option 213
- copy destination attribute 162
- copy frequency attribute 160
- copy group name attribute 160
- copy groups 158
 - archive 158
 - backup 158
- copy mode attribute
 - absolute 162
 - modified 162
- copy serialization attribute 161
- copy type attribute 160
- customer support
 - contacting xvi

D

- data
 - restoring 122
- date format
 - specifying 214
- dateformat option 214
- default
 - management class 158
 - policy domain 157
- default client user options file
 - creating and modifying 44
 - example of 44
- default domain
 - excluding domains from backup 93, 227
- defaultserver option 217
- delete
 - file space 95, 140
 - specifying description for 219
 - specifying NAS or client objects 202
- delete access command 417
- delete archive command 418
- delete backup command 420
- delete group command 425
- deletefiles option 218

- deleting
 - individual archives from server file space 147, 418
 - individual backups from server file space 95, 420
 - description option 219
 - detail option 159, 221
 - diagnostics
 - options 187
 - directories
 - assigning management class for 222
 - excluding 60
 - excluding from backup processing 246
 - incremental backup processing overview 88
 - processing during incremental-by-date 90
 - specifying on command line 398
 - dirmc option 222
 - dirsonly option 223
 - disability
 - accessibility xvi
 - keyboard xvi
 - shortcut keys xvi
 - disaster recovery 140
 - disk recovery 140
 - disk space requirements
 - AIX client 4
 - client 3
 - HP-UX Itanium 2 client 6
 - HP-UX PA-RISC client 5
 - Linux for zSeries client 9
 - Linux on POWER 8
 - Linux x86/x86_64 client 7
 - Solaris client 11
 - z/OS UNIX System Services client 10
 - diskbuffsize option 224
 - diskcachelocation option 225
 - display
 - policy information 159
 - displaying
 - archive information 442
 - online help 76
 - restartable restore sessions 459
 - scheduled events 460
 - session information 461
 - domain
 - back up using the GUI 92
 - backing up automounted file systems 228
 - include for image backup 231
 - include for incremental backup 227
 - include for NAS image backup 232
 - specifying drives in the default 92
 - domain option 227
 - domain.image option 231
 - domain.nas option 232
 - DSM_CONFIG 48
 - adding to .cshrc file 50
 - pointing to client user options file 44, 48
 - using on Solaris 48
 - DSM_DIR
 - adding to .cshrc file 50
 - pointing to dsm.sys file 48
 - pointing to executable files 48
 - pointing to resource files 48
 - set for image or NAS backup or restore 48
 - dsm_log 49
 - DSM_LOG
 - adding to .cshrc file 50
 - set to point to dsmerror.log, dsmwebcl.log, dsmsched.log 48
 - dsm.opt
 - customizing 45
 - setting options in 45
 - dsm.opt file
 - creating 44
 - creating and modifying 43
 - example of 44
 - required options for 43
 - specifying a drive specification using wildcards 62
 - dsm.opt.smp file 44
 - dsm.smp file
 - copying to dsm.opt 43
 - location 43
 - dsm.sys
 - setting options in 45
 - dsm.sys file
 - creating 43
 - example of 43
 - dsm.sys.smp file 43
 - dsmerror.log 49
 - set DSM_LOG to point to 48
 - DSMI_CONFIG environment variable
 - API, UNIX and Linux 50
 - DSMI_DIR environment variable
 - API, UNIX and Linux 50
 - DSMI_LOG environment variable
 - API, UNIX and Linux 50
 - dsmtca executable file
 - set DSM_DIR to point to 48
 - dynamic and shared serialization 161
- ## E
- editor option 234
 - enablearchiveretentionprotection option 235
 - enablelanfree option 237
 - encryption
 - multiple clients under a single node name 197
 - of file data 84
 - saving encryption key password 240
 - encryption processing
 - encryption methods available 84
 - excluding files from 246
 - include files for 269
 - query systeminfo command 462
 - encryptiontype option 84, 239
 - encryptkey option 240
 - enhanced query schedule 151
 - enhanced query schedule command 460
 - environment prerequisites
 - AIX client 3
 - HP-UX Itanium 2 client 5
 - HP-UX PA-RISC client 4
 - Linux for zSeries client 9
 - Linux on POWER 7
 - Linux x86/x86_64 client 6
 - Solaris client 10
 - z/OS UNIX System Services client 10
 - environment variables
 - DSM_CONFIG 48
 - DSM_DIR 48
 - DSM_LOG 48
 - LANG 48
 - setting API 50
 - setting Bourne and Korn shell 49
 - setting C shell 50
 - error log 48

- error log (*continued*)
 - controlling the size 242
 - pruning 244
 - specifying path and file name 243
- error processing, summary of options 185
- errorlogmax option 242
- errorlogname option 243
- errorlogretention option 244
- estimate function 93
- event logging
 - scheduler 153
- event-based policy retention protection
 - archive 166
 - backup 166
- exclude options 246
 - exclude.archive 60
 - exclude.attribute.symblink 60
 - exclude.backup 60
 - exclude.compression 60
 - exclude.dir 60
 - exclude.file 60
 - exclude.file.backup 60
 - exclude.fs 60
 - exclude.image 60
 - processing 65
 - testing 64
 - wildcard characters 62
- exclude.image option 60
- excluding files
 - using wildcard characters 62
 - wildcard characters 62
- executables
 - return codes from 155
- expire command 427
- extended permissions
 - archive 147
 - backup 91

F

- file space
 - delete 95, 140, 423
 - determining fsID 221
 - performing an image backup 406
 - specifying NAS or client objects 202
- file spaces
 - excluding 60
- file specification
 - maximum allowed on commands 399
- file system
 - define virtual mount point for 382
- file systems
 - ACL support for 84
 - excluding from backup processing 246
 - GPFS, multi-node cluster environment 86, 228
 - image backup of 102
 - QFS, restrictions 85
 - supported 84
- filelist option 250
- filename option 252
- files
 - archive a list of 145, 250
 - archive using commands 147
 - archived, overriding management class 164
 - archives, how managed 150
 - archiving 143, 402
 - archiving more than one file specification 144

- files (*continued*)
 - assigning management classes 116
 - authorizing another user to restore or retrieve 138
 - backing up hard-linked 117
 - backing up open 119
 - backup, how managed 83
 - binding management classes to 165
 - compressing during archive or backup 211
 - definition of changed 87
 - delete after archive 218
 - delete individual archives from server file space 147, 418
 - delete individual backups from server file space 95, 420
 - encryption 84
 - excluding groups 62
 - how stored 114
 - including groups 62
 - managing growth during compression 210
 - maximum file size for operations 86
 - performing large restore operations 123
 - processing include-exclude 65
 - query archive information 442
 - query backup information 444
 - query user access 441
 - restore, point-in-time 126
 - restore, using commands 138
 - restore/retrieve to another workstation 139
 - restoring 121, 122
 - restoring files belonging to another node 139
 - restoring hard-linked 117
 - restoring sparse 118
 - retrieve archived 149
 - retrieve using commands 149
 - retrieving files belonging to another node 139
 - revoke access 417
 - sorting list of 76
- filesonly option 253
- firewall
 - establishing communications through 54, 262, 366
 - specify whether server or client initiates sessions through 344
 - specifying TCP/IP ports for the Web client 392
 - using Web client through 392
- followsymbolic option 254
- format
 - summary of options 184
- fromdate option 255
- fromnode option 256
- fromowner option 257
- fromtime option 258
- full incremental
 - comparing with incremental-by-date 90
 - comparing with journal-based, incremental-by-date 90
 - definition 87
 - description 87
 - when to use 90
- fuzzy backup 162

G

- getting started
 - changing your password 69
 - client scheduler 69
 - command line session 69
 - displaying online help 69
 - ending a session 69
 - GUI session 69
 - sorting file lists 69

- getting started (*continued*)
 - web client session 69
- globally unique identifier 58
- GPFS
 - storage pools 133
- GPFS file system
 - multi-node cluster environment 86, 228
- GPFS filesystem
 - scheduling 99
- graphical user interface
 - changing password 75
 - delete individual files or images from server file
 - space 420
 - displaying active and inactive backup versions 76, 121
 - displaying online help 76
 - displaying processing status 112
 - enabling for local backup set restore 282
 - enabling local backupset 130
 - ending a session 77
 - performing image backup 106
 - starting a session 70
 - using to back up objects 92
- group
 - backup, specifying full or differential 291
- group backup
 - display active and inactive objects 267
 - display all members of 347
 - overview 96
 - specify name of group 259
 - specify virtual file space name for 381
- groupname option 259
- groups option 260
- GUI
 - ending a session 77
 - overriding management class during archive 164
 - performing point-in-time restore 126
 - starting a session 70
- GUID (globally unique identifier) 58
- GUID commands
 - create 58
 - help 58
 - new 58
 - quiet 58
 - show 58
 - write 58
- guitreeviewafterbackup option 261

H

- HACMP cluster
 - scheduling 100
- HACMP takeover environment
 - installing Tivoli Storage Manager 511
- hard links
 - archive and retrieve 148
 - backing up 117
 - restoring 117
- hard mounts, NFS 118
- hardware requirements
 - AIX client 4
 - HP-UX Itanium 2 client 6
 - HP-UX PA-RISC client 5
 - Linux for zSeries client 9
 - Linux on POWER 8
 - Linux x86/x86_64 client 7
 - Solaris client 11
 - z/OS UNIX System Services client 10

- help
 - displaying online 76
 - Internet resources 76
 - online forum 76
 - service and technical support 76
- help command 429
- HP-UX client
 - communication methods 5
- HP-UX Itanium 2 client
 - client components 5
 - communication methods 6
 - disk space requirements 6
 - hardware requirements 6
 - increasing default limit of data segment size 23
 - installing 22
 - memory requirements 6
 - operating system requirements 6
 - uninstalling 24
- HP-UX PA-RISC client
 - client components 4
 - disk space requirements 5
 - hardware requirements 5
 - increasing default limit of data segment size 20
 - memory requirements 5
 - operating system requirements 5
 - uninstalling 21
- HP-UX PA-RISC clients
 - installing 18
- httpport option 262

I

- ifnewer option 263
- image
 - restoring 127
 - using chkdsk to repair 128
 - using fsck to repair 128, 476
- image backup
 - considerations 102
 - deleting 420
 - excluding files from 246
 - include files for; assign management class to 269
 - incremental-by-date image backup 106
 - perform 102
 - point-in-time restore 407
 - revoke access 417
 - specifying selective or incremental 291
 - static, dynamic, snapshot 102
 - using incremental mode 407
 - using the GUI 106
 - using with file system incremental 106, 407
 - using with incremental backup 105, 407
 - using with incremental-by-date 105
 - volume device type support 103
- image to file
 - restoring 129
- imagnetofile option 264
- imagetype option 265
- inactive backup versions
 - displaying 76, 121, 444
 - restoring 121
- inactive option 267
- inexcl option 268
- include option
 - management class 163
 - processing 65
 - wildcard characters 62

- include-exclude list
 - backup, used during 83
 - creating 59
 - query order of processing 455
 - size restriction 62, 66
 - testing 64
 - include-exclude options file
 - bottom-up processing 65
 - overview 83
 - specifying path and file name of 268
 - to manage archives 150
 - Unicode-enabled file spaces 268
 - include-exclude options list
 - creating 59
 - include-exclude processing
 - options for 60
 - overview 60
 - incrbydate option 274
 - incremental backup
 - back up new and changed files with modification date later than last backup 274
 - by date 94
 - description 87
 - GPFS, multi-node cluster environment 86, 228
 - new and changed files 87
 - of directories
 - processing overview 88
 - optimizing memory during 81
 - overview 87
 - specify memory-conserving algorithm for 289
 - symbolic links 116
 - using client Java GUI 92
 - using command line 93
 - using the client command line 93
 - using with image backup 105, 407
 - incremental command 430
 - journal-based backup 431
 - incremental image backup 407
 - incremental option 275
 - incremental-by-date
 - comparing with incremental 90
 - comparing with incremental, journal-based 90
 - description 90
 - of directories
 - processing overview 90
 - overview 87
 - using command line 94
 - when to use 90
 - incremental-by-date backup
 - using client Java GUI 92
 - using the client command line 93
 - using with image backup 105
 - installation
 - backup-archive client 13
 - installation requirements
 - AIX client 3
 - client 3
 - HP-UX Itanium 2 client 5
 - HP-UX PA-RISC client 4
 - Linux for zSeries client 9
 - Linux on POWER 7
 - Linux x86/x86_64 client 6
 - Solaris client 10
 - z/OS UNIX System Services client 10
 - installation steps
 - Tivoli Storage Manager on HP-UX PA-RISC clients 18
 - Tivoli Storage Manager on Linux for zSeries clients 28
 - installation steps (*continued*)
 - Tivoli Storage Manager on Linux x86/x86_64 clients 25
 - Tivoli Storage Manager on Sun Solaris clients 37
 - installing
 - AIX client 14
 - HP-UX Itanium 2 client 22
 - HP-UX PA-RISC clients 18
 - Linux for zSeries client 28
 - Linux on POWER client 31
 - Linux x86/x86_64 client 25
 - overview 40
 - Solaris client 34
 - installing from server CD-ROM 12
 - installing Tivoli Storage Manager
 - CD-ROM, installing from 12
 - HACMP takeover environment 511
 - interactive mode 396
 - interactive session
 - ending 435
 - starting 71, 435
 - using 435
- ## J
- Java GUI
 - configuration restrictions 70
 - journal configuration file
 - how to configure 517
 - journal daemon 56
 - journal configuration file settings 517
 - journal-based backup 89, 431
 - comparing with incremental, incremental-by-date 90
 - performing traditional full incremental, instead of 299, 431
 - specifying configuration settings 517
 - when to use 90
- ## K
- keyboard
 - accessibility xvi
 - disability xvi
 - shortcut keys xvi
- ## L
- LAN-free data movement
 - enabling 81, 237
 - enabling communications for 276, 279
 - options 81
 - prerequisites 81
 - specifying shared memory port for 278
 - lanfreecommmethod option 276
 - lanfreeshmport option 278
 - lanfreetcpport option 279
 - lanfreetcpserveraddress option 280
 - LANG environment variable
 - setting language locale 46
 - language locales
 - supported 46
 - last access date
 - specifying whether to update during backup or archive 88, 316
 - latest option 281
 - restore latest backup version 138

- Linux for zSeries client
 - client components 9
 - communication methods 10
 - disk space requirements 9
 - hardware requirements 9
 - installing 28
 - memory requirements 9
 - operating system requirements 9
 - uninstalling 30
- Linux Logical Volume Manager
 - snapshot image backup of volumes 102
- Linux on POWER client 8
 - client components 8
 - communication method 9
 - disk space requirements 8
 - hardware requirements 8
 - installing 31
 - memory requirements 8
 - operating system requirements 8
 - uninstalling 33
- Linux x86/x86_64 client
 - client components 6
 - communication methods 7
 - disk space requirements 7
 - hardware requirements 7
 - installing 25
 - memory requirements 7
 - operating system requirements 7
 - uninstalling 27
- local backupset
 - enabling GUI for local restore 130
- local snapshot
 - associating local snapshot with server file space 101
- localbackupset option 282
- location option 283
- LOFS
 - platform support for 228
- LOFS through automounter
 - platform support for 228
- log
 - DSM_LOG environment variable 243, 333
 - error log, controlling the size 242
 - error log, pruning 242
 - error log, specifying path and file name 242
 - errorlogname option 243
 - errorlogretention option 243
 - schedlogname option 333, 486
 - schedlogretention option 333, 486
 - specifying path and file name 243, 333, 486
- logical volume
 - image backup of 102
 - restoring 127, 129
- loop command 435
- loopback file systems
 - platform support for 228
- LVM
 - bring up an application after LVM starts 313
 - quiesce an application before LVM starts 321

M

- macro command 437
- mailprog option 284
- makesparsefile option 285
- managedservices option 286
- management class
 - assigning 116

- management classes
 - assigning to directories 164, 222
 - assigning to files 163
 - binding archive files to 145
 - binding to files 165
 - default 158
 - displaying 159
 - displaying information about 456
 - how Tivoli Storage Manager uses 83
 - overriding during archive processing 164
 - overriding the default 163
 - processing 163
 - questions to consider 162
 - selecting for files 162
 - specifying with include option 163
 - using management class, example 163
- maxcmdretries option 288
- memory
 - optimizing when constrained 81
- memory requirements
 - AIX client 4
 - HP-UX Itanium 2 client 6
 - HP-UX PA-RISC client 5
 - Linux for zSeries client 9
 - Linux on POWER 8
 - Linux x86/x86_64 client 7
 - Solaris client 11
 - z/OS UNIX System Services client 10
- memoryefficientbackup option 289
- messages
 - displaying on screen 379
 - stop displaying 323
- migrated files
 - restoring 121
- migrating from earlier versions of the product 1
- migration
 - web client 2
 - web client language files 2
- mode option 291
- mode parameter 160
- modes
 - batch 396
 - interactive (loop) 396
- modified mode 160, 162
- monitor option 293
- monitor process command 438
- multiple sessions
 - backup objects, using 112

N

- NAS
 - backing up file systems 108
 - deleting file spaces 95, 140, 423
 - query node command 457
 - restore file systems 134, 479
 - restore NAS command 479
- NAS file systems backup
 - command line 110
 - web client GUI 109
- nasnodename option 294
- Network Attached Storage (NAS)
 - assigning management class to file systems 269
 - excluding files from backup 246
 - monitoring backup or restore operations 293
 - specifying for query 376
 - specifying full or differential backup 291

- Network Attached Storage (NAS) *(continued)*
 - specifying node name for operations 294
 - specifying whether to save Table of Contents for each file system backup 371
- Network Attached Storage (NAS) file server
 - backup file systems 108, 410
 - cancel backup and restore processes 415, 438
 - deleting file spaces 95, 140, 423
 - display file spaces on server 449
 - display nodes for which admin ID has authority 457
 - Querying file system images belonging to 444
 - restore file systems 134, 479
- Network Data Management Protocol (NDMP) 12
- Network File System (NFS)
 - backup file systems 111
- NFS
 - backing up file systems 111
 - hard mounts 118
 - soft mounts 118
- nfstimeout option 118, 295
- NLSPATH environment variable
 - displaying help browser menu in your language locale 47
 - to display help browser menu in your language locale 46
- no query restore 124
- node
 - specifying type to query 376
- node name 43
- Node name field 139
- nodename option 297
- nojournal option 299
- noprompt option 300
- numberformat
 - specifying 301
- numberformat option 301

O

- online forum
 - Tivoli Storage Manager 77
- online help
 - displaying 76
 - online forum 76
 - service and technical support 76
- online Help forum 76
- open registration
 - permissions 57
 - using 57
- operating system requirements
 - AIX client 4
 - clients 3
 - HP-UX Itanium 2 client 6
 - HP-UX PA-RISC client 5
 - Linux for zSeries client 9
 - Linux on POWER 8
 - Linux x86/x86_64 client 7
 - Solaris client 11
 - z/OS UNIX System Services client 10
- optfile option 303
- option 356
- options
 - archive, summary 173
 - archmc 195
 - archsymbkfile 196
 - asnodename 197
 - authorization, summary 184
 - automount 199
 - backup, summary 173

- options *(continued)*
 - backupsetname 200
 - central scheduling, summary 182
 - changingretries 201
 - class 202
 - clusternode 203
 - collocatebyfilespec 204
 - command processing, summary 184
 - commmethod 206
 - commrestartduration 208
 - commrestartinterval 209
 - communication, summary 169
 - compressalways 210
 - compression 211
 - console 213
 - dateformat 214
 - defaultserver 217
 - deletefiles 218
 - description 219
 - detail 221
 - diagnostics 187
 - dirmc 222
 - dirsonly 223
 - diskbuffsize 224
 - diskcachelocation 225
 - domain 227
 - domain.image 231
 - domain.nas 232
 - editor 234
 - enablearchiveretentionprotection 235
 - enablelanfree 237
 - encryptiontype 84, 239
 - encryptkey 240
 - errorlogmax 242
 - errorlogname 243
 - errorlogretention 244
 - exclude
 - exclude.archive 60, 246
 - exclude.attribute.symlink 60, 246
 - exclude.backup 60, 246
 - exclude.compression 60, 246
 - exclude.dir 60, 246
 - exclude.encrypt 246
 - exclude.file 60, 246
 - exclude.file.backup 60, 246
 - exclude.fs 60, 246
 - exclude.fs.nas 246
 - exclude.image 60, 246
 - wildcard characters 62
 - filelist 250
 - filename 252
 - filesonly 253
 - followsymbolic 254
 - format, summary 184
 - fromdate 255
 - fromnode 256
 - fromowner 257
 - fromtime 258
 - general rules when entering with commands 188
 - groupname 259
 - groups 260
 - guitreeviewafterbackup 261
 - httpport 262
 - ifnewer 263
 - imagetofile 264
 - imagetype 265
 - inactive 267

- options (*continued*)
 - includ 268
 - include 269
 - management class, specifying 163
 - wildcard characters 62
 - include.archive 269
 - include.attribute.symlink 269
 - include.backup 269
 - include.compression 269
 - include.encrypt 269
 - include.file 269
 - include.fs.nas 269
 - include.image 269
 - incrbydate 274
 - incremental 275
 - lanfreecommmethod 276
 - lanfreeshmport 170, 278
 - lanfreetcpport 279
 - lanfreetcpserveraddress 280
 - latest 281
 - localbackupset 282
 - location 283
 - mailprog 284
 - makesparsefile 285
 - manageservices 286
 - maxcmdretries 288
 - memoryefficientbackup 289
 - mode 291
 - monitor 293
 - nasnodename 294
 - nfstimeout 295
 - nodename 297
 - nojournl 299
 - noprompt 300
 - numberformat 301
 - optfile 303
 - order of processing (precedence) 187
 - password 304
 - passwordaccess 305
 - passworddir 307
 - pick 308
 - pitdate 309
 - pittime 310
 - postnschedulecmd 311
 - postschedulecmd 311
 - postsnapshotcmd 313
 - preschedulecmd 314
 - preschedulecmd 314
 - preservelastaccessdate 316
 - preservepath 318
 - presnapshotcmd 321
 - querschedperiod 322
 - quiet 323
 - removeoperandlimit 324
 - replace 325
 - resourceutilization 327
 - restore and retrieve, summary 179
 - retryperiod 329
 - revokeremoteaccess 330
 - schedcmddisabled 331
 - schedcmduser (server defined only) 154
 - schedlogmax 332
 - schedlogname 333
 - schedlogretention 334
 - schedmode 336
 - schedrestretrdisabled 338
 - scrolllines 339

- options (*continued*)
 - scrollprompt 340
 - servername 342
 - sessioninitiation 344
 - shmport 346
 - showmembers 347
 - skipacl 348
 - skipaclupdatecheck 349
 - snapshotcachesize 350
 - snapshotroot 351
 - specifying in commands 187
 - srvprepostscheddisabled 354
 - srvprepostsnapdisabled 355
 - subdir 356
 - tapeprompt 358
 - tcpadminport 360
 - tcpbuffsize 361
 - tcpaddress 362
 - tcpclientaddress 363
 - tcpclientport 364
 - tcpnodelay 365
 - tcpport 366
 - tcpserveraddress 367
 - tcpwindowsize 368
 - timeformat 369
 - toc 371
 - todate 373
 - totime 374
 - transaction processing, summary 186
 - txnbytelimit 375
 - type 376
 - users 377
 - v2archive 378
 - verbose 379
 - verifyimage 380
 - virtualfsname 381
 - virtualmountpoint 382
 - virtualnodename 384
 - wasexphome 386
 - washome 387
 - wasndhome 388
 - wasnode 389
 - wastype 390
 - wasuser 391
 - Web client, summary 187
 - webports 392

P

- parameters
 - yes and no, alternatives 194
- partial incremental
 - definition 87
 - incremental by date, running 94
- password
 - changing 75, 495
 - number of characters 75
 - setting 304
 - setting for client 57
 - specifying directory location to store encrypted password file 307
 - specifying whether to generate automatically or set as user prompt 305
 - using 71
 - valid characters 75
- password option 304
- passwordaccess option 305

- passworddir option 307
- performance
 - improving speed of backups, restores, archives, retrieves 170
 - transaction options 186
 - transaction processing 375
- permissions
 - access, saving standard and extended 147
 - backup 91
- pick option 308
- pitdate 309
- pittime option 310
- plug-in library
 - for image or NAS backup or restore 48
- point-in-time restore 126
 - image backup 407
- policies, storage management 157
- policy domains
 - default policy domain 157
 - standard policy domain 157
- policy sets
 - active policy set 157
- portable media
 - restoring backupsets 130
- postschedulecmd option 311
- postsnapshotcmd option 313
- Preferences editor
 - excluding domains from back up 92
- preschedulecmd option 314
- preschedulecmd option 314
- preservelastaccessdate option 316
- preservepath option 318
- Presnapshotcmd option 321
- preview command 439
- previewing
 - excluded files 439
 - included files 439
- processing options
 - authorization 184
 - backup and archive 173
 - central scheduling 182
 - communication 169
 - diagnostics 187
 - error processing 185
 - format 184
 - node option 171
 - overview 169
 - restore and retrieve 179
 - server and node 171
 - setting 45
 - specifying in commands 187
 - transaction processing 186
 - using 169
 - Web client 187
- processing symbolic links 248, 271
- processing time
 - estimating 93
- proxied session restrictions 98, 146

Q

- QFS file system
 - restrictions 85
- query
 - backups, establish point-in-time 309, 310
 - based on date and time of backup, archive 255, 258
 - display active and inactive objects 267

- query (*continued*)
 - files for another node 256
 - group
 - command 451
 - display members of 347
 - include-exclude list 455
 - nodes to which client has proxy authority 145
 - nodes to which client has proxy node authority 97
 - process directories only; not files 223
 - specify end date and time to search for backups or archives 373, 374
 - specify location of backup set for 283
 - specifying amount of information that displays on screen 339
 - specifying description for 219
 - specifying NAS or client objects 202
 - specifying scrolling preferences after displaying information on screen 340
 - specifying whether to include subdirectories for 356
 - system information 462
 - WAS
 - Application Server 389, 390, 464
 - Network Deployment Manager 389, 390, 464
 - Network Deployment Manager and Application Server 389, 390, 464
- query access command 441
- query archive command 442
- query backup command 444
- query backupset command 447
- query filespace command 449
- query group command 451
- query image command 453
- query inxcl command 455
- query mgmtclass command 159, 456
- query node command 457
- query options command 458
- query restore command 459
- query schedule
 - enhanced 151
- query schedule command 460
- query schedule command, enhanced 460
- query session command 461
- query systeminfo command 462
 - encryption processing 462
- query was command 464
- querschedperiod option 322
- quiet option 323

R

- raw logical volume
 - image backup of 102
 - restoring 127
- rebinding files to a different management class 165
- recall commands
 - limitations 400
- Recommendations
 - default diskbuffsize setting 224
 - define cluster resources in domain statement 513
 - discontinue use of largecommbuffers xxvi, 186
 - do not write on volume while backup is in progress 102
 - HACMP node name 513
 - long user and group names 86
 - long user names 377
 - path names for setting environment variables 44
 - set up SSL in the Integrated Solutions Console (ISC) environment 56

- Recommendations (*continued*)
 - setting the WAS security password 496
 - specifying diskbuffsize xxvi
 - target nodes and traditional nodes 97
 - to simplify retrieve 80
 - traditional scheduler system resources 503
 - using dynamic imagetype 102
- registering
 - client with server 57
 - using closed registration 57
 - using open registration 57
- removeoperandlimit option 324
- replace option 325
- resourceutilization option 327
- restart restore command 465
 - restart interrupted restore 138
- restartable restore 124
- restartable restore sessions, display 459
- restore
 - a list of files 250
 - active version 121
 - another user's files 257
 - authorizing another user 138
 - backup sets
 - supported tape devices 470
 - backups, establish point-in-time 309, 310
 - backupsets
 - overview 130
 - based on date and time of backup 255, 258
 - create list of backup versions to 308
 - data using command line 122
 - disk 140
 - display active and inactive objects 267
 - estimating processing time 93
 - files 114, 121
 - files and directories 122
 - files belonging to another node 139
 - files for another node 256
 - files only 253
 - files to another workstation 139
 - from portable media
 - overview 130
 - group
 - command 474
 - GUI, displaying active and inactive versions 76
 - hard links 117
 - image 127
 - considerations 476
 - enable detection of bad sectors on target volume 380
 - to a file 264
 - using chkdsk tool to repair 128
 - using DSM_DIR to point to plug-in library 48
 - using fsck tool to repair 128, 476
 - image to file 129
 - image, suppress confirmation prompt 300
 - improving speed using share memory 170
 - inactive version 121
 - local backup set via GUI 282
 - local backupset via GUI 130
 - logical volume 127, 129
 - maximum file size 86
 - most recent backup version 281
 - NAS
 - using DSM_DIR to point to plug-in library 48
 - NAS file systems 134
 - NAS file systems using command line 135
 - NAS file systems using web client 134
- restore (*continued*)
 - no query 124
 - performing large operations 123
 - point-in-time 126
 - point-in-time, using command line 127
 - point-in-time, using GUI 126
 - primary tasks 121
 - process directories only; not files 223
 - processing status window 122
 - raw logical volume 127
 - regulate resources server and client use during
 - processing 327
 - replace existing file with latest backup 263
 - sorting file list 76
 - sparse files 118
 - specify end date and time to search for backups 373, 374
 - specify location of backup set for 283
 - specifying whether to include subdirectories for 356
 - specifying whether to prompt before overwriting existing files 325
 - specifying whether to wait for tape to mount if it is required 358
 - starting a Web client session 72
 - summary of options 179
 - symbolic links 116
 - UNIX and Linux restrictions 466
 - UNIX restrictions 116
 - to different workstation 384
 - using commands 138
 - WAS
 - Application Server 389, 390, 481
 - Network Deployment Manager 389, 390, 481
- restore backupset command 470
- restore command 466
 - performing large operations 123
- restore group command 474
- restore image command 476
- restore NAS command 479
- restore was command 481
- restrictions
 - within a proxied session 98, 146
- Restrictions
 - asnodename option 197
 - specifying full path with Client Acceptor Daemon 303
- retain extra versions attribute 161
- retain only versions attribute 161
- retain versions attribute 162
- retention grace period
 - archive 159, 166
 - backup 159, 166
- retrieve
 - a list of files 250
 - archive copies 148
 - archived files using commands 149
 - authorizing another user 138
 - based on date and time of archive 255, 258
 - files belonging to another node 139
 - files for another node 256
 - files only 253
 - files to another workstation 139
 - hard links 148
 - improving speed using share memory 170
 - maximum file size 86
 - primary tasks 143
 - process directories only; not files 223
 - replace existing file with latest archive if existing file is newer 263

- retrieve (*continued*)
 - running 149
 - sorting file list 76
 - specify end date and time to search for archives 373, 374
 - specifying description for 219
 - specifying whether to include subdirectories for 356
 - specifying whether to prompt before overwriting existing files 325
 - specifying whether to wait for tape to mount if it is required 358
 - starting a Web client session 72
 - summary of options 179
 - symbolic links 147
 - to different workstation 384
- retrieve command 484
- retryperiod option 329
- return codes for operations 155
- revokeremoteaccess option 330
- root user
 - tasks 41
- root user tasks
 - creating default client user options file 44
 - setting up 43

S

- SAN
 - restoring backup sets using 470
 - using for LAN-free data movement 81
- SAN File System FlashCopy images
 - back up 352
- schedcmddisabled option 331
- schedcmduser option (server defined only) 154
- schedlogmax option 332
- schedlogname option 333
- schedlogretention option 334
- schedmode option 336
- schedrestretrdisabled option 338
- schedule command 486
- schedule log
 - specifying number of days to keep entries and whether to save pruned entries 334
 - specifying path and file name to store schedule log information 333
- schedule log, controlling the size 332
- schedule log, pruning 332
- schedule log, specifying path and file name 332
- scheduled (automated) backups
 - closing files before back up 119
 - displaying scheduled work 151, 153
 - options for 151
 - overview 151, 499
 - process commands after back up 311
 - process commands before back up 314
 - restart applications after back up 119
 - starting 52
- scheduled commands
 - enabling-disabling 154
- scheduled events, displaying 460
- scheduled services
 - defining schedules for uid other than 0 154
 - disabling scheduled commands 331
 - restrictions for NAS file systems 108
- schedulelog
 - controlling the size 332
- scheduler
 - configuring 51

- scheduler (*continued*)
 - displaying scheduled work 151, 153
 - event logging 153
 - managed by client acceptor daemon 286
 - options for 151
 - overview 151, 499
 - resolving memory retention after scheduled backups 286
 - specify whether server or client initiates sessions through firewall 344
 - specifying number of hours between contacts to server for scheduled work 322
 - specifying number of minutes between attempts to process scheduled commands 329
 - specifying polling mode or prompted mode 336
 - specifying whether to disable execution of restore or retrieve operations 338
 - starting 52
 - troubleshooting 508
- scheduler information
 - completed work 510
 - scheduled work 510
- scheduler log
 - log file 510
 - schedlogname option 510
 - schedlogretention option 510
 - understanding the 510
- Scheduler Service Configuration Utility 52
- scheduler wizard 52
- scheduling
 - client node proxy backup 97, 98
 - GPFS filesystem 99
 - HACMP cluster 100
- scrolllines option 339
- scrollprompt option 340
- selective backup
 - overview 87, 91, 94
 - symbolic links 116
 - using client Java GUI 92
 - using command line 94
 - using the client command line 93
- selective command 488
- serialization
 - copy serialization
 - dynamic 161
 - shared static 161
 - static 161
- server
 - establishing communications through firewall 54
 - establishing communications with 43
 - identify to begin a stanza containing options for 342
 - specifying name of server to contact for services 342
 - specifying TCP/IP address of Tivoli Storage Manager server 367
 - specifying TCP/IP port address for 366
- server and node options
 - summary 171
- servername option 342
- service and technical support 76
- session information, displaying 461
- sessioninitiation option 344
- set access command 491
 - restore-retrieve authorization 138
- set event command 493
- set password command 495
- set waspassword command 496
- setting
 - user privileges 74

- setting environment variables
 - API, UNIX and Linux
 - DSMI_CONFIG 50
 - DSMI_DIR 50
 - DSMI_LOG 50
- setting language locale 46
- setting up
 - required root user tasks 43
- shared dynamic serialization 161, 201
- Shared Memory communication method 276
 - options 170
- Shared memory communication protocol 170
- shared static serialization 161, 201
- shell scripts
 - return codes from 155
 - using commands in 155
- shmport option 346
- shortcut keys
 - accessibility xvi
 - disability xvi
 - keyboard xvi
- showmembers option 347
- skipacl option 348
- skipaclupdatecheck option 349
- snapshotcachesize option 350
- snapshotroot option 351
- socket files
 - skipped during restore 121
- soft mounts, NFS 118
- software requirements
 - AIX client 4
 - HP-UX Itanium 2 client 6
 - HP-UX PA-RISC client 5
 - Linux for zSeries client 9
 - Linux on POWER 8
 - Linux x86/x86_64 client 7
 - Solaris client 11
 - z/OS UNIX System Services client 10
- Solaris
 - installing Solaris 10 in zones 37
- Solaris client
 - client components 10
 - communication methods 12
 - disk space requirements 11
 - hardware requirements 11
 - installing 34
 - memory requirements 11
 - operating system requirements 11
- sparse files
 - backing up 118
 - restoring 118
 - specifying how to restore or retrieve 285
- special file systems 84, 115
- srvprepostscheddisabled option 354
- srvprepostsnapdisabled option 355
- standard management class
 - copy destination 162
 - copy frequency 160
 - copy group name 160
 - copy mode
 - absolute 162
 - modified 162
 - copy serialization 161
 - copy type 160
 - default values 159
 - retain extra versions 161
 - retain only version 161
- standard management class (*continued*)
 - retain versions 162
 - versions data deleted
 - active versions 160
 - inactive versions 160
 - versions data exists 160
- standard policy domain 157
- starting
 - automatically 75
 - overview 40
- starting a session
 - batch mode 71
 - interactive mode 71
- static serialization 161
- static, shared serialization 161
- storage
 - displaying restartable restore sessions 459
- Storage Agent
 - using for LAN-free data movement 81, 237
- storage area network
 - restoring backup sets using 237, 470
 - using for LAN-free data movement 81, 237
- storage management policies 157
 - assigning management classes to files 116
 - copy groups 158
 - default management class 157
 - display on backup-archive client or Web client GUI 116
 - include-exclude list 158
 - management classes 158
 - policy domains
 - default 157
 - standard 157
 - policy sets
 - active policy set 157
- storage pools
 - GPFS 133
- subdirectories
 - archive 145
 - include in backup 94
- summary of changes for version 5.3.0 xxiii
- summary of changes for version 5.3.2 xix
- support
 - gathering system information for 213, 252, 462
- supported language locales 46
- Swing-enabled browser
 - necessary to run Web client 72
- symbolic links
 - archiving and retrieving 147
 - backing up 248, 271
 - backing up and restoring 116
 - exclude from backup 64, 248, 271
 - exclude options 248, 271
 - processing 248, 271
 - restoring 254
 - UNIX and Linux restrictions 466
 - UNIX restrictions 116
- system information
 - gathering 213, 252, 462
- system objects
 - exclude from backup processing 247

T

- tapeprompt option 358
- tasks
 - assigning management classes to directories 164
 - client scheduler, starting 75

- tasks (*continued*)
 - closed registration 57
 - display management classes 159
 - GUI, override management class 164
 - open registration 57
 - password, change 75
 - root user 41
 - sessions, ending 75
- TCP/IP communication method
 - options 170
- tcpadminport option 360
- tcpbuffsize option 361
- tcpcadaddress option 362
- tcpclientaddress option 363
- tcpclientport option 364
- tcpnodelay option 365
- tcpserveraddress option 367
- tcpwindowsize option 368
- time format
 - specifying 369
- timeformat option 369
- tivguid 58
- Tivoli products and documentation
 - feedback survey xvi
 - ordering xvi
- Tivoli Storage Manager
 - client components
 - AIX client 3
 - HP-UX Itanium 2 client 5
 - HP-UX PA-RISC client 4
 - Linux for zSeries client 9
 - Linux x86/x86_64 client 6
 - Solaris client 10
 - z/OS UNIX System Services client 10
 - communication methods
 - AIX client 4
 - HP-UX client 5
 - HP-UX Itanium 2 client 6
 - Linux for zSeries client 10
 - Linux on POWER client 9
 - Linux x86/x86_64 client 7
 - Solaris client 12
 - z/OS UNIX System Services client 10
 - considerations when upgrading from ADSM AIX 3.1 16
 - hardware, disk space, memory requirements
 - AIX client 4
 - and z/OS UNIX System Services client 10
 - HP-UX Itanium 2 client 6
 - HP-UX PA-RISC client 5
 - Linux for zSeries client 9
 - Linux on POWER client 8
 - Linux x86/x86_64 client 7
 - Solaris client 11
 - Linux on POWER 8
 - migrating from earlier versions of the product 1
 - online forum 77
 - password 71
 - software requirements
 - AIX client 4
 - HP-UX Itanium 2 client 6
 - HP-UX PA-RISC client 5
 - Linux for zSeries client 9
 - Linux on POWER client 8
 - Linux x86/x86_64 client 7
 - Solaris client 11
 - z/OS UNIX System Services client 10

- Tivoli Storage Manager client
 - authentication 69
- Tivoli Storage Manager on HP-UX PA-RISC clients
 - installation steps 18
- Tivoli Storage Manager on Linux for zSeries clients
 - installation steps 28
- Tivoli Storage Manager on Linux x86/x86_64 clients
 - installation steps 25
- Tivoli Storage Manager on Sun Solaris clients
 - installation steps 37
 - uninstalling 40
- Tivoli Storage Manager password
 - using 70
- toc option 371
- todate option 373
- totime option 374
- transaction processing 375
 - summary of options 186
- TSM.PWD file
 - HP-UX restriction 240
- tsmjbbd.ini
 - configuring 517
- txnbytelimit option 375
- type option 376

U

- uninstalling
 - AIX client 17
 - HP-UX Itanium 2 client 24
 - HP-UX PA-RISC client 21
 - Linux for zSeries client 30
 - Linux on POWER client 33
 - Linux x86/x86_64 client 27
 - Tivoli Storage Manager on Sun Solaris clients 40
- UNIX
 - file systems, ACL support 84
 - restrictions
 - restoring symbolic links 116
 - saving standard access permissions 147
- UNIX and Linux
 - restrictions
 - restoring symbolic links 466
- UNIX clients
 - considerations when upgrading from ADSM AIX 3.1 16
- user privileges
 - setting 74
- users option 377

V

- v2archive option 378
- verbose option 379
- verifyimage option 380
- versions data
 - deleted attribute 160
 - deleted parameter 161
 - exists attribute 160
 - exists parameter 161
- virtual mount point, setting 92
- virtualfsname option 381
- virtualmountpoint option 382
- virtualnodename option 384
 - restore/retrieve to another workstation 139

W

- wasexphome option 386
- washome option 387
- wasndhome option 388
- wasnode option 389
- wastype option 390
- wasuser option 391
- Web client
 - enable to run in a Swing-enabled browser 72
 - establishing communications through firewall 54, 262
 - managed by client acceptor daemon 286
 - restrict administrator from accessing client running Web client 330
 - restrictions for NAS file systems 108
 - specifying TCP/IP port address for 262
 - starting 72
 - summary of options 187
 - supported browsers 72
 - unsupported functions 79
 - using through a firewall 392
- web client GUI
 - NAS file systems backup 109
- webports option 392
- WebSphere Application Server (WAS)
 - back up WebSphere Application Server-Express 386
 - query
 - Application Server 464
 - Network Deployment Manager 464
 - restore
 - Application Server 413, 481
 - Network Deployment Manager 413, 481
 - set WebSphere node name, user name, password 496
 - specify override base install path for Network Deployment Manager 388
 - specify override base install path for Application Server 387
 - specifying full or differential backup 291
- wildcard characters
 - include or exclude groups of files 62
 - specifying a drive specification in dsm.opt 62
 - to include or exclude groups of files 62
 - using 400
 - using with commands 119
 - using with file specifications 119

Z

- z/OS UNIX System Services client
 - client components 10
 - communication methods 10
 - disk space requirements 10
 - hardware requirements 10
 - memory requirements 10
 - operating system requirements 10



Program Number: 5608-HSM
5608-ISM
5608-ISX
5608-SAN

Printed in USA

SC32-0145-00



Spine information:

IBM Tivoli Storage Manager for UNIX and Linux **Version 5.4**

Backup-Archive Clients Installation and User's Guide

