



**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 635.

**Edition notice**

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

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

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

© 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

<b>Tables</b> . . . . .	<b>ix</b>
-------------------------	-----------

<b>About this book</b> . . . . .	<b>xi</b>
----------------------------------	-----------

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

<b>Summary of changes for Tivoli Storage Manager</b> . . . . .	<b>xix</b>
--	------------

Technical changes for Version 5.5 - November 2007 . . . . .	xix
Technical changes for Version 5.4 - January 2007 . . . . .	xx
Technical changes for Version 5.3 - March 2006 . . . . .	xxii
Technical changes for Version 5.3 - October 2005 . . . . .	xxii
Technical changes for Version 5.3 - December 2004 . . . . .	xxiii

<b>Chapter 1. Installing Tivoli Storage Manager</b> . . . . .	<b>1</b>
---	----------

Migrating from earlier versions . . . . .	1
Upgrade path for clients and servers . . . . .	1
Considerations for migrating between processor architectures (x32, x64, IA64) . . . . .	3
Additional migration information . . . . .	4
Upgrading Open File Support (OFS) or online image. . . . .	4
Client environment requirements . . . . .	5
Client components . . . . .	5
Hardware requirements. . . . .	5
Software requirements . . . . .	6
Communications methods . . . . .	7
Supported features . . . . .	7
NDMP support requirements (Extended Edition only) . . . . .	8
Client installation media on product CD . . . . .	8
Installing the software on your workstation . . . . .	9
Silent installation . . . . .	10
Upgrading, modifying, or reinstalling the client . . . . .	12
Installation troubleshooting . . . . .	13

<b>Chapter 2. Configuring Tivoli Storage Manager</b> . . . . .	<b>15</b>
--	-----------

Creating and modifying the client options file (required) . . . . .	15
Creating a shared directory options file (optional) . . . . .	17
Creating multiple options files . . . . .	17
Setting options in the client options file . . . . .	18
Setting environment variables . . . . .	19
Configuring the Web client . . . . .	20
Configuring the client scheduler . . . . .	21

Starting the client scheduler . . . . .	23
Configuring Tivoli Storage Manager client/server communication across a firewall . . . . .	23
Configuring Tivoli Storage Manager client/server communication with Secure Socket Layer . . . . .	26
Configuring the journal engine service . . . . .	28
Configuring online image backup support . . . . .	28
Configuring Open File Support (OFS) . . . . .	29
Registering your workstation with a server (required) . . . . .	30
Using closed registration . . . . .	30
Using open registration . . . . .	31
Creating an include-exclude list (optional) . . . . .	31
Using include-exclude options . . . . .	32
Controlling compression, encryption, and adaptive subfile backup processing . . . . .	39
Testing an include-exclude list with the Preview command . . . . .	40
Processing include and exclude options . . . . .	41
Processing rules when using UNC names . . . . .	42

<b>Chapter 3. Getting started</b> . . . . .	<b>45</b>
---	-----------

Considerations for client services . . . . .	45
Using Tivoli Storage Manager Backup-Archive client as a member of the Backup Operators group . . . . .	45
Considerations before you start using a Backup Operators group account for backup, archive, restore, and retrieve operations . . . . .	47
Permissions required to restore files that use adaptive subfile backup . . . . .	47
Permissions required to back up, archive, restore or retrieve files on cluster resources . . . . .	47
Tivoli Storage Manager client authentication . . . . .	47
Starting a GUI session. . . . .	48
Using your Tivoli Storage Manager password . . . . .	49
Setup wizard . . . . .	49
Starting a command-line session . . . . .	49
Using batch mode . . . . .	49
Using interactive mode . . . . .	50
Displaying Euro characters in command-line prompt. . . . .	50
Using options on the DSMC command . . . . .	51
Starting a Web client session. . . . .	51
Setting user privileges . . . . .	53
Starting the client scheduler automatically . . . . .	53
Changing your password. . . . .	53
Sorting file lists using the Tivoli Storage Manager GUI . . . . .	54
Displaying online help . . . . .	55
Ending a session. . . . .	55
Online forums . . . . .	55
Other sources of online help. . . . .	56

<b>Chapter 4. Backing up your data</b> . . . . .	<b>57</b>
--	-----------

Planning your backups . . . . .	57		Understanding how deleted file systems are handled . . . . .	98
When to back up and when to archive files. . . . .	58		Backing up removable media . . . . .	99
Pre-backup considerations . . . . .	59		Backing up fixed drives . . . . .	99
LAN-free data movement. . . . .	59		Backing up NTFS file spaces. . . . .	99
Configure memory-constrained systems to run incremental backups . . . . .	60		Using Universal Naming Convention names . . . . .	100
Configure systems with large numbers of files to run incremental backups . . . . .	60		Backing up Microsoft Dfs files. . . . .	101
Performing a backup with limited bandwidth . . . . .	61			
Using an include-exclude list to control processing. . . . .	62		<b>Chapter 5. Restoring your data. . . . .</b>	<b>105</b>
Maximum file size for operations . . . . .	63		Duplicate file names . . . . .	105
How Tivoli Storage Manager handles long user and group names . . . . .	64		Restoring files using Universal Naming Convention names. . . . .	106
Performing an incremental, selective, or incremental-by-date backup . . . . .	64		Restoring an active or inactive backup . . . . .	106
Supported file systems . . . . .	64		Restoring files and directories . . . . .	107
Full and partial incremental backup . . . . .	65		Restoring data using the GUI . . . . .	107
Incremental-by-date backup . . . . .	68		Restoring data using the command line . . . . .	107
Comparing full incremental, partial incremental, incremental-by-date, and journal-based backups	68		Restoring migrated files . . . . .	111
Selective backup. . . . .	69		Restoring system objects. . . . .	111
Open file support for backup operations. . . . .	69		Restoring Windows system state . . . . .	111
Backing up data using the GUI. . . . .	71		Restoring Automated System Recovery (ASR) files (Windows XP, Windows Server 2003) . . . . .	113
Backing up data using the command line . . . . .	73		Performing a Windows XP or Windows Server 2003 system recovery . . . . .	114
Deleting backup data . . . . .	75		Operating system functions but requires complete system restore . . . . .	114
Group backup: Backing up files from one or more file spaces . . . . .	76		Complete recovery required, including operating system re-installation . . . . .	114
Backup with client-node proxy support . . . . .	76		Restoring Microsoft Dfs trees and files . . . . .	120
Enabling multiple node operation from the GUI	77		Restoring an image . . . . .	120
Setting up encryption . . . . .	78		Performing an image restore using the GUI . . . . .	121
Scheduling with client-node proxy support. . . . .	78		Performing an image restore using the command line . . . . .	122
Associating a local snapshot with a server file space	79		Restoring data from a backupset . . . . .	122
Backing up system objects (Windows XP) . . . . .	79		Restoring an entire or partial backupset . . . . .	123
Backing up Windows system state. . . . .	80		Restoring backupsets using the GUI. . . . .	123
Backing up Automated System Recovery (ASR) files (Windows XP, Windows Server 2003). . . . .	82		Restoring backupsets using the command-line client . . . . .	124
Performing an image backup . . . . .	82		Performing point-in-time restores. . . . .	125
Before you perform an image backup. . . . .	83		Restoring NAS file systems. . . . .	126
Utilizing image backup to perform file system incremental backup. . . . .	84		Restoring NAS file systems using the Web client	126
Performing an image backup using the GUI . . . . .	86		Restoring NAS files and directories using the Web client . . . . .	127
Performing an image backup using the command line . . . . .	87		Restoring NAS file systems using the command line. . . . .	129
Backing up NAS file systems . . . . .	87		Restoring NetApp CIFS shares . . . . .	130
Support for CDP Persistent Storage Manager (PSM) . . . . .	88		Restoring VMware Consolidated Backup data . . . . .	130
Backing up NAS file systems using the Web client GUI. . . . .	88		Restoring the WebSphere Application Server . . . . .	131
Backing up NAS file systems using the command line . . . . .	90		Restore: Additional considerations . . . . .	132
Backup and restore of NAS file servers using CIFS. . . . .	91		Authorizing another user to restore or retrieve your files. . . . .	132
Using VMware Consolidated Backup . . . . .	92		Restoring or retrieving files from another client node . . . . .	133
Backing up NetApp CIFS share definitions . . . . .	95		Restoring or retrieving your files to another workstation . . . . .	134
Backing up the WebSphere Application Server. . . . .	95		Restoring or retrieving files to another type of workstation . . . . .	134
Displaying backup processing status . . . . .	95		Restoring a disk in the event of disk loss . . . . .	134
Backup: Additional considerations. . . . .	97		Deleting file spaces . . . . .	135
Backing up open files . . . . .	97			
Understanding how files are managed . . . . .	98			

## Chapter 6. Archiving and retrieving your data . . . . . 137

Archiving files . . . . .	137
Archiving files using open file support . . . . .	138
Archiving data using the GUI . . . . .	138
Archiving data using the command line . . . . .	139
Archiving data with client node proxy . . . . .	140
Deleting archive data . . . . .	142
Retrieving archives . . . . .	142
Retrieving archives using the GUI . . . . .	143
Retrieving archives using the command line . . . . .	143

## Chapter 7. Automating tasks . . . . . 145

Specifying scheduling options . . . . .	145
Displaying information about scheduled work . . . . .	145
Displaying information about completed work . . . . .	147
Event logging . . . . .	148
Enabling or disabling scheduled commands . . . . .	150
Accessing Windows network drives during a scheduled backup . . . . .	150
Mobile dial-up support . . . . .	151
Client return codes . . . . .	151

## Chapter 8. Understanding storage management policies . . . . . 153

Using policy domains and policy sets . . . . .	153
Using management classes and copy groups . . . . .	153
Displaying information about management classes and copy groups . . . . .	155
Copy group name . . . . .	156
Copy type . . . . .	156
Copy frequency . . . . .	156
Versions data exists . . . . .	156
Versions data deleted . . . . .	156
Retain extra versions . . . . .	156
Retain only version . . . . .	157
Copy serialization . . . . .	157
Copy mode . . . . .	158
Copy destination . . . . .	158
Retain versions . . . . .	158
Selecting a management class for files . . . . .	159
Assigning a management class to files . . . . .	159
Overriding the management class for archived files . . . . .	160
Selecting a management class for directories . . . . .	160
Binding management classes to files . . . . .	161
Rebinding backup versions of files . . . . .	162
Using a retention grace period . . . . .	162
Using event-based policy retention protection . . . . .	163
Archiving files on a data retention server . . . . .	163

## Chapter 9. Using processing options 165

Overview of processing options . . . . .	165
Communication options . . . . .	166
TCP/IP options . . . . .	166
Named Pipes option . . . . .	167
Shared memory options . . . . .	167
Node options . . . . .	167
Backup and archive processing options . . . . .	168
Restore and retrieve processing options . . . . .	176
Scheduling options . . . . .	179

Format and language options . . . . .	180
Command processing options . . . . .	181
Authorization options . . . . .	181
Error processing options . . . . .	181
Transaction processing options . . . . .	182
Web client options . . . . .	183
Diagnostics options . . . . .	183
Using options with commands . . . . .	183
Entering options with a command . . . . .	184
Initial command-line-only options . . . . .	190
Client options reference . . . . .	190
Activatekey . . . . .	192
Archmc . . . . .	193
Asnodename . . . . .	194
Asrmode . . . . .	196
Auditlogging . . . . .	197
Auditlogname . . . . .	200
Autofsrename . . . . .	202
Backupsetname . . . . .	204
Casesensitiveaware . . . . .	205
Changingretries . . . . .	207
Class . . . . .	208
Clientview . . . . .	209
Clusterdiskonly . . . . .	210
Clusternode . . . . .	214
Collocatebyfilespec . . . . .	215
Commmethod . . . . .	217
Commrestartduration . . . . .	219
Commrestartinterval . . . . .	220
Compressalways . . . . .	221
Compression . . . . .	222
Computername . . . . .	224
Console . . . . .	225
Dateformat . . . . .	227
Deletefiles . . . . .	229
Description . . . . .	230
Detail . . . . .	232
Dfsbackupmntpnt . . . . .	233
Dirmc . . . . .	234
Dirsonly . . . . .	235
Disablenqr . . . . .	236
Diskbuffsize . . . . .	237
Diskcachelocation . . . . .	238
Domain . . . . .	239
Domain.image . . . . .	241
Domain.nas . . . . .	242
Editor . . . . .	244
Enable8dot3namesupport . . . . .	245
Enablearchiveretentionprotection . . . . .	247
Enablelanfree . . . . .	249
Encryptiontype . . . . .	250
Encryptkey . . . . .	251
Errorlogmax . . . . .	253
Errorlogname . . . . .	254
Errorlogretention . . . . .	255
Exclude options . . . . .	257
Filelist . . . . .	262
Filename . . . . .	264
Filesonly . . . . .	266
Fromdate . . . . .	267
Fromnode . . . . .	268

Fromtime . . . . .	269	Schedmode . . . . .	355
Frsprimaryrestore . . . . .	270	Schedrestretretrdisabled . . . . .	357
Groupname . . . . .	271	Scrolllines . . . . .	358
Guitreeviewafterbackup . . . . .	272	Scrollprompt . . . . .	359
Httpport . . . . .	273	Sessioninitiation . . . . .	361
Ifnewer . . . . .	274	Shmport . . . . .	363
Imagegapsize . . . . .	275	Showmembers . . . . .	364
Imagetofile . . . . .	276	Skipntpermissions . . . . .	365
Inactive . . . . .	277	Skipntsecuritycrc . . . . .	366
Incllexcl . . . . .	278	Snapshotcachelocation . . . . .	367
Include options . . . . .	280	Snapshotcachesize . . . . .	368
Processing NAS file systems . . . . .	285	Snapshotfsidleretries . . . . .	369
Processing WebSphere Application Server file systems . . . . .	285	Snapshotfsidlewait . . . . .	371
Incrbydate . . . . .	286	Snapshotproviderfs . . . . .	373
Incremental . . . . .	287	Snapshotproviderimage . . . . .	374
Incrthreshold . . . . .	288	Snapshotroot . . . . .	375
Journalpipe . . . . .	290	Srvprepostscheddisabled . . . . .	378
Lanfreecommmethod . . . . .	291	Srvprepostsnapdisabled . . . . .	379
Lanfreeshmport . . . . .	293	Ssl . . . . .	380
Lanfreetcpport . . . . .	294	Subdir . . . . .	381
Lanfreetcpserveraddress . . . . .	295	Subfilebackup . . . . .	383
Language . . . . .	296	Subfilecachepath . . . . .	384
Latest . . . . .	298	Subfilecachesize . . . . .	385
Localbackupset . . . . .	299	Tapeprompt . . . . .	386
Location . . . . .	300	Tcpadminport . . . . .	388
Managedservices . . . . .	302	Tcpbuffsize . . . . .	389
Maxcmdretries . . . . .	304	Tcpcadaddress . . . . .	390
Memoryefficientbackup . . . . .	305	Tcpclientaddress . . . . .	391
Mode . . . . .	307	Tcpclientport . . . . .	392
Monitor . . . . .	309	Tcpnodelay . . . . .	393
Namedpipename . . . . .	310	Tcpport . . . . .	394
Nasnodename . . . . .	311	Tcpserveraddress . . . . .	395
Nodename . . . . .	312	Tcpwindowsize . . . . .	396
Nojournal . . . . .	314	Timeformat . . . . .	397
Noprompt . . . . .	315	Toc . . . . .	399
Numberformat . . . . .	316	Todate . . . . .	401
Optfile . . . . .	318	Totime . . . . .	402
Password . . . . .	319	Txnbytelimit . . . . .	403
Passwordaccess . . . . .	320	Type . . . . .	404
Pick . . . . .	322	Usedirectory . . . . .	405
Pitdate . . . . .	323	V2archive . . . . .	406
Pittime . . . . .	324	Verbose . . . . .	407
Postschedulecmd/Postnschedulecmd . . . . .	325	Verifyimage . . . . .	408
Postsnapshotcmd . . . . .	327	Virtualfsname . . . . .	409
Preschedulecmd/Prenschedulecmd . . . . .	329	Virtualnodename . . . . .	410
Preservelastaccessdate . . . . .	331	Vmchost . . . . .	412
Preservepath . . . . .	333	Vmcpw . . . . .	413
Presnapshotcmd . . . . .	336	Vmcuser . . . . .	414
Queryschedperiod . . . . .	338	Vmllist . . . . .	415
Quiet . . . . .	339	Wasexphome . . . . .	416
Replace . . . . .	340	Washome . . . . .	417
Resetarchiveattribute . . . . .	342	Wasndhome . . . . .	418
Resourceutilization . . . . .	344	Wasnode . . . . .	419
Retryperiod . . . . .	347	Wastype . . . . .	420
Revokeremoteaccess . . . . .	348	Wasuser . . . . .	421
Runasservice . . . . .	349	Webports . . . . .	422
Schedcmddisabled . . . . .	350		
Schedlogmax . . . . .	351		
Schedlogname . . . . .	352		
Schedlogretention . . . . .	353		
		<b>Chapter 10. Using commands . . . . .</b>	<b>423</b>
		Starting and ending a client command session . . . . .	427
		Process commands in batch mode . . . . .	427
		Process commands in interactive mode . . . . .	428



Entering client commands . . . . .	428	Query VM . . . . .	519
Command name . . . . .	428	Query WMI . . . . .	520
Options . . . . .	428	Restart Restore . . . . .	521
Parameters . . . . .	429	Restore . . . . .	522
File specification syntax . . . . .	429	Restore ASR . . . . .	528
Using wildcard characters . . . . .	431	Restore Backupset . . . . .	529
Entering commands . . . . .	431	Restore Complusbdb . . . . .	532
Client commands reference . . . . .	431	Restore Eventlog . . . . .	533
Archive . . . . .	433	Restore Group . . . . .	534
Backup ASR . . . . .	436	Restore Image . . . . .	536
Backup Certserverdb . . . . .	437	Restore NAS . . . . .	540
Backup Complusbdb . . . . .	438	Restore Registry . . . . .	542
Backup Eventlog . . . . .	439	Restore Sysfiles . . . . .	544
Backup Group . . . . .	440	Restore Systemobject . . . . .	545
Backup Image . . . . .	442	Restore Systemstate . . . . .	547
Backup NAS . . . . .	447	Restore WAS . . . . .	549
Backup Registry . . . . .	450	Restore WMI . . . . .	552
Backup Sysfiles . . . . .	451	Retrieve . . . . .	553
Backup Systemobject . . . . .	452	Schedule . . . . .	557
Backup Systemstate . . . . .	453	Selective . . . . .	559
Backup WAS . . . . .	455	Set Access . . . . .	563
Backup VM . . . . .	457	Set Event . . . . .	566
Backup WMI . . . . .	459	Set Password . . . . .	568
Cancel Process . . . . .	460	Set Waspassword . . . . .	569
Cancel Restore . . . . .	461		
Delete Access . . . . .	462	<b>Appendix A. Using the Tivoli Storage</b>	
Delete Archive . . . . .	463	<b>Manager central scheduler . . . . .</b>	<b>571</b>
Delete Backup . . . . .	465	Overview of the Tivoli Storage Manager scheduler . . . . .	571
Delete Filespace . . . . .	468	Handling spaces in file names in schedule	
Delete Group . . . . .	470	definitions . . . . .	572
Expire . . . . .	472	Preferential start times for certain nodes . . . . .	572
Help . . . . .	474	Understanding scheduler processing options . . . . .	573
Incremental . . . . .	475	Handling return codes from preschedulecmd	
Loop . . . . .	482	and postschedulecmd scripts . . . . .	574
Macro . . . . .	484	Using the client acceptor to manage scheduler	
Monitor Process . . . . .	485	services versus the traditional scheduler services . . . . .	575
Preview . . . . .	486	Setting the client scheduler process to run as a	
Query Access . . . . .	487	background task and start automatically at boot	
Query Archive . . . . .	488	time . . . . .	575
Query ASR . . . . .	490	Changing processing options used by the scheduler	
Query Backup . . . . .	491	service . . . . .	577
Querying NAS file system images . . . . .	494	Managing multiple schedule requirements on one	
Query Backupset . . . . .	495	system . . . . .	578
Query Complusbdb . . . . .	497	Scheduled backups of Windows network drives . . . . .	579
Query Eventlog . . . . .	498	Managing different scheduling requirements for	
Query Filespace . . . . .	499	Windows system objects . . . . .	579
Query Group . . . . .	501	Restarting the scheduler process on a remote	
Query Image . . . . .	503	system . . . . .	579
Query Indexcl . . . . .	505	Using the scheduler on clustered systems . . . . .	580
Query Mgmtclass . . . . .	506	Troubleshooting the Tivoli Storage Manager	
Query Node . . . . .	507	scheduler . . . . .	581
Query Options . . . . .	508	Investigating abnormal scheduled events . . . . .	582
Query Registry . . . . .	509	Checking the server activity log . . . . .	582
Query Restore . . . . .	510	Inspecting the Client Schedule Log . . . . .	582
Query Schedule . . . . .	511	Understanding the scheduler log . . . . .	583
Query Session . . . . .	512		
Query Sysfiles . . . . .	513	<b>Appendix B. Using the Client Service</b>	
Query Systeminfo . . . . .	514	<b>Configuration Utility . . . . .</b>	<b>585</b>
Query Systemobject . . . . .	516	Installing the backup-archive scheduler service . . . . .	585
Query Systemstate . . . . .	517	Using the Client Service Configuration Utility . . . . .	585
Query WAS . . . . .	518		

Setting up a schedule to back up system objects (Windows XP) . . . . .	588
Using the Dsmcutil command . . . . .	589
Dsmcutil commands . . . . .	590
Dsmcutil options . . . . .	596

<b>Appendix C. Journal service configuration . . . . .</b>	<b>601</b>
JournalSettings stanza . . . . .	602
JournalExcludeList stanza . . . . .	603
JournalFileSystemSettings stanza . . . . .	604
Overriding stanzas . . . . .	608

<b>Appendix D. Configuring the backup-archive client in a cluster server environment. . . . .</b>	<b>609</b>
Installing the backup-archive client on the cluster nodes . . . . .	609
Configuring the backup-archive client to process local nodes . . . . .	609
Configuring the backup-archive client to process cluster disk resources. . . . .	610
Step 1: Identify the cluster groups to manage	610
Step 2: Configure the client options file. . . . .	611

Step 3: Configure the scheduler service. . . . .	612
Step 4: Creating a generic service resource for failover . . . . .	613
Configuring the Web client in a cluster environment. . . . .	615
Configuring the Web client to process cluster disk resources . . . . .	615
Frequently asked questions. . . . .	620

<b>Appendix E. ASR supplemental information . . . . .</b>	<b>625</b>
ASR for Windows XP and Windows 2003 . . . . .	625
ASR questions and answers . . . . .	625
ASR diagnostic tips . . . . .	628
ASR in depth . . . . .	629
Reference information . . . . .	633

<b>Notices . . . . .</b>	<b>635</b>
Trademarks . . . . .	636

<b>Glossary . . . . .</b>	<b>639</b>
---------------------------	------------

<b>Index . . . . .</b>	<b>651</b>
------------------------	------------

## Tables

1. Typographical conventions . . . . .	xii	44. Transaction processing options . . . . .	182
2. IBM Tivoli Storage Manager client and server publications . . . . .	xiv	45. Web client options . . . . .	183
3. Optimized option default values . . . . .	xxv	46. Diagnostics options . . . . .	183
4. Windows disk storage requirements. . . . .	6	47. Client command options. . . . .	185
5. Supported features on Windows platforms . . . . .	7	48. Options that are valid on the initial command line only . . . . .	190
6. Tivoli Storage Manager server Installation Guides . . . . .	13	49. <i>Clusternode</i> and <i>clusterdisksonly</i> combinations. . . . .	212
7. File path and name limits. . . . .	18	50. System services components and corresponding keywords. . . . .	259
8. Option for excluding directories. . . . .	34	51. Entering commands . . . . .	423
9. Options for controlling backup, archive, image, system object, and system state processing . . . . .	35	52. Commands . . . . .	423
10. Wildcard and other special characters. . . . .	38	53. Wildcard characters . . . . .	431
11. Specifying a drive specification using wildcards . . . . .	39	54. Archive command: Related options . . . . .	434
12. Using wildcard characters with include and exclude patterns . . . . .	39	55. Backup Group command: Related options . . . . .	441
13. Options for controlling compression, encryption, and adaptive subfile backup processing . . . . .	40	56. Backup Image command: Related options . . . . .	445
14. Using UNC names with include and exclude patterns. . . . .	43	57. Backup NAS command: Related options . . . . .	448
15. Required user security rights for Tivoli Storage Manager backup and restore services . . . . .	46	58. Backup WAS command: Related options . . . . .	455
16. Working with your files using the Tivoli Storage Manager GUI . . . . .	54	59. Backup VM command: Related options . . . . .	457
17. Backup: Primary tasks . . . . .	57	60. Delete Archive command: Related options . . . . .	463
18. Planning your backups. . . . .	57	61. Delete Backup command: Related options . . . . .	465
19. Maximum file size . . . . .	64	62. Delete Filespace command: Related options . . . . .	469
20. Command line backup examples . . . . .	73	63. Delete Group command: Related options . . . . .	470
21. Comparing incremental image backup methods . . . . .	86	64. Expire command: Related options. . . . .	472
22. NAS options and commands. . . . .	90	65. Incremental command: Related options . . . . .	480
23. Client command line informational messages . . . . .	96	66. Query Archive command: Related options . . . . .	488
24. UNC examples . . . . .	101	67. Query ASR command: Related options . . . . .	490
25. Restore: Primary tasks . . . . .	105	68. Query Backup command: Related options . . . . .	491
26. Command line restore examples . . . . .	107	69. Query Backupset command: Related options . . . . .	495
27. NAS options and commands . . . . .	129	70. Query Complusedb command: Related options . . . . .	497
28. Archive and retrieve: Primary tasks . . . . .	137	71. Query Eventlog command: Related options . . . . .	498
29. Command line archive examples . . . . .	139	72. Query Filespace command: Related options . . . . .	499
30. Command line examples of retrieving archives . . . . .	143	73. Query Group command: Related options . . . . .	501
31. Client return codes and their meanings . . . . .	151	74. Query Image command: Related options . . . . .	503
32. Default values in the standard management class . . . . .	155	75. Query Mgmtclass command: Related options . . . . .	506
33. TCP/IP options. . . . .	166	76. Query Node command: Related options . . . . .	507
34. Named Pipes communication option . . . . .	167	77. Query Options command: Related options . . . . .	508
35. Shared memory communication options . . . . .	167	78. Query Registry command: Related options . . . . .	509
36. node options. . . . .	167	79. Query Sysfiles command: Related options . . . . .	513
37. Backup and archive processing options . . . . .	168	80. Query Systeminfo command: Related options . . . . .	515
38. Restore and retrieve processing options . . . . .	176	81. Query Systemobject command: Related options . . . . .	516
39. Scheduling options. . . . .	179	82. Query Systemstate command: Related options . . . . .	517
40. Format and language options . . . . .	180	83. Query WAS command: Related options . . . . .	518
41. Command processing options . . . . .	181	84. Query VM command: Related options . . . . .	519
42. Authorization options. . . . .	181	85. Query WMI command: Related options . . . . .	520
43. Error processing options . . . . .	181	86. Restore command: Related options . . . . .	525
		87. Restore ASR command: Related options . . . . .	528
		88. Restore Backupset command: Related options . . . . .	530
		89. Restore Complusedb command: Related options . . . . .	532
		90. Restore Eventlog command: Related options . . . . .	533
		91. Restore Group command: Related options . . . . .	534
		92. Restore Image command: Related options . . . . .	538
		93. Restore NAS command: Related options . . . . .	541
		94. Restore Registry command: Related options . . . . .	542

95. Restore Sysfiles command: Related options	544	99. Restore WMI command: Related options	552
96. Restore Systemobject command: Related options . . . . .	545	100. Retrieve command: Related options . . . . .	554
97. Restore Systemstate command: Related options . . . . .	547	101. Schedule command: Related options . . . . .	557
98. Restore WAS command: Related options	550	102. Selective command: Related options . . . . .	561
		103. Scheduled client event status codes . . . . .	581

---

## About this book

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

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

- A *server program* that permits systems to perform as a backup and archive server for distributed workstations and file servers. See “Related publications” on page xiv 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 xiv 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.
- On Windows® platforms, an *ODBC driver* is available that allows you to use a relational database product such as Lotus® Approach to query the database and display the results. Also see *IBM Tivoli Storage Manager for Windows Administrator’s Guide*, SC32-0121.
- A *Web backup-archive client* that permits an authorized administrator, help desk person, or end user to perform backup, restore, archive, and retrieve services using a Web browser on a remote system. See “Starting a Web client session” on page 51 for more information.

---

## 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 15.

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

Unless otherwise specified, references to Windows refer to all supported Windows operating systems.

---

## IBM Tivoli Storage Manager Web site

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

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

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

---

## Conventions used in this book

This book uses the following typographical conventions:

*Table 1. Typographical conventions*

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

---

## Reading syntax diagrams

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

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

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

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

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

**Abbreviations:**

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

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

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



**Symbols:**

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

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

**Variables:**

Italicized lowercase items (*var\_name*) denote variables.

In this example, you can specify a *var\_name* when you enter the KEYWORD command.



**Repetition:**


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



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



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



**Notes:**

- 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.

Syntax diagram description	Example
<p><b>Optional choice:</b></p> <p>When an item is below the line, that item is optional. In the first example, you can choose A or nothing at all.</p> <p>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.</p>	
<p><b>Defaults:</b></p> <p>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.</p> <p>In this example, A is the default. You can override A by choosing B or C. You can also specify the default explicitly.</p>	
<p><b>Repeatable choices:</b></p> <p>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.</p> <p>In this example, you can choose any combination of A, B, or C.</p>	
<p><b>Syntax fragments:</b></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><b>The fragment name:</b></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
<i>Client publications</i>	
IBM Tivoli Storage Manager Messages	SC32-0140
IBM Tivoli Storage Manager for NetWare Backup-Archive Clients Installation and User's Guide	SC32-0144
IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide	SC32-0145
IBM Tivoli Storage Manager for Macintosh Backup-Archive Clients Installation and User's Guide	SC32-0143
IBM Tivoli Storage Manager Using the Application Programming Interface	SC32-0147



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

Publication title	Order number
<i>Server publications</i>	
<i>IBM Tivoli Storage Manager for AIX Installation Guide</i>	SC32-0134
<i>IBM Tivoli Storage Manager for AIX Administrator's Guide</i>	SC32-0117
<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 Sun Solaris Administrator's Guide</i>	SC32-0120
<i>IBM Tivoli Storage Manager for Windows Installation Guide</i>	SC32-0138
<i>IBM Tivoli Storage Manager for Windows Administrator's Reference</i>	SC32-0127
<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, 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](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 of 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.

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

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

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

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

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

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

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

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

## Reporting a problem

Have the following information ready when you report a problem:

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

## Internet

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

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

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

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

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

A newsgroup, *listserv@marist.edu*, is implemented by a third party. IBM supports this newsgroup on a best-effort basis only. See “Online forums” on page 55 for more information.



---

## Summary of changes for Tivoli Storage Manager

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

---

### Technical changes for Version 5.5 - November 2007

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

#### Open file support and online image backup are now available for Windows 64-bit operating systems

Open file support and online image backup are now supported on Windows 64-bit operating systems, and Volume Shadowcopy Service (VSS) is now a supported snapshot provider on both 32-bit and 64-bit Windows versions, excluding Windows XP.

With this support, the *fileleveltype* and *imagetype* options are being replaced by *snapshotproviderfs* and *snapshotproviderimage*, respectively.

See the following sections for information about using these new options:

- “Snapshotproviderfs” on page 373
- “Snapshotproviderimage” on page 374

#### VMware Consolidated Backup

VMware Consolidated Backup allows you to back up multiple virtual machines to be offloaded to a dedicated physical host. VMware Consolidated Backup also allows backups in a LAN-free environment.

See the following sections for more information about using VMware Consolidated Backup:

- “Using VMware Consolidated Backup” on page 92
- “Vmchost” on page 412
- “Vmcpw” on page 413
- “Vmcuser” on page 414
- “Vmlist” on page 415
- “Backup VM” on page 457
- “Query VM” on page 519

#### Client-node proxy support

Backup time can be reduced 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. 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 the following sections for more information about using *asnodename*:

- “Backup with client-node proxy support” on page 76
- “Archiving data with client node proxy” on page 140

- “Asnodename” on page 194

#### **Transparent encryption support**

Support has been added for transparent encryption, using the *encryptkeygenerate* option. Refer to “Encryptkey” on page 251 for more information about this option.

#### **Secure socket layer (SSL) support**

Support has been added for Secure socket layer (SSL), which allows secure communications between the Tivoli Storage Manager client and server. Refer to “Configuring Tivoli Storage Manager client/server communication with Secure Socket Layer” on page 26 and “Ssl” on page 380 for more information.

#### **TCP/IP Version 6 support**

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

#### **Auditlogging and auditlogname options**

The *auditlogging* option generates an audit log with an entry for each file that is processed during an incremental, selective, archive, restore, or retrieve operation. Refer to “Auditlogging” on page 197 for more information about this option.

The *auditlogname* option specifies the path and file name where you want to store audit log information. This option applies when audit logging is enabled. Refer to “Auditlogname” on page 200 for more information about this option.

#### **Maximum file path and file name extension**

The maximum number of bytes for a file name and file path combined is 8440. However, the file name itself cannot exceed 256 bytes and the path leading to the file cannot exceed 8184 bytes. Refer to Table 7 on page 18 for detailed information.

#### **System state and system services**

All system services components are now backed up as part of system state. This eliminates the need to schedule a separate backup for system services or specifying systemservices in the domain. You will not be able to back up individual system services components; however, individual components can be restored. Refer to these sections for more information: “Backup Systemstate” on page 453 and “Restore Systemstate” on page 547 for more information.

---

## **Technical changes for Version 5.4 - January 2007**

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

#### **Registry format and location changes**

The format of the Tivoli Storage Manager passwords stored in the Windows Registry and the registry location where passwords are stored has changed. Your existing registry keys 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, “Step 4: Creating a generic service resource for failover” on page 613 and “Step 5: Creating a generic service resource for failover” on page 618 for more information about the changes.

### **System dllcache directory now in the boot partition**

The system dllcache directory is now included in the boot partition backup of Windows 2003 and Windows Vista systems. See “Backing up Windows system state” on page 80 for more information.

### **Support for *srvprepostscheddisabled*, *srvprepostsnapdisabled*, and *schedrestretrdisabled* options**

Three new options: *srvprepostscheddisabled*, *srvprepostsnapdisabled*, 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 “Srvprepostscheddisabled” on page 378, “Srvprepostsnapdisabled” on page 379, and “Schedrestretrdisabled” on page 357 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 305 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 238 for information about how to use this option.

### **Tivoli Storage Manager Express upgrade to Tivoli Storage Manager enterprise**

Tivoli Storage Manager Express users can now upgrade to Tivoli Storage Manager enterprise. This upgrade allows you to take advantage of one or more of the advanced features available in Tivoli Storage Manager. After upgrading, you can either continue using the Tivoli Storage Manager Express interface (*clientview=express*), or you can use the standard Tivoli Storage Manager interface, which contains the enhancements necessary to support the backup set functionality found in Tivoli Storage Manager Express. See “Backupsetname” on page 204 and “Clientview” on page 209 for more information about the enhancements.

### **Windows Vista support**

Tivoli Storage Manager supports backing up and restoring system files on Windows Vista. See “Backup Systemstate” on page 453 and “Restore Systemstate” on page 547 for more information.

ASR is not supported at this time.

### **Automatic propagation of permissions to child objects upon restore**

Tivoli Storage Manager provides the ability to back up and restore Windows 2003 system state and system services using the Backup Operator account. See “Considerations for client services” on page 45 for more information.

### **Dropped support for Windows 2000**

Beginning with this release of Tivoli Storage Manager, Windows 2000 is no longer supported. The supported Windows operating systems are:

- Windows XP
- Windows 2003
- Windows 2003 R2
- Windows Vista

### Disabled query diskinfo command

Beginning with this release of Tivoli Storage Manager, the **query diskinfo** command is no longer available.

### Removed support for server-free

Beginning with this release of Tivoli Storage Manager, server-free support is no longer available.

---

## Technical changes for Version 5.3 - March 2006

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

### NetApp CIFS Share Definitions

The Tivoli Storage Manager Windows Client now supports backing up NetApp CIFS Share Definitions, including share permissions. See “Backing up NetApp CIFS share definitions” on page 95 for more information.

---

## 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 163, “Enablearchiveretentionprotection” on page 247 and “Set Event” on page 566 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 57, and Chapter 10, “Using commands,” on page 423 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 GUI directory tree shows detailed information of included and excluded objects.

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

See Chapter 2, “Configuring Tivoli Storage Manager,” on page 15, Chapter 4, “Backing up your data,” on page 57, Chapter 6, “Archiving and retrieving your data,” on page 137, and Chapter 10, “Using commands,” on page 423 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 145 and Chapter 10, "Using commands," on page 423 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, **dsmtree**, 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.

#### **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 75 for more information.

#### **LVSA and OFS support**

Open File Support (OFS) is now available on Windows 2003 32-bit operating systems. See "Configuring Open File Support (OFS)" on page 29 for more information about installing and configuring LVSA and OFS.

---

## **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 GUI to the Tivoli Storage Manager and Tivoli Home Pages**

From the backup-archive client 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 165 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 15, Chapter 9, “Using processing options,” on page 165, and Chapter 10, “Using commands,” on page 423 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 57, and Chapter 10, “Using commands,” on page 423 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 Windows and the backup-archive client for Windows. The server and the client must be configured for the shared memory protocol, and they must be on the same server system. 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 Windows storage agent for communication with the Windows backup-archive client.

### **Journal-based backup enhancements**

Multiple journal-based backup sessions are now possible. See Chapter 4, “Backing up your data,” on page 57 and Chapter 10, “Using commands,” on page 423 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 GUI directory tree shows detailed information of included and excluded objects.

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

See Chapter 2, “Configuring Tivoli Storage Manager,” on page 15, Chapter 4, “Backing up your data,” on page 57, Chapter 6, “Archiving and retrieving your data,” on page 137, and Chapter 10, “Using commands,” on page 423 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 145 and Chapter 10, “Using commands,” on page 423 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.

### Tivoli Storage Manager Administration Center

The hyperlink from the new Java based Tivoli Storage Manager Administration Center to a Tivoli Storage Manager client system provides the administrator ID and encrypted password to the Web client through a Java portlet. This enables the launching of the Web client 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 20 and the *IBM Tivoli Storage Manager Administrator’s Guide* for your operating system, for more information.

### Web client enhancements

You can now perform the following functions from 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.

### 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 75 for more information.

### Single drive support for Open File Support (OFS) or online image backups

You can now perform OFS or online image backups on systems with a single NTFS-based C: drive. You can also easily verify if the LVSA is functioning properly or find the root cause of any failure by checking the Windows event log, without turning on any trace facilities. See Chapter 9, “Using processing options,” on page 165 and Chapter 10, “Using commands,” on page 423 for more information.

### Optimized option default values

Table 3. Optimized option default values

Option	Change	Page
<i>diskbuffsize</i>	Allows you to specify I/O buffer size (in kilobytes) that the client can use to optimize backup, archive, or HSM client performance.  <b>Recommendation:</b> Use the default value instead of specifying the <i>diskbuffsize</i> option.	237

Table 3. Optimized option default values (continued)

Option	Change	Page
<i>largecommbuffers</i>	<p>This option has been replaced by the <i>diskbuffsize</i> option. At this time, <i>largecommbuffers</i> will continue to be accepted by the Tivoli Storage Manager client in order to ease the transition to the new option. However, the value specified by <i>largecommbuffers</i> will be ignored in favor of the <i>diskbuffsize</i> setting.</p> <p><b>Recommendation:</b> Discontinue the use of <i>largecommbuffers</i> because future releases of Tivoli Storage Manager might not accept this option.</p>	237
<i>tcpadminport</i>	<p>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.</p>	388
<i>tcpbuffsize</i>	<p>The default value was changed from 31 kilobytes to 32 kilobytes.</p>	389
<i>tcpnodelay</i>	<p>The default value was changed from <i>no</i> to <i>yes</i>. <i>tcpnodelay yes</i> disables the TCP/IP Nagle algorithm. This algorithm is used to reduce the number of small segments sent across the network, but in some environments this might negatively impact Tivoli Storage Manager performance.</p> <p><b>Recommendation:</b> Use the default of <i>yes</i>, unless you fully understand the effects of the TCP Nagle algorithm on network transmissions and how its use affects the performance of Tivoli Storage Manager in your environment.</p>	393

Table 3. Optimized option default values (continued)

Option	Change	Page
<i>tcpwindowsize</i>	The default value was changed from 32 kilobytes to 63 kilobytes.	396
<i>txnbytelimit</i>	The default value was changed from 2048 kilobytes to 25600 kilobytes.	403



---

## 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 backup-archive client graphical user interface (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 Version 5.5 from a previous version.

#### Upgrade path for clients and servers

As part of a migration plan from Tivoli Storage Manager Version 5.4 to Tivoli Storage Manager Version 5.5, Tivoli Storage Manager clients and servers can be upgraded at different times. This configuration is supported as long as the older version is supported. To help prevent disruption to your backup and archive activities during the migration, follow these guidelines:

- A Tivoli Storage Manager Version 5.4 client can perform backup, restore, archive, and retrieve functions to a Tivoli Storage Manager Version 5.5 server.
- A Tivoli Storage Manager Version 5.5 client can perform backup, restore, archive, retrieve, and query functions to a Tivoli Storage Manager Version 5.4 server.
- If you back up or archive data from a Tivoli Storage Manager Version 5.3 or Version 5.4 client to any Tivoli Storage Manager server, you can restore or retrieve that data using a Tivoli Storage Manager Version 5.5 client.
- If you back up or archive data from a Tivoli Storage Manager Version 5.5 client, you can restore or retrieve that data using a Tivoli Storage Manager Version 5.4 client.
- All administrative clients can administer Tivoli Storage Manager Version 5.4 and Version 5.5 servers.
- If you are upgrading from client version 5.2.0.x and online image backup or Open File Support is installed, then you must uninstall the 5.2.0.x client prior to installing the new client package. Otherwise subsequent online image or Open File Support backup operations will not function correctly. If you are upgrading from a client version earlier than 5.2, or from client version 5.2.2.0 or higher, then the upgrade can be performed without first uninstalling the existing client.
- If you are migrating from Tivoli Storage Manager Version 5.3, the format of the journal database is changed from past client versions. The journal will be reset after starting the Tivoli Storage Manager 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 passwords stored in the Windows Registry has changed. If you install the Version 5.4.0 or higher IBM Tivoli Storage Manager API client or Tivoli Data Protection, which is using the Tivoli

Storage Manager API, and you attempt to use the Version 5.3 or earlier IBM Tivoli Storage Manager Backup-Archive client, you will be unable to start the Tivoli Storage Manager Backup-Archive Client Scheduler service or client acceptor daemon service.

Your existing registry will be migrated to the new format during the first session with a Tivoli Storage Manager server after you install the Tivoli Storage Manager Version 5.4.0 or higher client. Once the registry keys have been migrated, previous versions of the Tivoli Storage Manager client will not process the updated keys 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 keys and the older client:

1. Before installing the new version, export and save the following registry subtree: `\HKLM\SOFTWARE\IBM\ADSM\CurrentVersion\BackupClient\Nodes`.
  2. After installing the new version, run a backup and enter the encryption key password when prompted.
  3. 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.
  4. If you need to go back to an older version, follow these steps:
    - a. Uninstall the client.
    - b. Restore the saved registry subtree.
    - c. Install the older version.
- The registry location where passwords are stored has changed in Version 5.4, so the Generic Resource Service for cluster failover must be updated to correct the Registry Replication value. Refer to “Step 4: Creating a generic service resource for failover” on page 613 and “Step 5: Creating a generic service resource for failover” on page 618 for more information.
  - The method for processing system state data has changed in Tivoli Storage Manager Version 5.5 such that system state (and system service) backup from prior clients is supported but no longer recommended. When you use the Tivoli Storage Manager Version 5.5 client you will generate new system state backups using the new methods. You *cannot* perform the following operations:
    - Generate a backup set with system state data. If you use the system state data backed up with the Tivoli Storage Manager Version 5.5 client to generate a backup set, you must be connected to a Tivoli Storage Manager Version 5.3.6 or higher, Version 5.4.1 or higher, or Version 5.5.0 server.
    - Restore system state and system services file spaces that were backed up by a Tivoli Storage Manager Version 5.4.x or lower client.
    - Use a Tivoli Storage Manager Client prior to Version 5.5 to restore system state backed up by a Tivoli Storage Manager Client Version 5.5 or above.
    - The Windows client can be regressed from Tivoli Storage Manager Version 5.5 to Tivoli Storage Manager Version 5.4 without any impact, except that system state backed up by the Tivoli Storage Manager Version 5.5 client cannot be restored by the Tivoli Storage Manager Version 5.4 client. If the system had not yet been backed up by the Tivoli Storage Manager Version 5.5 client, but was still the version backed up at the Tivoli Storage Manager Version 5.4 level, the Tivoli Storage Manager Version 5.4 client would be able to restore the system state; The system state would not be restorable by the Tivoli Storage Manager Version 5.4 client if the system state had already been backed up by the Tivoli Storage Manager Version 5.5 client.



- |           – Specify `systemservices` in the domain statement (for example, `domain`
- |           `systemservices`).
- |           – Use the **backup systemservices** command.
- |           – Use the **restore systemservices** command in normal production or recovery
- |           scenarios. Instead, use **restore systemstate** <service name> to restore a
- |           particular system service.
- |           – Use the **query systemservices** command.
- |           – Use the **show systemservices** command.

## | **Considerations for migrating between processor architectures**

### | **(x32, x64, IA64)**

|           When migrating between the Tivoli Storage Manager processor architectures (x32,

|           x64, IA64), consider the following:

- |           • The restore of DFS links backed up from 32-bit Windows environments to 32-bit
- |           Windows environments, is not supported.
- |           • The restore of DFS links backed up from 64-bit Windows environments to 32-bit
- |           Windows environments, is not supported.
- |           • System state or system object data is incompatible between different processor
- |           architectures. Therefore, when migrating client data between different processor
- |           architectures, it is recommended that you avoid backing up the system object or
- |           system state data from one architecture into the same file spaces containing the
- |           system object or system state data from another architecture. To avoid
- |           commingling the data, you can do either of the following before the migration:
  - |           – Ask your Tivoli Storage Manager administrator to rename the existing system
  - |           object or system state file spaces. Once the migration is complete and you
  - |           have new backups of the system object or system state data, your Tivoli
  - |           Storage Manager administrator can delete the renamed file spaces. The file
  - |           spaces to rename include one or more of the following:
    - |           – SYSTEM OBJECT
    - |           – SYSTEM STATE
    - |           – SYSTEM SERVICES
    - |           – ASR

|           Ask your Tivoli Storage Manager administrator to delete the existing file

|           spaces (listed above) for your node.

- |           • Renamed file spaces remain on the server and managed as stabilized file spaces.
- |           *These file spaces contain all the original data, which can still be restored until the file*
- |           *spaces are deleted. You can find additional information in the IBM Tivoli Storage*
- |           *Manager for Windows Administrator's Guide.*
- |           • When an existing file space is renamed during Unicode conversion, any access
- |           rules defined for the file space remain applicable to the original file space. New
- |           access rules must be defined to apply to the new Unicode file space, if necessary.
- |           • When backing up files to a file space that is not Unicode-enabled, the
- |           Unicode-enabled client skips the files and directories with names containing
- |           characters from a code page other than the current locale.
- |           • If files and directories with names containing characters from a code page other
- |           than the current locale were previously *backed up with a client that was not*
- |           *Unicode-enabled*, Tivoli Storage Manager will *expire* them in the file space that is
- |           not Unicode-enabled. However, the Unicode-enabled client can back up or
- |           archive these files to a Unicode-enabled file space.

- When migrating from a client that is not Unicode-enabled to a Unicode client, file names with double-byte characters mixed with single-byte characters might be restored in mixed single-byte character set (SBCS) format. This is only a display issue.
- When a client performs a selective backup of an object and the original file space is renamed, the new Unicode-enabled file space contains only that object. Tivoli Storage Manager will back up all other directories and file spaces during the next full incremental backup.
- The Unicode-enabled client has a file and directory name size limitation of 504 bytes. Depending on the Unicode characters (2-byte, 3-byte, or 4-byte) used, this translates to a maximum of 126 to 252 characters.

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

## Additional migration information

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

To view the non-English online help from the Web Client applet, you must install the language versions of the help files on the agent, the system where the Tivoli Storage Manager Backup-Archive client was installed. If the language versions are not installed or are not available, the online help 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 the *IBM Tivoli Storage Manager for Windows Administrator's Reference*, SC32-0127.

See the `client_message.chg` file in the client package for a list of new and changed messages since the Tivoli Storage Manager Version 5.4 release.

## Upgrading Open File Support (OFS) or online image

Beginning with Tivoli Storage Manager Version 5.5, the OFS and online image installation features have been replaced with a single Logical Volume Snapshot Agent (LVSA) feature. This feature will not be selected by default unless you are upgrading from an installation where OFS or online image was being used. This feature allows you to install LVSA for use during both online image and OFS operations, but those features will not be automatically enabled. To enable OFS, the *snapshotproviderfs* option must be set in the `dsm.opt` file, and to enable online image, the *snapshotproviderimage* option must be set in the `dsm.opt` file.

Volume Shadowcopy Service (VSS) is also now supported for OFS and online image operations. You can enable VSS by setting the *snapshotproviderfs* and *snapshotproviderimage* options in the `dsm.opt` file. If you use VSS, you do not need to install LVSA.

The Setup wizard allows you to select **NONE**, **VSS**, or **LVSA** for each of the OFS and online image functions. If **LVSA** is selected and it is not already installed on your system, it will be installed.

If you are migrating from a previous version of the Tivoli Storage Manager client where you were using the LVSA for OFS or online image, and you decide during the installation to continue to use the LVSA, then you do not need to explicitly set the *snapshotproviderfs* or *snapshotproviderimage* options. Since you do not need to set these options, it is easier to install the new client on a large number of systems, because the *dsm.opt* file will not need to be updated to continue to use the OFS or online image functions.

---

## Client environment requirements

This section contains client environment information, Tivoli Storage Manager client components, and hardware and software requirements for the Windows XP, Windows Vista and Windows Server 2003 platforms.

**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>

## Client components

- Backup-archive client command line and GUI
- Administrative client command line
- Tivoli Storage Manager API (32-bit and 64-bit)
- Web backup-archive client
- Open Database Connectivity (ODBC) driver

### Notes:

1. The ODBC driver is installed with a separate package. Refer to the README.ODBC.TXT file for installation instructions. This file is available in the Tivoli Storage Manager ODBC Driver install package. Also see *IBM Tivoli Storage Manager for Windows Administrator's Guide*, SC32-0121.
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 Windows Administrator's Reference*, SC32-0127.

## Hardware requirements

- One of the following:
  - An Intel® Pentium® or compatible processor, or multiprocessor based computer
  - An Itanium® (IA64) processor or multiprocessor based computer
  - An AMD64/EM64T (x64) processor or multiprocessor based computer
- Memory: 128 MB
- A minimum screen resolution of 800 by 600 pixels is required to display the Tivoli Storage Manager Backup-Archive Client GUI and Web GUI.

The disk storage requirements for several installation options are listed below.

*Table 4. Windows disk storage requirements*

Installation type	Total MB (x32)	Total MB (IA64)	Total MB (x64)
Base typical install	92 MB	107 MB	85 MB
Base full install	100 MB	125 MB	100 MB
Chinese language (PRC)	11 MB	10 MB	10 MB
Chinese language (Taiwan)	10 MB	10 MB	10 MB
Czech language	11 MB	10 MB	11 MB
German language	11 MB	11 MB	11 MB
French language	11 MB	11 MB	12 MB
Hungarian language	13 MB	10 MB	11 MB
Italian language	11 MB	13 MB	13 MB
Japanese language	11 MB	11 MB	11 MB
Korean language	11 MB	11 MB	11 MB
Polish language	12 MB	10 MB	11 MB
Portuguese language (Brazil)	11 MB	11 MB	12 MB
Russian language	13 MB	10 MB	11 MB
Spanish language (traditional)	12 MB	12 MB	13 MB

The following hardware is required for Windows Vista:

- A computer with 1 gigahertz or higher processor clock speed is recommended; 800 megahertz minimum required (single or dual processor system)
- Intel, AMD, or VIA processors
- Memory: 512 MB of RAM or higher
- Disk space: 20 gigabytes with 15 gigabytes of free space
- DirectX 9 capable graphics card
- Windows Display Driver Model (WDDM) driver-capable graphics card
- A video adapter capable of supporting the Longhorn Display Driver Model (LDDM) drivers used in Windows Vista
- A DVD drive.

When selecting a "Custom" installation, the online installation instructions will display how much space is required for the selected installation options, however, this number might not be accurate. Refer to the above chart for necessary disk space needed for installation. The online installation instructions will also display (correctly) how much space is available on the hard disk.

## Software requirements

This client is supported on the following Windows versions:

- Windows XP Professional SP2 or higher
- Windows Server 2003 (all editions, 32-bit and 64-bit)
- Windows Server 2003 R2 (all editions, 32 and 64 bit)

- Windows Vista (all editions, 32-bit and 64-bit)

**Note:** Support for Windows Server 2003 R2 is the same as for Windows Server 2003. Thus the client does not provide support for Windows Server 2003 R2's new features.

## Communications methods

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

## Supported features

The following table shows which features are supported on the various Windows platforms.

*Table 5. Supported features on Windows platforms*

Features	XP 32-bit	XP 64-bit x64	2003 32-bit	2003 64-bit x64	2003 64-bit IA64	Vista 32-bit	Vista 64-bit x64
Journal-based backup	yes	yes	yes	yes	no	yes	yes
Online image backup	yes	yes	yes	yes	yes	yes	yes
Offline image backup	yes	yes	yes	yes	yes	yes	yes
System object support with APIs from previous versions	yes	yes	n/a	n/a	n/a	n/a	n/a
System state support with Volume Shadowcopy Services (VSS)	no	no	yes	yes	yes	yes	yes
LAN-free operations	yes	yes	yes	yes	yes	no	no

Table 5. Supported features on Windows platforms (continued)

Features	XP 32-bit	XP 64-bit x64	2003 32-bit	2003 64-bit x64	2003 64-bit IA64	Vista 32-bit	Vista 64-bit x64
Automated System Recovery (ASR)	yes	yes	yes	yes	yes	no	no
Open File Support (OFS)	yes	yes	yes	yes	yes	yes	yes

**Notes:**

1. Only LVSA is available for 64-bit x64 Windows XP platforms.
2. LVSA is not signed, therefore the LVSA driver is included with the Tivoli Storage Manager Client package, but it is not installed by default. If your organization does not permit using unsigned drivers, you can use VSS.

---

## 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
- 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 87 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

The Tivoli Storage Manager Windows client is included on the desktop client installation CD in the setup directory structure.

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

- Install directly from the CD.
- Create client images to install.
- Transfer installable files from the CD to a target workstation.

You can copy all of the clients to your server workstation so that client workstations can get the files from the `x:\tsmcli` directory. A sample command for Windows is:

- `xcopy h:\setup\* x:\ /s`

**Note:** All of the examples in this chapter use the h drive as the CD or mounted drive. Substitute h with the CD drive of your system.

---

## Installing the software on your workstation

**Attention:** You will need to reboot the system if any instance of the client is running during the install. Stop all instances of the Tivoli Storage Manager client (services, interactive clients, etc.) before performing the install. If the Logical Volume Snapshot Agent (LVSA) component is selected for install then a reboot will be required in order to install or update the `tsmlvsa.sys` filter driver.

Follow these steps to install the software on your Windows XP, Windows Vista, or Windows Server 2003 workstations.

1. Insert the CD that contains the Tivoli Storage Manager Windows client into your CD drive. If you have autorun enabled, the installation dialog should start when the CD loads. If the installation dialog does not start automatically, you can start it manually. Select **Run** from the Start menu and at the prompt, type:  
`x:\setup`  
Where *x* is your CD drive. Click **OK**.
2. Follow the instructions displayed on the screen.

**Note:** If files from a previous version are in the installation directory, the Windows installer presents these options: **Modify**, **Repair**, and **Remove**. To install a new version of the product, first remove the currently installed version using the **Remove** option. To add a component that was not initially installed, select the **Modify** option.

### Installation setup types:

There are two setup types:

- Choosing **Typical** installs the minimum necessary to provide normal backup and archive functions. This includes the Backup-Archive Client, the API Runtime files, and the Web Client.
- Choosing **Custom** takes you to the Custom Setup window. From this window, you can click on any program feature icon to modify that feature if it is not mandatory for installation. You can select from the following program features:
  - Backup-Archive Client
    - Backup-Archive Client GUI Files (Mandatory; cannot be deselected)
    - Backup-Archive Client Web Files (Mandatory; cannot be deselected)
    - Client API Runtime Files (Mandatory; cannot be deselected)
    - Client API SDK Files (Optional)
    - Administrative Client Command Line Files
    - Logical Volume Snapshot Agent (LVSA) (Optional; not enabled by default)
  - Language Support

**Attention:** The Tivoli Storage Manager client now makes use of language packs for non-English language support. Each supported language has its own installation package that must be installed in order to use Tivoli Storage Manager in a supported, non-English language. The Tivoli Storage Manager client is a prerequisite for installing a Tivoli Storage Manager Client Language Pack.

### Notes:

- a. The Backup-Archive Client, the API, and the Web Client are interdependent. If you select the Backup-Archive Client, you must also select the API. Similarly, if you select the Web client, you must also select the Backup-Archive Client and the API.

- b. The Backup-Archive Client component includes the client scheduler files. See “Starting the client scheduler” on page 23 for more information about the client scheduler.
- c. The installer displays the exact amount of disk space that is required for each program feature. Ensure that there is enough disk space on the destination drive for the files you choose to install. The installation program will not install to a destination drive with insufficient disk space.
- d. If you do not have a dsm.opt file, a setup wizard is launched automatically when you start the GUI. The setup wizard can help you configure an initial options file. See “Creating and modifying the client options file (required)” on page 15 for detailed information about configuring an initial options file.
- e. If you installed the Web client, see “Starting a Web client session” on page 51 for configuration information.
- f. For more information on using the ODBC driver, see the README.ODBC.TXT file that accompanies the ODBC driver package. For more information on the Administrative client, see *IBM Tivoli Storage Manager for Windows Administrator’s Reference*, SC32-0127.

## Silent installation

The Tivoli Storage Backup-Archive Client installation program supports silent, unattended installations. By placing a customized version of the **msiexec** command (which invokes the Microsoft® Software Installer) in a script or batch file, you can easily perform installations on multiple Windows machines. The following is a sample command to install the backup-archive command line client, GUI client, Web client, API, and Administrative command line client. You might need to customize this example to run correctly on your system. Note that while the command is physically spread across multiple lines on this page, it is one single command. Also note that this example should be modified when installing a Tivoli Storage Manager Backup-Archive Client Language Pack.

```
msiexec /i "Z:\tsm_images\TSM_BA_Client\IBM Tivoli Storage Manager Client.msi"
RebootYesNo="No" REBOOT="Suppress" ALLUSERS=1 INSTALLDIR="d:\program files\tivoli\
tsm" ADDLOCAL="BackupArchiveGUI,BackupArchiveWeb,ApiRuntime,AdministrativeCmd,LVSA",
TRANSFORMS=1033.mst /qn /! *v "c:\log.txt"
```

The command and its parameters are as follows:

### **msiexec**

Invokes the Microsoft Software Installer (MSI) program.

**/i** Installs the specified source package (replace with **/x** to uninstall the package).

**"Z:\tsm\_images\TSM\_BA\_Client\IBM Tivoli Storage Manager Client.msi"**

This is the complete path to the source package. **Z:** is the CD or network drive containing the installation image. You will need to substitute the appropriate drive letter for your system.

**RebootYesNo="No" REBOOT="Suppress"**

Under certain conditions, a system reboot might be necessary for the installation to complete successfully. This option causes the installation program to not reboot the system if circumstances would otherwise cause the reboot to occur. While this option is convenient, it should be used with caution because suppressing the reboot might cause the program to behave in an unpredictable manner. The most common reason that a reboot would be required is if the installation was an upgrade to an existing Tivoli Storage Manager client, and the installation was performed while the client



programs were running. Therefore, we recommend that you shut down all Tivoli Storage Manager client programs and services before performing the installation.

**ALLUSERS=1**

Specifies that the package is for all users. This option is required.

**INSTALLDIR="d:\program files\tivoli\tsm"**

Specifies the destination path. If you have already installed this product or a previous version of this product on your workstation, you should use the current installation directory as the destination path for this package.

**ADDLOCAL="BackupArchiveGUI,BackupArchiveWeb,ApiRuntime,**

**AdministrativeCmd,LVSA"**

Specifies the features to install. Specify all the components on a single line within quotes, separated by commas, with no spaces before or after the commas. The following backup-archive features are available:

Windows 32-bit client features	Windows 64-bit client features	Feature description
BackupArchiveGUI	BackupArchiveGUI	Graphical User Interface
BackupArchiveWeb	BackupArchiveWeb	Backup-archive Web client
ApiRuntime	Api32Runtime Api64Runtime	API Runtimes
ApiSdk	ApiSdk	API SDK
AdministrativeCmd	AdministrativeCmd	Administrative Command Line
LVSA	LVSA	Online image backup and OFS
Windows 32-bit client Language Pack Features	Windows 64-bit client Language Pack Features	Feature description
LanguageFiles	LanguageFiles	Language specific files

**TRANSFORMS=1033.mst**

Specifies which language transform to use. The following language transforms are available:

Transform	Language
1028.mst	CHT Traditional Chinese
1029.mst	CSY Czech
1031.mst	DEU German
1033.mst	ENG English
1034.mst	ESP Spanish
1036.mst	FRA French
1038.mst	HUN Hungarian
1040.mst	ITA Italian
1041.mst	JPN Japanese
1042.mst	KOR Korean
1045.mst	PLK Polish
1046.mst	PTB Portuguese
1049.mst	RUS Russian

Transform	Language
2052.mst	CHS Simplified Chinese

**/qn** Specifies to perform the installation silently.

**/l\*v "c:\log.txt"**

Specifies verbose logging and the name and location of the log file.

The installation process creates a Tivoli Storage Manager folder in the programs folder of the Windows Start menu. You can start Tivoli Storage Manager by clicking one of the icons in this folder.

You can also install a predefined (custom) dsm.opt file, using the following instructions.

- Installing a predefined (custom) dsm.opt file
  - Users who would like to install a predefined dsm.opt file should place a copy of the dsm.opt file in the `...\CONFIG` directory located within the install image, for example:
 

```
tsm_images\TSM_BA_Client\baclient\Program Files\Tivoli\TSM\config
```
  - The file must be named dsm.opt.
- The install program will copy the predefined dsm.opt file to the `..\BACLIENT` directory when BOTH of the following conditions are met:
  - dsm.opt does NOT exist in the `..\BACLIENT` directory. The install program will not copy over an existing dsm.opt file.
  - dsm.opt exists in the install image's `..\CONFIG` directory, as described above.

## Upgrading, modifying, or reinstalling the client

The reinstall will replace any previous version it finds in the installation directory. However, the reinstall will not replace any files that were not created during the install process, such as dsm.opt, dsmerror.log, dsmsched.log, and dsmwebcl.log. In other words, if you have already set up an options file for your client, you do not need to do this again just because you reinstalled the client. Instead, the newly reinstalled client will use these existing files.

**Note:** It is recommended that you stop Tivoli Storage Manager client services before performing the upgrade.

If you want to install the client files in a directory other than the one where you previously installed them, you should uninstall the previous version before installing the new version. To do this:

1. Open the Windows **Start** menu and click **Settings** → **Control Panel**, then choose **Add/Remove Programs**.
2. Select the Tivoli Storage Manager client from the list of installed programs.
3. Click the **Add/Remove** button and click **OK**. **For components that were installed using Version 4.1 or higher:** Click the **Next** button until you see a panel containing the **Modify**, **Repair**, and **Remove** options.
4. To remove all components, click **Remove** and then click the **Next** button.
5. To remove individual components, select **Modify** and then click the **Next** button. The Custom Setup panel displays: *Select the component or components you want to remove and select **This feature will not be available**.*
6. Complete and close the setup wizard.
7. To fix missing or corrupt files, shortcuts, and registry entries, select the **Repair** option.

**Note:** If you do not restart Windows prior to installing the new client, some of the new client's files might be removed the next time you reboot. See "Installing the software on your workstation" on page 9 for instructions on reinstalling.

## Installation troubleshooting

If you are upgrading from a previous version of Tivoli Storage Manager and there are Tivoli Storage Manager services running (Client Acceptor, Scheduler, etc.) the following error might be encountered during install:

```
Error 1303. The installer has insufficient privileges to access this directory:  
(Install Drive):\Program Files\Tivoli\TSM\baclient\plugins. The installation  
cannot continue. Log on as an administrator or contact your system administrator.
```

Once this error occurs, you must quit the install. After quitting the install, the previous version is no longer installed. Stop running Tivoli Storage Manager service(s) and retry the installation.

### Installing LVSA using the setup wizard

It is necessary to have the client installation files available to the machine when installing the LVSA using the setup wizard. This can be either CD or the expanded FTP image. The LVSA install using the GUI setup wizard makes use of the client installation file "IBM Tivoli Storage Manager Client.msi", which is included in the client installation files. The default location for the client installation files is C:\tsm\_images. This is a limitation at this time.

### Removing online image backup or open file support during the upgrade

If you currently have Online Images Support and Open File support installed and want to remove either or both when upgrading to the latest Tivoli Storage Manager Client, deselecting the features during the upgrade installation will not work. To remove either or both features, you will need to go to the **Utilities -> Setup Wizard** and remove them using the **Help me configure....** options.

If you receive a message during setup indicating that you do not have enough free space to install on a drive that is known to have sufficient free space, you might be able to proceed by updating an entry in the setup.ini file that is included with the install package. If your package is on CD, you will need to copy the entire package to a hard drive so that you can edit setup.ini. Once setup.ini is available on your hard drive, open it up in your favorite text editor and change the line:

```
DiskSpace=8000 ;DiskSpace requirement in KB
```

to

```
DiskSpace=0000 ;DiskSpace requirement in KB
```

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 6. 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

*Table 6. Tivoli Storage Manager server Installation Guides (continued)*

<b>Publication title</b>	<b>Order number</b>
<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` file used by the previous client installation is 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 this file for configuration information.

*Required* configuration tasks include the following:

Task	Page
Creating and modifying the client options file (required)	15
Registering your workstation with a server (required)	30

*Optional* configuration tasks include the following:

Task	Page
Creating a shared directory options file (optional)	17
Creating multiple options files	17
Setting environment variables	19
Configuring the Web client	20
Configuring the client scheduler	21
Configuring the journal engine service	28
Configuring online image backup support	28
Configuring Open File Support (OFS)	29

---

### Creating and modifying the client options file (required)

- 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
- From the client GUI main window, select **Edit** → **Preferences**.
- Use your favorite text editor.

Use the client options file to identify the Tivoli Storage Manager server to contact for services, and communication protocol necessary to establish communications with that server. The client options file can also include authorization options, backup and archive processing options, scheduling options, and other options that govern the client's behavior. When you run the Tivoli Storage Manager client GUI,

if it does not detect an options file, it will automatically launch the configuration wizard to help you create one. If you want to manually start the wizard later, perhaps to modify your options file, follow these steps:

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.

You can also modify your options file manually. The install process creates a sample client options file called `dsm.smp`. If you accepted the default installation path, this file is in the `Program Files\Tivoli\TSM\config` directory. The `dsm.smp` file is a generic configuration file that contains communication parameters that are normally used in a Windows environment. You can use this file as a template when customizing your own options file. First, copy the `dsm.smp` file to your `Program Files\Tivoli\TSM\baclient` directory and rename it `dsm.opt`. Open it in Notepad or your favorite plain text editor and use the following guidelines to customize it for your environment.

One of the options file's most important purposes is to specify the communication protocol necessary to establish communications between your backup client and the backup server. Use the *commmethod* option to specify a communication protocol. For example, to use the TCP/IP communication protocol, enter:

```
COMMMethod TCPIP
```

You must also specify the TCP/IP server address of the backup server you will connect to using the *tcpserveraddress* option. For example, you can enter the following:

```
TCPServeraddress dsmchost.endicott.ibm.com
```

You can also specify what you want to call your workstation by using the *nodename* option. When you register with the backup server during setup, this is the name the server will use for your workstation. If you do not specify a node name, the server will identify your workstation by its host name. You can type **hostname** at a command prompt to see your workstation's host name. For more information about registration requirements, see "Registering your workstation with a server (required)" on page 30.

You can also use the options file to specify authorization, backup and archive processing, and scheduling options. The options specified in your options file are in effect every time you start Tivoli Storage Manager from either the GUI or the command line. However, you can override options temporarily by using the command line.

Your backup server can override your options by using server defined and enforced client option sets. See Chapter 9, "Using processing options," on page 165 for more information about option processing.

**Notes:**

1. See "Client options reference" on page 190 for detailed information about the supported options for your platform.

2. See “Communication options” on page 166 for supported communication protocols for your Windows client.

Once you have created an options file, you can use one of the following methods to edit your options file from the GUI. Method 2 includes basic options only, while Method 1 allows more extensive tailoring (has a lot more options available).

- Method 1:
  1. Open the **Edit** menu and select **Preferences**.
  2. Make any necessary changes, then click **OK** to save those changes.
- Method 2:
  1. Open the **Utilities** menu and select **Setup Wizard**.
  2. Select the **Help me configure the TSM Backup Archive Client** checkbox.
  3. Click the **Next** button.
  4. Select the **Update my options file** radio button.

## Creating a shared directory options file (optional)

A server administrator can generate client options files in a shared directory. Windows XP, Windows Vista, and Windows Server 2003 clients can access the shared directory and use the files there to create their own client options file. For more information see *IBM Tivoli Storage Manager for Windows Installation Guide*, SC32-0138.

## Creating multiple options files

You can create multiple options files if you have to work with multiple servers, or find that you need multiple sets of parameters to perform backup or archive tasks.

Suppose you want to back up your files to one server (server a), and archive files to another (server b). Instead of editing the `dsm.opt` file each time you want to connect to a different server, create two options files. For example, create the options files `a.opt` for server a, and `b.opt` for server b. Then do one of the following:

1. From a command prompt, use the **copy** command to replace the `dsm.opt` file with the appropriate options file before you begin a backup or archive session. For example:

```
copy a.opt dsm.opt
dsm
```
2. Use the **optfile** option on the command line to specify the options file you want to use. For example:

```
dsm -optfile=b.opt
```
3. Define the `DSM_CONFIG` environment variable to specify the options file to use:

```
SET DSM_CONFIG=C:\Program Files\Tivoli\TSM\baclient\b.opt
```

If you are running from the command line, the `DSM_DIR` and `DSM_LOG` environment variables might also need to be configured as follows:

1. Define the `DSM_DIR` environment variable to point to the directory where all other executable files reside:

```
SET DSM_DIR=C:\Program Files\Tivoli\TSM\baclient
```
2. Define the `DSM_LOG` environment variable to point to the directory where `dsmerror.log` resides:

```
SET DSM_LOG=C:\Program Files\Tivoli\TSM\baclient
```

**Note:** The directory path where the client executable files are located must be included in the PATH environment variable or you must enter a fully qualified path.

## Setting options in the client options file

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

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

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

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

```
compression yes
nodename      client_a
```

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

```
verbose
ve
```

Follow these additional rules when entering options in your client 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.
- Enter one or more blank spaces between parameters.
- The maximum number of bytes for a file name and file path combined is 8440. However, the file name itself cannot exceed 256 bytes and the path leading to the file cannot exceed 8184 bytes. Furthermore, directory names (including the directory delimiter) within a path are limited to 256 bytes. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.

For a complete understanding of file path names and limits, refer to Table 7.

Table 7. File path and name limits

MBCS encoding	Path		File name
	New limit	Legacy limit	Limit
1	8184	1016	256
2	4092	508	128
3	2728	338	85

In the table, MBCS encoding has these meanings:

### Basic Latin

For example, standard US English characters, numbers, symbols, and control characters that are traditionally represented in 7-bit ASCII have a 1:1 ratio of bytes to characters.



### Latin extensions

For example, Latin characters with tildes, grave or acute accents, and so on, as well as Greek, Coptic, Cyrillic, Armenian, Hebrew and Arabic characters typically have a 2:1 ratio of bytes to characters.

### Chinese, Japanese, Korean, Vietnamese

These characters and other East Asian language characters typically have a 3:1 ratio of bytes to characters.

If you update the client options file while a session is active, you must restart the session to pick up the changes. If you use the Setup wizard to make changes, the changes are effective immediately. If you are not using the TSM Client Acceptor to manage the scheduler, you will need to restart the scheduler as well.

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 508 for more information.

---

## Setting 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 if you need to run in either of the following environments:

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 318 for more information.

You need to set four environment variables:

**PATH** This is the default search path the operating system uses to locate executable files. Set this to include the fully qualified paths of the client installation directories.

### DSM\_CONFIG

Set this environment variable to the fully qualified path and file name of the client options file.

### DSM\_DIR

Set this environment variable to the directory where the client message file dsc\*.txt is located.

### DSM\_LOG

Set this environment variable to the directory where the log files should reside.

Ensure that the environment variables meet the following guidelines:

1. Include the directory where the executable files (for example, dsm.exe) reside in the current PATH environment variable. If you accepted the default installation directory using the C: drive, you can set this from a command prompt by typing:

```
SET PATH=%PATH%;C:\Program Files\Tivoli\TSM\baclient
```

2. Specify the fully-qualified path name of your client options file (dsm.opt) using the DSM\_CONFIG environment variable:  

```
SET DSM_CONFIG=C:\Program Files\Tivoli\TSM\baclient\dsm.opt
```
3. Define the DSM\_DIR environment variable to point to the directory where the Tivoli Storage Manager client message file dsc\*.txt is located:  

```
SET DSM_DIR=C:\Program Files\Tivoli\TSM\baclient
```

---

## Configuring the Web client

You can use the GUI Setup wizard or command line to install and configure the Web client.

**To install and configure the Web client using the GUI Setup wizard, perform the following steps:**

1. From the backup-archive client GUI's main window, open the **Utilities** menu and select **Setup Wizard**.
2. Select the **Help me configure the TSM Web Client** check box.
3. Click the **Next** button and then follow the instructions on the screen.

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

1. Ensure that you specify *passwordaccess generate* in the client options file (dsm.opt). For more information on *passwordaccess*, see "Passwordaccess" on page 320.

2. Install the Client Acceptor Service by entering the following command:  

```
dsmcutil install cad /name:"TSM CAD" /node:nodename /password:password /autostart:yes
```

where *nodename* and *password* are your Tivoli Storage Manager node name and password. TSM CAD is an example name, you can use any name you want. The default name is TSM Client Acceptor.

3. Install the Remote Client Agent Service by entering the following command:  

```
dsmcutil install remoteagent /name:"TSM AGENT" /node:nodename /password:password /partnername:"TSM CAD"
```

where *nodename* and *password* are your Tivoli Storage Manager node name and password. TSM AGENT is an example name, you can use any name as long as it is different from the CAD name. The default name is TSM Remote Client Agent. The /partnername option value must match the name of the CAD service. The default is name is TSM Client Acceptor.

4. Start the client acceptor service (CAD) by entering: net start "TSM CAD" on the command line, or do the following:
  - a. Open the Windows **Start** menu and select **Settings**→ **Control Panel**.
  - b. Double-click **Administrative Tools** and then double-click **Services**.
  - c. In the Services window, right-click **TSM CAD** and select **Start** from the pop-up menu.

The Tivoli Storage Manager Remote Client Agent service must not be started manually. It is automatically started by the Tivoli Storage Manager Client Acceptor service 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 also manages the Tivoli

Storage Manager scheduler. See Chapter 9, “Using processing options,” on page 165 for more information about these options.

The Administration Center, which is a component available through the Integrated Solutions Console, allows you to securely hyperlink to the client machine. The web client GUI is launched without the administrator needing to sign on again. In order for this to be fully secured, SSL needs to be enabled for the Integrated Solutions Console server. Refer to *IBM Tivoli Storage Manager for Windows Installation Guide* for information on configuring SSL for the Integrated Solutions Console, and for adding an administrator to the Administration Center.

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.

For more information about Storage Manager client services and their options, see Appendix B, “Using the Client Service Configuration Utility,” on page 585.

5. Add `MANAGEDServices webclient` to `dsm.opt`.
6. 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 273 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 1.5 or 1.4.x (where  $x \geq 1$  or higher). See “Starting a Web client session” on page 51 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, you can perform the following tasks:

- Display information about available schedules.
- Display information about work that the schedule has completed.
- Modify scheduling options in the client options file (`dsm.opt`). See “Scheduling options” on page 179 for more information.

See Chapter 7, “Automating tasks,” on page 145 for more information. See Appendix A, “Using the Tivoli Storage Manager central scheduler,” on page 571 for supplemental information about the Tivoli Storage Manager central scheduler.

The Tivoli Storage Manager Client Acceptor service (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 run immediately or the scheduler exits. The CAD restarts the scheduler when it is time to run the scheduled event. This reduces the number of background processes on your workstation and resolves memory retention problems that can occur when running the scheduler service without CAD management. **Recommendation:** Use the client acceptor daemon to manage the client scheduler. “Configuring the client scheduler” on page 21 for more information.

**Attention:** You cannot use the `dsmcad` for scheduling when you set the `sessioninitiation` option to `serveronly`. Refer to “Sessioninitiation” on page 361 for more information.

Use the `managedservices` option in your client options file (`dsm.opt`) to specify whether the CAD manages the scheduler. See “Managedservices” on page 302 for more information.

**Note:** See Appendix D, “Configuring the backup-archive client in a cluster server environment,” on page 609 for information on how to set up the Tivoli Storage Manager backup-archive client scheduler in a Microsoft Cluster Server (MSCS) or Veritas Cluster Server (VCS) environment with the `managedservices` option.

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

1. Select **Utilities** → **Setup Wizard** → **Help me configure the TSM Client Scheduler** from the Tivoli Storage Manager main window. The TSM Scheduler Wizard window appears.
2. Select the **Install a new or additional scheduler** task and press the **Next** button. The TSM Scheduler name and location window displays.
3. Specify the name of the Scheduler. Select the **Local Machine** and **Use the Client Acceptor daemon (CAD) to manage the schedule** options. Click the **Next** button. The **Select the names of the Web services** window displays.
4. Specify the name of the client acceptor service that you want to manage the scheduler. If the CAD is already installed for use by the Web client, select that CAD from the drop down list. Otherwise, type the name you want to give the CAD, for example, TSM Client Acceptor. Click the **Next** button.
5. Follow the instructions on the screen to complete the configuration of the options file name, the HTTP port the CAD uses, authentication information, service login options, log file names, and immediate start option.

**Notes:**

- a. If `sessioninitiation` option is set to `serveronly` in your client options file (`dsm.opt`), the client setup wizard and scheduler service might be unable to initiate authentication with the Tivoli Storage Manager server. To avoid this problem, ensure that the **Contact the TSM Server to validate password** checkbox on the TSM Authentication page is unchecked.
- b. In the Service login options window, select the **Automatically when Windows boots** option to have the service started automatically when Windows boots so that your schedules will be run.

#### Notes:

1. You can also use the **Scheduler Service Configuration** utility. The **Scheduler Service Configuration** utility must be run from an account that belongs to the Administrator/Domain Administrator group. You can start multiple client scheduler services on your system. For additional information, see Appendix B, "Using the Client Service Configuration Utility," on page 585.
2. For more information about scheduling options, changing the scheduling mode, specifying the TCP/IP address or port number, or running commands before or after a schedule, see "Scheduling options" on page 179.
3. If you are having the CAD manage your schedule, start the CAD service, but do not start the scheduler service. The scheduler service will be started and stopped automatically by the CAD service as needed. If you are not managing the schedule with the CAD, then you can use the Services Control Panel or the net start command to start the Scheduler service.

**Note:** Running the client scheduler on the command line is not recommended since it does not run the scheduler as a background service.

4. To setup a schedule to back up system objects, see "Setting up a schedule to back up system objects (Windows XP)" on page 588 for more information.
5. See Chapter 7, "Automating tasks," on page 145 for information about the following tasks:
  - Starting the client scheduler
  - Modifying scheduling options in the client options file
  - Mobile dial-up support
  - Displaying information about available schedules
  - Displaying information about work that the schedule has completed

---

## Starting the client scheduler

To start the Tivoli Storage Manager client scheduler, use the Services Control Panel or the net start command.

**Note:** Running the client scheduler on the command line is not recommended since it does not run the scheduler as a background service.

When you start the client scheduler, it runs continuously until you close the window, shut down your system, or log out of your system. If you are running the Scheduler Service, the scheduler runs until the system is shutdown or you explicitly stop it using the services control panel. For information on scheduling options, see Chapter 9, "Using processing options," on page 165.

See the following section for information about starting the client scheduler:

- "Setting the client scheduler process to run as a background task and start automatically at boot time" on page 575

---

## 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 394 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 422 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 “*Tcpadminport*” on page 388 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 GUI, and Web client GUI will still attempt to initiate sessions, but are blocked by the Tivoli Storage Manager server for nodes that have the *sessioninitiation* option set to *serveronly*.
4. When installing Tivoli Storage Manager scheduler using the setup wizard or `dsmcutil`, and the Tivoli Storage Manager server is behind a firewall, the node password will not get stored on the client machine. As a result, the scheduler service might be unable to authenticate to the server when the server contacts the client to run a schedule. In this case, you can run the scheduler from the command line (`dsmc schedule`), wait until a scheduled operation starts, and enter the password for your node when prompted. After you enter the password for your node, restart the scheduler service. You can also use the following **dsmcutil** command to write the password into the registry:  
`dsmcutil updatepw /node:nnn /password:ppp /validate:no`  
If *sessioninitiation* option is set to *serveronly* in your client options file (`dsm.opt`), the client setup wizard and scheduler service will be unable to initiate authentication with the Tivoli Storage Manager server. To avoid this problem, when configuring the client scheduler using the setup wizard, ensure that the **Contact the TSM Server to validate password** checkbox on the TSM Authentication page is unchecked. See “Configuring the client scheduler” on page 21 for more information.  
A similar problem can occur if an encryption key is required for backup operations. In this case, you can run the scheduler from the command line (`dsmc schedule`), wait until a scheduled backup starts, and enter the encryption key when prompted. After the password and encryption key are updated, you must restart the scheduler.
5. 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 the client scheduler to run a schedule. This can happen when the *passwordaccess* is set to generate and the Tivoli Storage Manager server is behind a firewall and the encrypted password cannot be locally stored before the scheduler is started. To correct this problem, you need to run the scheduler from the command line (`dsmc schedule`), wait until a scheduled operation starts, and enter the password for your node when prompted.

**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 23 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 (23) 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 361 for more information about the *sessioninitiation* option.

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

- In *prompted* mode the Tivoli Storage Manager server needs to contact the client. In order to do this, some software 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<sup>®</sup> server across a firewall.

---

## Configuring Tivoli Storage Manager client/server communication with Secure Socket Layer

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

The following client components support SSL:

- Command-line client
- Administrative command-line client
- Backup-archive client GUI
- Client API

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

In order to enable SSL communication, you need to perform the following actions:

1. Obtain the Tivoli Storage Manager server certificate.
2. Create the local key database, if it has not already been created.
3. Add the server certificate to the local key database.
4. Add the *SSL* option to the client options file.
5. Select the correct server TCP port for SSL communications.

The Tivoli Storage Manager Server certificate has a fixed file name of *cert.arm* and it is stored on the server machine in the server instance directory, for example: *c:\program files\tivoli\tsm\server1\cert.arm*. If this file does not exist when server options *SSLTCPPORT* or *SSLTCPADMINPORT* are specified, it will be created and a self-signed SSL certificate will be generated and stored in it.



Each Tivoli Storage Manager server generates its own certificate. In order to set up the SSL connection to a particular server, you need to obtain its certificate (in the form of the cert.arm file) from the server administrator, and import it into the local client key database.

The **gsk7capicmd** command is provided by Global Security Kit (GSKit). Tivoli Storage Manager automatically installs GSKit in \Program Files\IBM\gsk7. However, if GSKit has been installed prior to Tivoli Storage Manager installation, it is possible that it is in some other location. You might have to obtain the GSKit location from the following registry key:

```
HKLM\SOFTWARE\IBM\GSK7\CurrentVersion\InstallPath
```

Follow these steps to set up the server certificate on the client:

1. Open a command window and change the directory to your Tivoli Storage Manager client directory, for example: `cd "c:\Program Files\Tivoli\TSM\baclient"`
2. Add the GSKit binary and library paths to the PATH environment variable, for example: `set PATH=C:\Program Files\IBM\gsk7\bin;C:\Program Files\IBM\gsk7\lib;%PATH%`
3. Create the local key database, if it does not exist: `gsk7capicmd -keydb -create -db dsmcert.kdb -pw <password> -stash`
4. Import the server certificate file cert.arm (assuming it is located in the current directory): `gsk7capicmd -cert -add -db dsmcert.kdb -pw <password> -label "TSM server <servername> self-signed key" -file cert.arm -format ascii -trust enable`

**Notes:**

1. An arbitrary password (provided by the user) is used to encrypt the key database. The password is automatically stored encrypted in the stash file (dsmcert.sth). The stash file is used by the Tivoli Storage Manager client to retrieve the key database password.
2. More than one server certificate can be added to the client key database so that the client can connect to different servers. Different certificates must have different labels. The label names are not important, but meaningful names should be used.
3. If you do not run the preceding commands from the Tivoli Storage Manager client directory, you need to copy dsmcert.kdb and dsmcert.sth into that directory.
4. On a 64-bit platform, GSKit is installed in \Program Files\IBM\gsk7\_64, although it can vary, so you should check the registry key. For a 64-bit platform, use the **gsk7capicmd\_64** command.

After the server certificate has been added to the client key database, add the *SSL yes* option to the client options file and update the value of the TCPSPORT option. It is important to understand that the server is normally set up for SSL connections on a different port. In other words, two ports are opened on the server: one port accepts regular non-SSL client connections, and another port accepts SSL connections only. You cannot connect to a non-SSL port with an SSL-enabled client, and vice versa.

Refer to "Ssl" on page 380 for more information about the *SSL* option.

---

## Configuring the journal engine service

Journal-based backup is supported for all Windows clients, except for clients running on Windows Server 2003 for Itanium-based Systems. If you install the *journal engine service* 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 page 66 for more information.

Journal-Based Backup is enabled by installing and configuring the Tivoli Storage Manager Journal Service. The Tivoli Storage Manager Journal Service can be installed with the GUI Setup wizard or with the **dsmcutil** command. Basic Journal Service configuration can be done with the GUI Setup wizard, more advanced configuration can be done by editing the Journal Service configuration file, *tsmjbbd.ini*. See Appendix B, “Using the Client Service Configuration Utility,” on page 585 for more information about using the **dsmcutil** command to install the Tivoli Storage Manager Journal Service.

To install and configure this service using the client Java GUI setup wizard, perform the following steps:

1. From the main window, open the **Utilities** menu and select **Setup Wizard**.
2. Select the **Help me configure the TSM Journal Engine** check box.
3. Select the task you want to perform. You can install a new TSM Journal Engine, update a previously installed TSM Journal Engine, or remove a previously installed TSM Journal Engine from your system.
4. Complete each panel in the wizard and click the **Next** button to continue. To return to a previous panel, click the **Back** button. To display help information for a panel, click the **Help** button.

Journal service configuration settings are stored in the journal configuration file *tsmjbbd.ini*. This file can be installed and configured with the GUI setup wizard or edited manually. See Appendix C, “Journal service configuration,” on page 601 for detailed information about configuring this file.

---

## Configuring online image backup support

If the Tivoli Storage Manager online image feature is configured, Tivoli Storage Manager performs a snapshot-based image backup, during which the real volume is available to other system applications. A consistent image of the volume is maintained during the online image backup. To configure online image backup, perform the following steps:

1. Select **Utilities** → **Setup Wizard** from the Tivoli Storage Manager GUI main window. The Client Configuration Wizard panel appears.
2. Select **Help me configure Online Image Support** and click the **Next** button. The TSM Online Image Support Wizard panel appears.
3. Select the Snapshot Provider you want to use and then click the **Next** button. The providers you can select from are **VSS** or **LVSA**. To disable online image support, select **None**.
  - a. If you select **None** and the LVSA is installed, and the LVSA is not being used for OFS, the **Uninstall the Logical Volume Snapshot Agent (LVSA)** checkbox is enabled. If selected, the LVSA storage filter driver will be uninstalled.
4. Click the **Finish** button to complete the setup.

5. Complete each panel in the wizard and click the **Next** button to continue. To return to a previous panel, click the **Back** button. To display help information for a panel, click the **Help** button.
6. If you selected as the provider in 3 on page 28, restart the system if prompted.

See “Performing an image backup” on page 82 for detailed information about performing online image backups.

To set preferences for open file support, use the Include-Exclude tab on the Tivoli Storage Manager Preferences editor. You can set these options for all volumes or for individual volumes using the *include.fs* option: *snapshotproviderfs*, *presnapshotcmd*, *postsnapshotcmd*. Also, if you selected the LVSA as your snapshot provider, the following additional options can be set: *snapshotcachelocation*, *snapshotcachesize*, *snapshotfsidleretries*, *snapshotfsidlewait*.

See “Client options reference” on page 190 for more information about these options.

LVSA is not signed, therefore the LVSA driver is included with the Tivoli Storage Manager Client package, but it is not installed by default. If your organization does not permit using unsigned drivers, you can use VSS.

---

## Configuring Open File Support (OFS)

The OFS feature can be selected at installation time or installed later using the Tivoli Storage Manager client GUI setup wizard. By default, the OFS feature is not installed. The install program or the setup wizard will perform all the steps necessary to install, set up, and enable the OFS feature so that the next backup or archive operation will attempt to take advantage of the open file support. You can also use the install program or setup wizard to install or remove the OFS feature. If the LVSA is chosen as the snapshot provider, the installation or removal of this feature might require a machine reboot.

If the client is unable to create a snapshot; for example, if the *snapshotfsidleretries* number is not met or the *snapshotcachelocation* option is not valid, failover to a non-OFS backup will occur on the volumes where this problem exists. If the LVSA is your snapshot provider, information and error messages might also be logged to the Windows event log.

If the Tivoli Storage Manager Open File Support feature is configured, Tivoli Storage Manager performs a snapshot-based file-level operation, during which the real volume is available to other system applications. A consistent image of the volume is maintained during the operation. To configure OFS, perform the following steps:

1. Select **Utilities** → **Setup Wizard** from the Tivoli Storage Manager GUI main window. The Client Configuration Wizard panel appears.
2. Select **Help me configure Online Image Support** and click the **Next** button. The TSM Online Image Support Wizard panel appears.
3. Select the Snapshot Provider you want to use and then click the **Next** button. The providers you can select from are **VSS** or **LVSA**. To disable online image support, select **None**.

- a. If you select **None** and the LVSA is installed, and online image is not being used for OFS, the **Uninstall the Logical Volume Snapshot Agent (LVSA)** checkbox is enabled. If selected, the LVSA storage filter driver will be uninstalled.
4. Click the **Finish** button to complete the setup.
5. Complete each panel in the wizard and click the **Next** button to continue. To return to a previous panel, click the **Back** button. To display help information for a panel, click the **Help** button.
6. If you selected in 3 on page 29, restart the system if prompted.

Considerations:

- After installing, updating, or removing the OFS feature, you are prompted to restart your system, if LVSA was selected as the snapshot provider.
- To set preferences for open file support, use the Include-Exclude tab on the Tivoli Storage Manager Preferences editor. You can set these options for all volumes or for individual volumes using the *include.fs* option: *snapshotproviderfs*, *presnapshotcmd*, *postsnapshotcmd*. Also, if you selected the LVSA as your snapshot provider, the following additional options can be set: *snapshotcachelocation*, *snapshotcachesize*, *snapshotfsidleretries*, *snapshotfsidlewait*. See “Client options reference” on page 190 for more information about these options.

IBM has done extensive testing with the LVSA driver and is confident in its stability and function, and it is fully supported by IBM. IBM is actively working with Microsoft on a driver signing process, but the driver is currently not signed. Since LVSA is not signed, it was not tested as part of the Windows certification process. The LVSA driver is included with the Tivoli Storage Manager client package but is not installed by default. If your organization has a strict policy which does not allow non-signed drivers to be installed, then you can use VSS as your online image provider. Note that offline image backups can be performed without installing the LVSA driver.

---

## Registering your workstation with a server (required)

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 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.

---

## Creating an include-exclude list (optional)

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 153.

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.

**Attention:** There are some system files that you should exclude. See “Excluding system files” on page 37 for more information.

Specify your include-exclude list in your client options file (`dsm.opt`). The include-exclude list can also go into a separate file, which is referred to by the **inclexcl** option. The include-exclude statements are not case-sensitive.

The client options file, `dsm.opt`, must be in a non-Unicode format. However, if you are using a separate include-exclude file, it can be in Unicode or non-Unicode format. See “Inclexcl” on page 278 for instructions on creating a Unicode include-exclude file.

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

See “Inlexcl” on page 278 for important detailed information about specifying an include-exclude file using the *inlexcl* 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. See “Considerations for Unicode-enabled clients” on page 278 for instructions on creating a Unicode include-exclude file.

You can also create an include-exclude list manually by performing the following steps:

1. Determine your include and exclude requirements.
2. **Recommendation:** Group your include-exclude options together in your client options file.
3. Enter your *include* and *exclude* statements using the appropriate include-exclude options as described in “Using include-exclude options.” Tivoli Storage Manager evaluates all *exclude.dir* statements *first* (regardless of their position within the include-exclude list), and removes the excluded directories and files from the list of objects available for processing. 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 c:\test\includefile.txt
exclude c:\test\...\*
```

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

```
exclude c:\test\...\*
include c:\test\includefile.txt
```

4. Save the file and close it.
5. Restart your Tivoli Storage Manager client and the scheduler and client acceptor services 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 client options file (dsm.opt). See table references for more information about each option.
- A minimum include-exclude list that excludes system files.
- Information on how to exclude remotely accessed files.
- A list of supported wildcard characters.
- Examples of how you might use wildcard characters with *include* and *exclude* patterns.

## Excluding directories

Use *exclude.dir* statements to exclude all files and sub-directories in the specified directory from processing. Tivoli Storage Manager evaluates all *exclude.dir* statements *first* (regardless of their position within the include-exclude list), and removes the excluded directories and files from the list of objects available for processing. The *exclude.dir* statements override all include statements that match the pattern.

Table 8. Option for excluding directories

Option	Description	Page
<i>exclude.dir</i>	<p>Excludes a directory, its files, and all its subdirectories and their files from backup processing. For example, the statement <code>exclude.dir c:\test\dan\data1</code> excludes the <code>c:\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.</p> <p><b>Notes:</b></p> <ol style="list-style-type: none"> <li>1. If you define an exclude statement without using a drive letter, such as <code>exclude.dir dirname</code>, this will exclude from processing any directory named <code>dirname</code> on any drive.</li> <li>2. The following examples illustrate valid <i>exclude.dir</i> statements: <ul style="list-style-type: none"> <li>Exclude directory <code>C:\MyPrograms\Traverse</code> and its files and subdirectories:  <pre>exclude.dir c:\MyPrograms\Traverse</pre> </li> <li>Exclude all directories below <code>c:\MyPrograms\Traverse</code>. Note that directory <code>C:\MyPrograms\Traverse</code> and the files immediately below <code>C:\MyPrograms\Traverse</code> will still be eligible for backup.  <pre>exclude.dir c:\MyPrograms\Traverse\*</pre> </li> <li>Exclude all directories whose names begin with <code>temp</code>, and are located within directory <code>x:\documents and settings</code> and its subdirectories, where <code>x:</code> is any drive.  <pre>exclude.dir "x:\documents and settings\...\temp"</pre> </li> <li>Exclude all directories whose names begin with <code>temp</code>, regardless of the drive or directory in which they reside:  <pre>exclude.dir temp*</pre> </li> </ul> </li> </ol> <p>The following example is invalid because it ends with a directory delimiter:  <pre>exclude.dir c:\MyPrograms\Traverse\</pre> </p> <li>3. Use the following statements to exclude drive <code>x:</code> altogether from backup processing. Note that the drive root (<code>x:\</code>) will still be backed up, but all other files and directories on <code>x:</code> will be excluded.  <pre>exclude x:\* exclude.dir x:\*</pre> </li> <li>4. An alternative method for excluding an entire drive from domain incremental backup is to use a domain statement to exclude the drive. For example:  <pre>domain -x:</pre> <p>This alternative still permits selective and explicit incremental backup processing of files on <code>x:</code>. For example:  <pre>dsmc s x:\ -subdir=yes dsmc i x: dsmc i x:\MyPrograms\ -subdir=yes</pre> </p> </li>	257



## Include-exclude statements for networked file systems

Include-Exclude statements that involve networked file systems (remote drives) must be written in the UNC format. For example, let's say that Z: is a mapped drive to a remote file system on `vista.sanjose.ibm.com`.

The old format would be to exclude `\dir\dir2` on the remote file system, as in this example:

```
EXCLUDE.DIR "Z:\dir1\dir2"
```

Here is an example of the new format using UNC:

```
EXCLUDE.DIR "\\vista.sanjose.ibm.com\d$\dir1\dir2"
```

The include-exclude statements written in the old format will not be recognized by Tivoli Storage Manager.

## Excluding directories from journal-based backup

There are two methods of excluding files and directories from a journal-based backup. One method is to add exclude statements to the client options file to prevent the files or directories from being backed up during backup processing. The other method is to add exclude statements to the journal configuration file `tsmjbbd.ini`, to prevent journal entries from being added for the files or directories, which prevents them from being processed during a journal-based backup.

**Note:** There is no correlation between the two exclude statements. It is preferable to place exclude statements in `tsmjbbd.ini` to prevent them from entering the journal database and being processed during a journal-based backup.

## Controlling backup, archive, image, system object, and system state processing

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

Table 9. Options for controlling backup, archive, image, system object, and system state processing

Option	Description	Page
<b>Backup processing</b>		
<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.	257
<i>include</i> <i>include.backup</i> <i>include.file</i>	Use these options to include files or assign management classes for backup processing.	280
<i>include.fs</i>	Use this option to set options on a file space-by-file space basis.	280
<b>Archive processing</b>		
<i>exclude.archive</i>	Excludes a file or group of files from archive services.	257
<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.	280
<b>Image processing</b>		

Table 9. Options for controlling backup, archive, image, system object, and system state processing (continued)

Option	Description	Page
<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 all Windows clients.	257
<i>exclude.fs.nas</i>	Excludes file systems on the NAS filer from an image backup when used with the <b>backup nas</b> command. If you do not specify a NAS node name, the file system identified applies to all NAS filers. The <b>backup nas</b> command ignores all other exclude statements including <i>exclude.dir</i> statements. This option is for all Windows clients.	257
<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 <b>backup image</b> command. The <b>backup image</b> command ignores all other include options. This option is valid for all Windows platforms.	280
<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 client options file (dsm.opt). See "Toc" on page 399 for more information. This option is valid for all Windows clients <i>only</i> .	280
<b>System object processing</b>		
<i>include.systemobject</i>	Assigns management classes for backup of Windows system objects. By default, Tivoli Storage Manager binds all system objects to the default management class. You cannot use this option to bind individual systemobject components to a different management class. You cannot use this option to include or exclude a system object from processing. This option is valid for Windows XP <i>only</i> . <b>Note:</b> Other include-exclude statements do not affect system object processing. It is unnecessary to explicitly include the registry staging directory <code>include c:\adsm.sys\...\*</code> for back up, to ensure that the registry is backed up properly. If you accidentally exclude a directory that is critical to a system object backup, the system object backup is not affected.	280
<i>exclude.systemobject</i>	Excludes individual system objects from backup services. Excluded system object types that you backed up previously are not expired during subsequent backups. This option only excludes the system object types that you specify from subsequent backups. This option is valid for Windows XP clients <i>only</i> .	257
<b>System state processing</b>		

Table 9. Options for controlling backup, archive, image, system object, and system state processing (continued)

Option	Description	Page
<i>exclude.systemservice</i>	Excludes individual system components from system state backup. Input can be the keyword or component name to be excluded from back up. This option is valid for Windows Server 2003 and Windows Vista <i>only</i> .	257
<i>include.systemstate</i>	Assigns management classes for back up of the Windows Server 2003 and Windows Vista system state. The default is to bind the system state object to the default management class.	280

## Excluding system files

**Recommendation:** Use the sample include-exclude list in the dsm.smp file as a starting point for your include-exclude list. This is the minimum recommended include-exclude list. The dsm.smp file is located in the config folder in the installation directory. If you accepted the defaults, the path to this file is C:\Program Files\Tivoli\TSM\config\dsm.smp

**Note:** If subfile backup is being used, the directory containing cache files should be excluded with the *exclude.dir* option.

There are exclude statements generated from a list defined by the Windows operating system in the Windows Registry. Those implicitly generated statements can be seen in the lines of output of the **query inclexcl** command with the source "operating system".

## Excluding files with UNC names

You can exclude remotely accessed files by specifying their universal naming convention (UNC) names in your exclude statement. For example, assume that local drive letter g is mapped to the remote share point

```
\\remote\books
```

You would like to exclude from backups all files at the root of this share point that have an extension of .txt. You could use either of the following commands:

```
exclude g:*.txt
exclude \\remote\books\*.txt
```

You cannot specify UNC names for removable drives such as CD, ZIP, or diskette. For example, the following command is *not valid*:

```
exclude \\ocean\a$\winnt\system32\...\*
```

## 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 10 on page 38. 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 431.

**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 10. Wildcard and other special characters

Character	Function
?	<p>The match one character matches any single character <i>except</i> the directory separator; it does not match the end of the string. For example:</p> <ul style="list-style-type: none"> <li>• The <b>pattern</b> <code>ab?</code>, <b>matches</b> <code>abc</code>, but <b>does not match</b> <code>ab</code>, <code>abab</code>, or <code>abzzz</code>.</li> <li>• The <b>pattern</b> <code>ab?rs</code>, <b>matches</b> <code>abfrs</code>, but <b>does not match</b> <code>abrs</code>, or <code>abllrs</code>.</li> <li>• The <b>pattern</b> <code>ab?ef?rs</code>, <b>matches</b> <code>abdefjrs</code>, but <b>does not match</b> <code>abefrs</code>, <code>abdefrs</code>, or <code>abefjrs</code>.</li> <li>• The <b>pattern</b> <code>ab??rs</code>, <b>matches</b> <code>abcdrs</code>, <code>abzzrs</code>, but <b>does not match</b> <code>abrs</code>, <code>abjrs</code>, or <code>abkkrs</code>.</li> </ul>
*	<p>The match-all character. For example:</p> <ul style="list-style-type: none"> <li>• The <b>pattern</b> <code>ab*</code>, <b>matches</b> <code>ab</code>, <code>abb</code>, <code>abxxx</code>, but <b>does not match</b> <code>a</code>, <code>b</code>, <code>aa</code>, <code>bb</code>.</li> <li>• The <b>pattern</b> <code>ab*rs</code>, <b>matches</b> <code>abrs</code>, <code>abtrs</code>, <code>abrsrs</code>, but <b>does not match</b> <code>ars</code>, or <code>aabrs</code>, <code>abrss</code>.</li> <li>• The <b>pattern</b> <code>ab*ef*rs</code>, <b>matches</b> <code>abefrs</code>, <code>abefghrs</code>, but <b>does not match</b> <code>abefr</code>, <code>abers</code>.</li> <li>• The <b>pattern</b> <code>abcd.*</code>, <b>matches</b> <code>abcd.c</code>, <code>abcd.txt</code>, but <b>does not match</b> <code>abcd</code>, <code>abdc</code>, or <code>abcdtxt</code>.</li> </ul>
\...	The match- <i>n</i> character matches zero or more directories.
\	<p>The directory separator character limits the scope of the search for the matching <i>n</i> characters and directories. If a pattern does not begin with a directory separator (or one does not follow the drive specification), a match-all directories is appended to the pattern. For example, these patterns are equivalent:</p> <pre>c:* c:\...\*</pre>
[	<p>The open character-class character begins the enumeration of a character class. For example:</p> <pre>xxx[abc] matches xxxa, xxxb, or xxxc.</pre>
-	<p>The character-class range includes characters from the first character to the last character specified. For example:</p> <pre>xxx[a-z] matches xxxa, xxxb, xxxc, ... xxxz.</pre> <p>This format should not be used to specify remote drives in an <i>exclude</i> statement.</p>
\	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.
:	<p>The drive separator character separates a file specification. The character <i>before</i> the colon identifies a drive letter. The characters <i>after</i> the colon identify file specification or pattern. For example:</p> <pre>d:\direct\file.nam</pre>

**Note:** Because a drive specification can consist of only one letter, you should not use more than one wildcard or a combination of a wildcards with a letter to designate a drive specification. The following patterns are not allowed, and if specified in the client options file (`dsm.opt`), will stop the client program immediately after it starts:

```
?*:\test.txt
```

```

*?:\...\pagefile.sys
H*:\test.*
*H:\test.txt
myvolume*:\
myvolume?*:\

```

If you are using UNC names, Table 11 shows how to correctly specify shared drives.

Table 11. Specifying a drive specification using wildcards

Incorrect	Correct
\\remote\*\...\*.*	\\remote\*\$\...\*.*
\\remote\?:\...\*.*	\\remote\?\$\...\*.*
\\remote\*\...\pagefile.sys	\\remote\*\$\...\pagefile.sys

## Examples using wildcards with include and exclude patterns

Table 12 contains examples of ways you might use wildcard characters with *include* and *exclude* patterns.

Table 12. Using wildcard characters with include and exclude patterns

Task	Pattern
Exclude all files during backup with an extension of <i>bak</i> , except those found on the d: drive in the dev directory.	exclude *.*.bak include d:\dev\*.bak
Exclude all files and directories in any tmp directory that might exist, <i>except</i> for the file d:\tmp\save.fil. Include this file.	exclude ?:\...\tmp\...\* include d:\tmp\save.fil
Exclude any .obj file for backup in any directory on the c: e: f: and g: drives.	exclude [ce-g]:\...\*.obj The c: e: f: and g: drives are local or removable.
Exclude the .obj files found in the root directory in the d: drive <i>only</i> .	exclude d:\*.obj
Exclude any file that resides under the tmp directory found on any drive.	exclude ?:tmp\...\*
Exclude the c:\mydir\test1 directory and any files and subdirectories under it.	exclude.dir c:\mydir\test1
Exclude all directories under the \mydir directory with names beginning with test.	exclude.dir c:\mydir\test*
Exclude all directories directly under the \mydir directory with names beginning with test, on any drive.	exclude.dir ?:\mydir\test*
Exclude the raw logical volume from image backup.	exclude.image c:\*
Exclude all directories and files on the local drives, except the c: drive.	exclude [abd-z]:\...\* exclude.dir [abd-z]:\...\*

## Controlling compression, encryption, and adaptive subfile backup processing

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

Table 13. Options for controlling compression, encryption, and adaptive subfile backup processing

Option	Description	Page
<b>Compression processing</b>		
<i>exclude.compression</i>	Excludes files from compression processing if <i>compression=yes</i> is specified. This option applies to backups and archives.	257
<i>include.compression</i>	Includes files for compression processing if <i>compression=yes</i> is specified. This option applies to backups and archives.	280
<b>Encryption processing</b>		
<i>exclude.encrypt</i>	Excludes files from encryption processing.	257
<i>include.encrypt</i>	Includes files for encryption processing.  The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received.  <b>Attention:</b> 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.	280
<b>Adaptive subfile backup processing</b>		
<i>exclude.subfile</i>	Excludes files from adaptive subfile backup processing. This option does not apply to archive processing. This option is valid for all Windows clients.	257
<i>include.subfile</i>	Includes files for adaptive subfile backup processing. This option does not apply to archive processing. This option is valid for all Windows clients.	280

## Testing an include-exclude list with the Preview command

You can preview the list of objects to be backed up or archived according to the include-exclude list, prior to sending any data to the server. The Tivoli Storage Manager client Java GUI directory tree shows detailed information of included and excluded objects. The directory tree windows in the Tivoli Storage Manager Client Java GUI allow you to select files and directories to include or exclude. You should use this **Preview** command to make sure that you include and exclude the correct files. The following is a sample scenario for using the include-exclude preview function.

Assume that you want to back up the files on your /Users/home file space. You bring up the Tivoli Storage Manager Client Java GUI and open the Backup tree. You can see all of the directories and files that have been excluded by your options file and other sources. You scroll down the tree and notice that all of the \*.o files in your /Volumes/home/mary/myobjdir will be backed up. You don't want to back up these files, so you right click a .o file, and choose "View File Details" from the popup menu. The dialog shows that these files are included. click the "Advanced" button and create a rule to exclude all .o files from the /home file space. A rule is created at the bottom of your options file. The current directory is refreshed in the Backup tree, and the .o files have the red 'X', meaning they are excluded. When you look at other directories, they show the new excludes that you have added. Press "Backup" and back up the files on your /home file space.

See “Preview” on page 486 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 client options file (dsm.opt). 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 client options file include-exclude list contains one or more *incl excl* options that specify include-exclude files, the include-exclude statements in these files are placed in the list position occupied by the *incl excl* option and processed accordingly.

When performing an incremental backup, Tivoli Storage Manager evaluates all *exclude.dir* statements *first*, and removes the excluded directories and files from the list of objects available for processing. See “Excluding directories” on page 33 and “Exclude options” on page 257 for more information about the *exclude.dir* option.

After evaluating all *exclude.dir* statements, Tivoli Storage Manager evaluates the 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 165 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 505 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 ?:\*.obj
include c:\foo\...\*.obj
exclude c:\foo\junk\*.obj
```

The file being processed is: c:\foo\dev\test.obj. Processing follows these steps:

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

### Example 2

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

```
exclude ?:\*.obj
include c:\foo\...\*.obj
exclude c:\foo\junk\*.obj
```

The file being processed is: c:\widg\copyit.bat. 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 c:\widg\copyit.bat is implicitly included and backed up.

### Example 3

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

```
exclude ?:\...\*.obj
include c:\foo\...\*.obj
exclude c:\foo\junk\*.obj
```

The current file being processed is: c:\lib\objs\printf.obj. 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 c:\lib\objs\printf.obj is not backed up.

## Processing rules when using UNC names

When processing files with UNC names, Tivoli Storage Manager uses the rules described above. In addition, the following rules apply.

### Explicit use of UNC names for remote drives

Tivoli Storage Manager will recognize explicit use of UNC names for remote drives. For example, in the list below the UNC name pattern can be substituted for the DOS pattern. Assume local drive letter r: is mapped to remote share point \\remote\c\$, s: is mapped to \\remote\share4, and t: is mapped to \\remote\share2.



Table 14. Using UNC names with include and exclude patterns

UNC name pattern	DOS pattern
\\remote\c\$\include\file.out	r:\include\file.out
\\remote\c\$\...\file.out	r:\...\file.out
\\remote\share4\exclude\*	s:\exclude\*
\\remote\share2\...\?.out	t:\...\?.out

### Conversion of DOS pathnames for fixed and remote drives

Tivoli Storage Manager converts DOS pathnames that are mapped to remote share points. For example, a remote share point that is mapped from

r:\test\...\exclude.out

to

\\remote\share\test\...\exclude.out

is converted. Remote share points that are not mapped are not converted. Files on removable media are not converted.

### Character class matching

The following examples show valid matches using character class:

```
\\remote[a-z]\share\file.txt
matches    \\remotea\share\file.txt
           \\remote\share[a-z]\file.txt
matches    \\remote\sharex\file.txt
           \\remote\share\file[a-z].txt
matches    \\remote\share\fileg.txt
```



---

## Chapter 3. Getting started

This chapter includes instructions for the following tasks:

Task	Page
Starting a GUI session	48
Starting a command-line session	49
Starting a Web client session	51
Starting the client scheduler automatically	53
Changing your password	53
Sorting file lists using the Tivoli Storage Manager GUI	54
Displaying online help	55
Ending a session	55

---

### Considerations for client services

This section explains the types of Tivoli Storage Manager Backup-Archive client operations that can be performed and the security rights that are needed.

#### Using Tivoli Storage Manager Backup-Archive client as a member of the Backup Operators group

The Backup Operators group is a built-in Windows group which can override security restrictions for the sole purpose of backing up or restoring files. Since this group has a limited set of user rights, some functions are not available to a member of the Backup Operators group. The following summarizes the Tivoli Storage Manager Backup-Archive client operations which can and cannot be performed by a member of Backup Operators:

- Can back up and restore files (see Table 15 on page 46)
- Can back up system state. Use the Administrator user to restore the system state.
- Can start the TSM Scheduler service
- Cannot start any other services (TSM Client Acceptor, TSM Remote Client Agent, and TSM Journal Service)
- Cannot install and configure client services
- Cannot use open file support (OFS)
- Cannot perform image backup and restore operations
- Cannot back up and restore Windows file shares

You must have local or domain administrator privileges to install and configure Tivoli Storage Manager client services.

The following table summarizes the user security rights needed for backup and restore.

Table 15. Required user security rights for Tivoli Storage Manager backup and restore services

Operating system	Account	What can I backup and restore?
Windows 2003	Member of Administrators group or member of Backup Operators group	<ul style="list-style-type: none"> <li>• All file and directory objects</li> <li>• System state data</li> </ul>
Windows Vista	Member of Administrators group or member of Backup Operators group	<ul style="list-style-type: none"> <li>• All file and directory objects</li> <li>• System state data (Backup Operators group cannot restore system state data)</li> </ul>
Windows 2003 or Windows Vista	All other accounts	<ul style="list-style-type: none"> <li>• All file and directory objects to which the account has access</li> <li>• System Access Control Lists (SACL) used for auditing entries will not be backed up or restored</li> </ul>
Windows XP	Member of Administrators group	<ul style="list-style-type: none"> <li>• All file and directory objects</li> <li>• System Object data</li> </ul>
Windows XP	All other accounts	<ul style="list-style-type: none"> <li>• All file and directory objects to which the account has access</li> <li>• System Access Control Lists (SACL) used for auditing entries will not be backed up or restored</li> </ul>
<p>In order to back up all file objects, directory objects, system state data with Windows 2003 and Windows Vista the account needs to have the "Backup files and directories" and "Restore files and directories" security settings. By default, the Administrators and Backup Operators groups are granted these security settings.</p>		

By default, Tivoli Storage Manager client services run under the local system account. However, the local system account does not have access to network mapped drives and does not have the same permissions and logon properties as a user that is logged into the system. If you experience discrepancies between a user initiated backup and a scheduled backup using the local system account, consider changing the services to run under the user account.

**Tip:** In addition to the appropriate user security rights, the Tivoli Storage Manager Backup-Archive client requires that the user has read permission to the root of any drive that needs to be backed up or restored. If you are using the system account to logon for the TSM Scheduler service, ensure that you grant the system account (SYSTEM) read access to the root of the drive. It is not sufficient to grant Everyone read access to the root of the drive.

Domain resources, such as network drives, can only be accessed by services configured to run under a domain authorized account using *dsmcutil* or the Service Control Panel Application.

## Considerations before you start using a Backup Operators group account for backup, archive, restore, and retrieve operations

- If you have already been using the Tivoli Storage Manager Backup-Archive client with an Administrators group account you might not be able to launch the client because you cannot open the log files (for example `dsmerror.log`). To alleviate this problem, you can grant the Backup Operators group Read and Write permissions to the log files or the directories containing these log files.
- If you have been performing Windows 2003 or Windows Vista system state backups with an account which is a member of the Administrators group and if you want to start doing backups or restores with an account which is a member of the Backup Operators group, you must delete the staging directory, `c:\adsm.sys`, before attempting backups or restores of the system state as a member of the Backup Operators group.
- If you have existing backups from a Tivoli Storage Manager V5.2 or previous client and you attempt an incremental backup of an existing file space with a member of the Backup Operators group, all of the data will appear as changed and it will be resent to the Tivoli Storage Manager Server.
- Members of the Backup Operators group might not be able to back up or restore file data that was encrypted by an Administrator account using the Windows encrypting file system (EFS).
- Members of the Backup Operators group do not have the proper authority to update the last access time for files that is encrypted with the Windows encrypting file system (EFS). If EFS files are restored by a member of the Backup Operators group, the last access time will not be preserved.

---

### Permissions required to restore files that use adaptive subfile backup

To restore files that were processed using adaptive subfile backup, you must either be the owner of the file or have read access. These permissions are in addition to those required to perform a normal restore. For more information about adaptive subfile backups, see “Performing a backup with limited bandwidth” on page 61.

---

### Permissions required to back up, archive, restore or retrieve files on cluster resources

To back up, restore, archive, or retrieve data residing on Microsoft Cluster Server (MSCS) or Veritas Cluster Server cluster resources, your Windows account must belong to the Administrators or Domain Administrators group or Backup Operators group. By default, Backup Operators do not have the user rights necessary to perform these tasks on a cluster node. However, Backup Operators can perform this procedure if that group is added to the security descriptor for the Cluster service. You can do that using Cluster Administrator or `cluster.exe`.

---

### Tivoli Storage Manager client authentication

When using the backup-archive GUI, command-line client, or the Web client, you can log on using a node name and password *or* administrative user ID and password. Tivoli Storage Manager prompts for your user ID and compares it to the configured node name. If they match, Tivoli Storage Manager attempts to authenticate the user ID as a node name. If the authentication fails or if the user ID does not match the configured node name, the client attempts to authenticate the user ID as an administrative user ID.

To use an administrative user ID with any of the backup-archive clients, the user ID must have one of the following authorities:

*System privilege*

Authority over the entire system. An administrator with system privilege can perform any administrative task.

*Policy privilege*

Authority over the node policy domain. Allows an administrator to manage policy objects, register client nodes, and schedule client operations for client nodes.

*Client owner*

Authority over the registered Tivoli Storage Manager client node. You can access the client through the Web client or backup-archive client. You own the data and have a right to physically gain access to the data remotely. You can back up and restore files on the same or different system, and you can delete file spaces or archive data.

*Client access*

To use the Web Client to back up and restore files on a remote client system, you must have an administrative user ID with client access authority over the node name for the remote client system. If you do not want Tivoli Storage Manager administrators with client access authority over your node name to be able to back up and restore files on your system, specify the *revokeremoteaccess* option in your client options file (dsm.opt). See “Revokeremoteaccess” on page 348 for more information.

Client access authority only allows Tivoli Storage Manager administrators to back up and restore files on remote systems. They do not have physical access to the data. That is, they cannot restore the data belonging to the remote system to their own systems. To restore data belonging to a remote system to your own system, you must possess at least client owner authority.

To determine what authority you have, you can use either of the following methods:

- From the main Tivoli Storage Manager GUI window, select **File** → **Connection Information**.
- Use the Tivoli Storage Manager server QUERY ADMIN command from the administrative command-line client. See the appropriate *Tivoli Storage Manager Administrator's Reference* for more information about the QUERY ADMIN command.

---

## Starting a GUI session

Start a GUI session using one of the following methods:

- Click the Windows **Start** button and select **Programs** → **Tivoli Storage Manager** → **Backup Archive GUI**.
- Click the Windows **Start** button and select **Run**, then enter the full path to the backup client executable file (**dsm.exe**).
- On the command line, change directory to the Tivoli Storage Manager installation directory and enter **dsm**.
- On Windows Vista, when you log on as an administrator, you are running as a standard user. If you open a command prompt without running as an administrator, you will be asked for consent to run **dsm**.

Tivoli Storage Manager locates the *client options* file (dsm.opt) and starts with the options specified in that file. See Chapter 2, “Configuring Tivoli Storage Manager,” on page 15 for more information about the client options file.

## 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 53.

## Setup wizard

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

---

## Starting a command-line session

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

**Note:** If all DSM environment variables are set, you can enter the **dsmc** command from any directory; otherwise, enter the fully qualified path.

- Open the Windows **Start** menu and select **Programs** → **Tivoli Storage Manager** → **Backup Client Command Line**.
- On the command line enter **dsmc** followed by the command, to execute a single command (*batch mode*).
- 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 c:\myfiles\file1.dat, enter the command:

```
dsmc archive c:\myfiles\file1.dat
```

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

```
dsmc incremental
```

Depending upon the current setting of your *passwordaccess* option, Tivoli Storage Manager might prompt you for your password before the command is processed in a batch mode session. See “Passwordaccess” on page 320 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>
```

On Windows Vista, when you log on as an administrator, you are running as a standard user. If you open a command prompt without running as an administrator, you will be asked for consent and another command prompt will be opened.

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 `c:\myfiles\file1.dat`, enter the command:  
`archive c:\myfiles\file1.dat`

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 429 for a discussion of how options are handled in interactive mode.

For more information about using the command-line client, see “Starting and ending a client command session” on page 427.

For more information about client services, see “Considerations for client services” on page 45.

See Chapter 10, “Using commands,” on page 423 for more information on how to start and use the command-line client.

## Displaying Euro characters in command-line prompt

To display the Euro character in the Windows command-line prompt (console window):

1. Contact your Microsoft Representative for the 858 code page (the file name is `c_858.nls`). Copy the file into your Windows system32 directory (for example, `C:\WINNT\system32`).
2. Edit the Windows Registry key:

**Note:** The Windows Registry editor is very unforgiving, as changes cannot be undone. Errors made in editing the Windows Registry can cause your system to malfunction, and can even render it unbootable. We recommend that you



exercise the utmost care in editing the Windows Registry. If you are unfamiliar with using the Windows Registry editor, then ask someone else who is familiar with the Windows Registry editor to help you.

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Nls\CodePage\850
```

and set it to value `c_858.nls`.

3. In your Regional Settings, select a Western European country (Germany, France, Italy, etc.) as your locale setting.
4. Exit and reboot the system.

Ensure that the console window font you use supports the Euro symbol (such as Lucida Console).

## Using options on the DSMC command

You can include options on the `dsmc` command. For example, suppose you have one workstation with node name `galaxy1`, and another workstation with node name `galaxy2`, and you want to restore `galaxy1`'s data to the `galaxy2` system. To recover a file from one workstation (`galaxy1`) while at the other workstation (`galaxy2`), you will need access to `galaxy1`. We recommend that you use the `set access` command to gain access.

For example, assume the file to be recovered on `galaxy1` is `c:\universe\saturn.planet`. The owner of `galaxy1` enters the following command:

```
dsmc set access archive c:\universe\saturn.planet galaxy2
```

When access is granted, you would retrieve the file by entering:

```
dsmc retrieve -fromnode=galaxy1 \\galaxy1\universe\saturn.planet c:\
```

**Note:** Access to another user's files can also be granted and gained using the GUI.

For more information about the `set access` command, see "Set Access" on page 563. For more information about the `fromnode` option, see "Fromnode" on page 268.

If you have more than one backup server in your organization, you can easily switch between them using a command-line option. To override the server specified in `dsm.opt`, you could use a command such as this:

```
dsmc -tcpserveraddress=myserver -node=mynode -tcpport=1599
```

---

## Starting a Web client session

After installing the Web client on your workstation (see "Configuring the Web client" on page 20) you can use the Web client to perform backup, archive, restore, and retrieve operations from any browser that is running Java Runtime Environment (JRE) 1.5 or 1.4.x, where *x* is greater than or equal to 1. The Web client facilitates the use of assistive devices for users with disabilities and contains improved keyboard navigation. The native look and feel of the platform running the browser is preserved. The Web client will use most desktop font and color settings when run in browsers on Windows platforms.

The Web client runs in the following browsers:

- Microsoft Internet Explorer 5.0 or higher with JRE 1.5 or 1.4.x, where *x* is greater than or equal to 1

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

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 1.5 or 1.4.x, where *x* is greater than or equal to 1 (unless otherwise noted), manually from the appropriate URL:

- For Windows:  
<http://java.sun.com/javase/index.jsp>

**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 system:

- **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.

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 HTTP port number defined on the client workstation; the default is 1581. For example:

```
http://myhost.mycompany.com:1581
```

**Note:** Entering a different URL or pressing the browser **Back** button during an operation disconnects the Web client and causes the current operation to end. 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 47 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 348 for more information about the *revokeremoteaccess* option. See “Tivoli Storage Manager client authentication” on page 47 for more information about access authorities.

---

## Starting the client scheduler automatically

You can start the client scheduler automatically when you start your workstation. If the Tivoli Storage Manager administrator has defined schedules for your node, starting the client scheduler permits you to automatically back up your workstation (or perform other scheduled actions). See Chapter 7, “Automating tasks,” on page 145 for more information about the client scheduler.

You can also use the Tivoli Storage Manager Client Acceptor service to manage the scheduler.

Tivoli Storage Manager supports remote network connections to the server. With a remote network connection, mobile users no longer need to dial-in to their company network when a backup is scheduled to run. Tivoli Storage Manager automatically establishes a connection before the scheduled backup occurs. If the connection fails, Tivoli Storage Manager reestablishes the connection before attempting the backup. See “Mobile dial-up support” on page 151 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 319 for additional password information.

- “Starting the client scheduler automatically” on page 53
- “Password” on page 319
- “Set Password” on page 568

## Sorting file lists using the Tivoli Storage Manager GUI

Table 16. 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"> <li>• Select one of the <b>Sort by</b> items from the <b>View</b> menu bar.</li> <li>• Click the appropriate column heading in the File List box.</li> </ul>
Display active and inactive backup versions	<ul style="list-style-type: none"> <li>• Click the <b>Display Active/Inactive Files</b> option from the <b>View</b> menu.</li> <li>• Click the <b>Display both active and inactive files</b> tool on the tool bar.</li> </ul>
Display only active backup versions	Click the <b>Display active files only</b> option from the <b>View</b> menu.
Selecting files to restore or retrieve.	<ul style="list-style-type: none"> <li>• Click the selection box next to the directory or file name that you want to restore or retrieve.</li> <li>• Highlight the files that you want to restore or retrieve and click the <b>Select Items</b> tool on the tool bar.</li> <li>• Highlight the files that you want to restore or retrieve and click the <b>Select Items</b> option from the <b>Edit</b> menu.</li> </ul>
Deselecting files	<ul style="list-style-type: none"> <li>• Click the checked selection box next to the directory or file name.</li> <li>• Highlight the files that you want to deselect and click the <b>Deselect Items</b> tool on the tool bar.</li> <li>• Highlight the files that you want to deselect and click the <b>Deselect Items</b> option from the <b>Edit</b> menu.</li> </ul>
Displaying file information	<ul style="list-style-type: none"> <li>• Highlight the file name, and click the <b>View File Details</b> button on the tool bar.</li> <li>• Highlight the file name, and select <b>File Details</b> from the <b>View</b> menu.</li> </ul>

### 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 125).
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 277 and “Pick” on page 322 for more information.

---

## Displaying online help

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

- **From the Tivoli Storage Manager 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 474 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 Web client:**
  - Open the **File** menu and select **Exit**.
  - Click the **X** icon in the upper right corner.
  - Open the **System** menu and select **Close**.
  - Press **Alt+F4**.
  - **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>

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/>

---

## Other sources of online help

An anonymous FTP server ([ftp.software.ibm.com/storage](ftp://ftp.software.ibm.com/storage)) is available where you can find maintenance and other Tivoli Storage Manager-related materials. Three other anonymous servers are unofficially maintained by non-IBM volunteers. These servers are:

[ftp.rz.uni-karlsruhe.de](ftp://ftp.rz.uni-karlsruhe.de) (mirror - Germany)

[ftp.wu-wien.ac.at](ftp://ftp.wu-wien.ac.at) (mirror - Austria)

[ftp.cac.psu.edu](ftp://ftp.cac.psu.edu) (mirror - Pennsylvania)

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

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

Also see "Contacting customer support" on page xv for product support information.

---

## Chapter 4. Backing up your data

Use the backup-archive client to store backup versions of your files on the Tivoli Storage Manager server. You can restore these backup versions if the original files are lost or damaged. This chapter discusses various ways you can back up your data.

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

- Estimate
- Preferences editor
- Setup wizard

See “Starting a Web client session” on page 51 for information on starting the Web client.

Unless otherwise specified, references to Windows refer to all supported Windows operating systems.

See Table 17 for a list of primary tasks.

*Table 17. Backup: Primary tasks*

Task	Page
Planning your backups	57
Pre-backup considerations	59
Performing an incremental, selective, or incremental-by-date backup	64
Deleting backup data	75
Group backup: Backing up files from one or more file spaces	76
Backing up system objects (Windows XP)	79
Backing up Windows system state	80
Backing up Automated System Recovery (ASR) files (Windows XP, Windows Server 2003)	82
Performing an image backup	82
Backing up NAS file systems	87
Using VMware Consolidated Backup	92
Backing up NetApp CIFS share definitions	95
Backing up the WebSphere Application Server	95

---

### Planning your backups

If you are a first-time user, or if you only back up files occasionally, you can use the checklist in Table 18 for preliminary planning.

*Table 18. Planning your backups*

—	Decide whether you want to back up or archive files. See “When to back up and when to archive files” on page 58 for more information.
---	---

Table 18. Planning your backups (continued)

—	See “Pre-backup considerations” on page 59 for important migration information, and how you might increase performance before backing up files and directories.
—	Create an include-exclude list to specify files and directories you want to exclude from backup services. See “Using an include-exclude list to control processing” on page 62 for more information.
—	Decide what type of backup you want according to your needs. See the following sections for more information: <ul style="list-style-type: none"><li>• “Performing an incremental, selective, or incremental-by-date backup” on page 64</li><li>• “Group backup: Backing up files from one or more file spaces” on page 76</li><li>• “Backing up system objects (Windows XP)” on page 79</li><li>• “Backing up Windows system state” on page 80</li><li>• “Backing up Automated System Recovery (ASR) files (Windows XP, Windows Server 2003)” on page 82</li><li>• “Performing an image backup” on page 82</li><li>• “Backing up NAS file systems” on page 87</li><li>• “Backing up the WebSphere Application Server” on page 95</li></ul>
—	For additional backup considerations, see “Backup: Additional considerations” on page 97.

## When to back up and when to archive files

When the backup-archive client backs up or archives a file, it sends a copy of the file and its associated attributes to the server; however, backups and archives have different goals.

Use backups to protect against unforeseen damage to your files, and use archives for maintaining more permanent versions of your files.

Backup data is managed by version, using predetermined policy based rules. Using these rules, the Tivoli Storage Manager administrator can control the following:

- The number of versions
- The number of days each additional backup copy is kept
- What happens to backup data versions when the file is deleted on the client system

Each copy of the file stored on the server is considered to be a separate and unique version of the file.

Archive is a powerful and extremely flexible mechanism for storing long term data. Archive data is kept for a specified number of days. Archive has no concept or support for versions. The user or Tivoli Storage Manager administrator is responsible for determining what files get added to an archive.

**Recommendation:** If archive is run multiple times against a file using the same archive description, a new copy of the file will be added to the archive each time that archive is run. To simplify retrieve, store only one copy of a file in each archive.

Backups protect against file damage or loss that could occur through accidental deletion, corruption, disk crashes, and so forth. The server maintains one or more



backup versions for each file that you back up. Older versions are deleted as newer versions are made. The number of backup versions the server maintains is set by your administrator.

**Note:** If you frequently create archives for the same data, consider using instant archives (backup sets) instead. Frequent archive operations can create a large amount of metadata in the server database increasing database growth and decreasing performance for operations such as expiration. See “Restoring data from a backupset” on page 122 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 137 for more information about archiving and retrieving files.

---

## Pre-backup considerations

The Windows client is Unicode-enabled. However, client versions before Version 4.2 were not enabled for Unicode. If you are backing up a system that had at one time used a client version older than Version 4.2, and the file spaces have not yet been migrated to Unicode, then you need to plan for the migration of file spaces to Unicode. This involves renaming your file spaces on the server and creating new Unicode-enabled file spaces on the server using the *autofsrename* option.

For additional information about working with Unicode-enabled file spaces, see the following references:

- “Autofsrename” on page 202
- “Considerations for Unicode-enabled clients” on page 278
- “Detail” on page 232
- “Query Filespace” on page 499
- “Restore” on page 522
- “Retrieve” on page 553

## LAN-free data movement

All Windows 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 using 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*

## 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 249 for more information.

### *lanfreecommmethod*

Specifies a communication protocol between the client and the Storage Agent. See "Lanfreecommmethod" on page 291 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 293 for more information.

### *lanfreetcpport*

Specifies the TCP/IP port number where the Storage Agent is listening. See "Lanfreetcpport" on page 294 for more information.

### *lanfreetcpserveraddress*

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 (dsm.opt) 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. See "Memoryefficientbackup" on page 305 for more information about this option.

## 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*. See “Memoryefficientbackup” on page 305 for more information about this option.
- 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, if you want to back up two file systems at the same time but their memory requirements exceed the limits of a single process, then start one instance of the client to back up one of the file systems, and start a second instance of the client to back up the other file system. See “Resourceutilization” on page 344 for more information about this option.
- Use the - *incrbydate* client option to perform an “incremental-by-date” backup. See “Incrbydate” on page 286 for more information about this option.
- 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. See “Exclude options” on page 257 for more information about this option.
- 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. See “Memoryefficientbackup” on page 305 for more information about this option.

## Performing a backup with limited bandwidth

If you plan to perform a backup over a network device that has limited bandwidth, such as a modem, you can help reduce network traffic by using *adaptive subfile backup*. Adaptive subfile backup requires a Version 4 or later server. Reducing traffic can also increase the speed of your backup. An adaptive subfile backup sends only changed portions of a file to the server during successive backup operations.

Perform the following steps to use adaptive subfile backup.

1. Ensure that the server allows this type of backup. Check with your system administrator.
2. Set the *subfilebackup*, *subfilecachepath* and *subfilecachesize* options in your client options file *dsm.opt*. For more information on adaptive subfile backup options, see:
  - “Controlling compression, encryption, and adaptive subfile backup processing” on page 39
  - “Subfilebackup” on page 383

**Note:** If you set the *subfilebackup* option to *yes* and no *exclude.subfile* statements exist, Tivoli Storage Manager considers all files for adaptive subfile backup processing.

- “Adaptive subfile backup processing” on page 260
  - “Subfilecachepath” on page 384
  - “Subfilecachesize” on page 385
3. Add an *exclude.dir* option to your client options file, which excludes the directory containing the subfile cache information. See “Exclude options” on page 257 for more information.

When you are ready to back up changes, adaptive subfile backup will occur as part of incremental and selective backups.

You can restore an adaptive subfile backup using the **restore** command. The server will restore the base file along with the updates from the changed (delta) subfile so that what you receive is the latest backed-up version of your file.

## Using an include-exclude list to control processing

There might be files on your system that you do not want to back up. These files might be operating system or application files that you can easily recover by reinstalling the program, or any other file that you can easily rebuild.

Use the *include* and *exclude* options in the client options file (*dsm.opt*) 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.

The include-exclude list might contain items specified by the server. To view the contents of your include-exclude list, use the **query inclexcl** command.

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 145 for more information on management classes and how to assign them to files.

For more information about setting up an include-exclude list, see “Creating an include-exclude list (optional)” on page 31.

You can also add include-exclude statements in the Tivoli Storage Manager Client GUI directory tree. The **preview** command allows you to see the resultant effects of

the currently defined include-exclude list without need of performing an actual backup of data. For more information about the **preview** command, see “Preview” on page 486.

## Encrypting data during backup or archive operations

### Caution

If the encryption key is not saved in the Windows Registry, and you have forgotten the key, your data will be unrecoverable.

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 250 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 280. For more information about the *exclude.encrypt* option, see “Exclude options” on page 257.

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. You can specify whether to save the encryption key password in the Windows Registry by using the *encryptkey* option.

Tivoli Storage Manager client encryption allows you to enter a value of up to 63 characters in length. This encryption password needs to be confirmed when encrypting the file for backup, and also needs to be entered when performing restores of encrypted files.

While restoring an encrypted file, Tivoli Storage Manager will prompt you for the key password to decrypt the file in the following cases:

- If the *encryptkey* option is set to Prompt.
- If the key supplied by the user in the above case does not match.
- If the *encryptkey* option is set to Save and the locally saved key password does not match the encrypted file.

For more information about the *encryptkey* option, see “Encryptkey” on page 251.

See “Backup: Additional considerations” on page 97 for additional information to consider when performing a backup. You do not need to understand this information to perform basic backups.

## Maximum file size for operations

Table 19 on page 64 specifies the maximum file sizes for backup, restore, archive, and retrieve operations.

Table 19. Maximum file size

File system	Max file size (in bytes)
FAT16	2,147,483,647 (2GB)
FAT32	4,294,967,295 (4GB)
NTFS	17,592,185,978,880 (16TB-64K)

## 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 set up schedules to automatically perform backups. See Chapter 7, “Automating tasks,” on page 145 for information about checking and running schedules that are available to you. The following sections discuss how to back up files without a schedule.

There are three types of incremental backup: *full*, *partial*, and *incremental-by-date*.

Incremental backup might recall migrated files, while selective backup and archive will always recall migrated files, if you do not use the *skipmigrated* option. Using incremental-by-date will not back up any migrated file. Refer to *IBM Tivoli Storage Manager for HSM for Windows Administration Guide* for details about using the *skipmigrated* option and incremental-by-date processing of migrated files.

## Supported file systems

Tivoli Storage Manager provides backup and archive services for all files residing in the following types of file systems:

- File Allocation Table (FAT)
- NT File System (NTFS)
- FAT 32

## Full and partial incremental backup

If you select entire drives, 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 drives 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 153 for more information about storage management policies.

To perform a full or partial incremental backup using the client 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
- Update the attributes on the server. Changes to read-only, hidden and compressed file attributes causes Tivoli Storage Manager to send only attribute changes to the server. Changes to sparse file, reparse point, and encrypted file attributes cause Tivoli Storage Manager to send the entire file to the server. The archive attribute is not examined by Tivoli Storage Manager in determining changed files.
- Back up files whose contents changed since the last backup. Changes include *any* of the following:
  - File size
  - Date or time of last modification
  - Extended Attributes
  - Access Control List
  - Sparse, reparse point or encrypted file attributes.
  - NTFS file security descriptors. These are the Owner Security Identifier (SID), Group SID, Discretionary Access Control List (ACL), and System ACL.

- 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 331 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 257.

### 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 153 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 153 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 153 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 153 for more information.

### Journal-based backup

Journal-based backup is supported for all Windows clients, except for clients running on Windows Server 2003 for Itanium-based Systems. A backup for a particular file system will be journal-based when the Tivoli Storage Manager Journal Service 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 engine service” on page 28 for more information on configuring file systems for journal-based backup.



To support journal-based backup, you must configure the journal engine service using the **dsmcutil** command or the client GUI setup wizard. See Appendix B, “Using the Client Service Configuration Utility,” on page 585 for more information about using the **dsmcutil** command to install the journal engine service. To install and configure this service using the client GUI setup wizard, see “Configuring the journal engine service” on page 28.

If you install the *journal engine service* 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. In order to successfully perform a journal-based backup, several conditions must be met. These include:

1. The journal service 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 engine service* 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; therefore the client 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, so the frequency attribute is ignored.

The journal engine service does not record changes in specific system files, such as the registry in the journal. Therefore, a journal-based backup will not back up this file. See the journal service configuration file, `tsmjbbd.ini`, in the client installation directory for excluded system files. See Appendix C, “Journal service configuration,” on page 601 for more information about the journal service configuration file, `tsmjbbd.ini`.

For more information about journal-based backup, refer to “Incremental” on page 475 and Appendix C, “Journal service configuration,” on page 601.

**Notes:**

1. Journal-based backup is supported on a Version 4.2 or higher Tivoli Storage Manager server.
2. Beginning with Tivoli Storage Manager Version 5.3, multiple journal-based backup sessions are possible.
3. When using anti-virus software, there are limitations to journal-based backup.

4. 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.*

To perform an incremental-by-date backup using the GUI, select the **Incremental (date only)** option from the *type of backup* pull-down menu or use the *incrbydate* option with the **incremental** command.

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.

**Tip:** Unlike incremental backups, incremental-by-date backups do not expire deleted files or rebind backup versions to a new management class if you change the management class.

## Comparing full incremental, partial incremental, 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 GUI, see “Backing up data using the GUI” on page 71 for more information. Use the **selective** command to perform a selective backup from the client command line. See “Selective” on page 559 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.
- Does not expire deleted files.
- Does not rebind backup versions to a new management class if you change the management class.

## Open file support for backup operations

There are two snapshot providers that can be used for open file support: LVSA and VSS (VSS is not supported on Windows XP). VSS is the recommended solution, since it utilizes Microsoft’s strategic snapshot solution. The LVSA is provided for customers migrating from prior versions of Tivoli Storage Manager, where LVSA was utilized, and as an alternative to VSS if there are issues with using VSS. Refer to “Upgrading Open File Support (OFS) or online image” on page 4 for additional information.

**Note:** To configure open file support, see “Configuring Open File Support (OFS)” on page 29.

Some applications can create files and open these files in a way that denies access to all other processes on a Microsoft Windows operating system. Although this is not a common practice, it is sometimes used by database vendors or other applications that might want to limit access to certain files. By restricting access to these files, backup products are prevented from backing up this data. These *locked* files are *not* the same as files which are *open* or *in use*. Tivoli Storage Manager, running without the OFS feature, can back up *open* or *in use* files, including files

that are open for reading or writing, files that are changing during the backup, executable and dll files that are running, log files that are being appended to, etc.

You can perform OFS or online image backups on machines with a single NTFS-based C: drive.

The following is the error message that is seen in the `dsmerror.log` when a Tivoli Storage Manager backup encounters one of these locked files without OFS support enabled:

```
ANS4987E Error processing '\\machine1\d$\dir1\lockedfile.xyz': the object is in use by another process
```

```
ANS1228E Sending of object '\\machine1\d$\dir1\lockedfile.xyz' failed
```

OFS should not be used for backing up locked Windows system files, such as system objects (Windows XP) and system state (Windows Server 2003 and Windows Vista). The Tivoli Storage Manager client has advanced features for backing up data contained within these files. The backup of the system data that is contained in these files requires additional processing and must be backed up in a group to allow for a successful restore. These files are excluded from the Tivoli Storage Manager file level backup. See the following sections for more information:

- “Backing up system objects (Windows XP)” on page 79
- “Backing up Windows system state” on page 80

For database applications that use certain files for transactional consistency (for example, a recovery log file), it might not be possible to back up and restore these files without database coordination. In these situations, do not back up these database files with the normal Tivoli Storage Manager file level backup. You can exclude these files from backup processing using an *exclude* or *exclude.dir* statement. Tivoli Storage Manager provides a number of Data Protection clients (IBM Tivoli Storage Manager for Databases, IBM Tivoli Storage Manager for Mail, IBM Tivoli Storage Manager for Application Servers, etc.) which provide this database coordination and backup along with other advanced database backup features. For a current list of Data Protection clients go to this Web site:

<http://www.ibm.com/software/tivoli/products/storage-mgr/product-links.html>

For private applications or other database products where a Data Protection client is not available, you can use the *preschedulecmd* option to signal the database or application to do one of the following:

- Take the steps necessary to move these files to a consistent and unopen state.
- Bring down the database before the file level backup is started.
- Program or script another method to back up this data and exclude these files from the file level backup. In these cases the OFS feature is not necessary since these files are no longer unavailable or locked by the application. Once the file level backup has completed, use the *postschedulecmd* option to bring the database back online or restart the application.

If the time it takes to complete the file level backup is too long to have the open files offline (for example, having the database offline or holding up transactions), use the OFS feature to create a point-in-time snapshot of the volume. In this case, use the *presnapshotcmd* and *postsnapshotcmd* options to signal the database or application to coordinate with the backup of these open files. The snapshot, which occurs between the pre-snapshot command and post-snapshot command, should only take a few seconds to create. This allows the database or application to resume operations quickly while still allowing Tivoli Storage Manager to perform a

full incremental backup of the volume, including the locked files. There are other situations where these application *locked files* can be safely backed up and restored on a file by file basis. In these situations, you can enable the OFS feature for that volume where the open files exist. Tivoli Storage Manager file level backup will then have access to these files and back them up using the Tivoli Storage Manager file level backup and archive operations.

For information about Tivoli Storage Manager Open File Support restrictions and issues, search for the **TSM Client Open File Support** document, which is Technote number 1248971, at the following Web site:

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

If open file support has been configured (see “Configuring Open File Support (OFS)” on page 29), Tivoli Storage Manager performs a snapshot backup or archive of files that are locked (or “in use”) by other applications. The snapshot allows the backup to be taken from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to the file system are not included in the backup. You can set the *snapshotproviderfs* parameter of the *include.fs* option to *none* to specify which drives do not use open file support.

To control an open file support operation you can specify these additional options in your dsm.opt file or as values of the *include.fs* option: *snapshotproviderfs*, *presnapshotcmd* and *postsnapshotcmd*. Additionally, when the LVSA is the snapshot provider, the following additional options can be specified: *snapshotcachelocation*, *snapshotcachesize*, *snapshotfsidleretries*, and *snapshotfsidlewait*. See Chapter 9, “Using processing options,” on page 165 for more information about these options.

#### Notes:

1. You can use the *include.fs* option to set snapshot options on a per file system basis.
2. For the LVSA, use the *snapshotcachelocation* option to relocate the cache if necessary. You can specify a snapshotcachelocation for a specific drive using the *include.fs* option.
3. Open file support is provided for both backup and archive. For backup, this includes incremental, incremental by date, selective, incremental image, adaptive subfile backup, and journal-based backup.
4. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with FAT, FAT32 or NTFS file systems. This support includes SAN-attached volumes that meet these requirements.
5. For the LVSA, if the client is unable to create a snapshot; for example, if the *snapshotfsidle* is not met or the *snapshotlocation* option is not valid, failover to a non-OFS backup will occur on the volumes where this problem exists.
6. To enable OFS support in a cluster environment, all machines in the cluster must have OFS configured. If LVSA is being used as the snapshot provider, ensure it is installed and configured on each machine in the cluster. You can also use VSS by setting the *snapshotproviderfs* option.

## Backing up data using the GUI

You can use Tivoli Storage Manager to back up specific files, a group of files with similar names, or entire directories. You can locate the files you want to back up by

searching or filtering. Filtering displays only the files that match the filter criteria for your backup. Files that do not match the filter criteria do not display.

To perform a GUI backup, use the following steps:

1. Click **Backup** from the GUI main window. The Backup window appears.
2. Expand the directory tree by clicking the plus sign **+**. To display files in a folder, click the folder icon. To search or filter files, click the **Search** icon from the tool bar.
3. Click the selection box for the object(s) you want to back up.
4. Select the type of backup from the pull down menu:
  - a. To run an incremental backup, select **Incremental (complete)**.
  - b. To run an incremental backup by date, select **Incremental (date only)**.
  - c. To run a selective backup, select **Always backup**.
  - d. To run an incremental backup without using the journal database, select **Incremental (without journal)**.

If you installed the journal engine service 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. This option lets you perform a traditional full incremental backup, instead of the default journal-based backup.

5. Click **Backup**. The Backup Task List window displays the backup processing status. When processing completes, the Backup Report window displays processing details.

Considerations:

- 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 or Web client GUI to view the backup policies defined by the Tivoli Storage Manager server for your client node. Select **Edit** → **Preferences** from the backup-archive client or Web client GUI and select the Include-Exclude tab in the Preferences editor to display your include-exclude list. See Chapter 7, “Automating tasks,” on page 145 and Chapter 8, “Understanding storage management policies,” on page 153 for more information about management classes.
- To modify specific backup options, click the **Options** button. Any options you change are effective during the current session *only*.
- To estimate the transfer time for your backup selections click the **Estimate** button. The estimated transfer is a rough calculation of the time it takes to transfer your data. This estimate is based on previous transfers of data between your workstation and the current Tivoli Storage Manager server, so you must run at least one backup operation first. The actual transfer time could be longer or shorter than the estimate due to factors like network traffic, system load on your workstation, or system load on the server. The Estimated Transfer Time field reads N/A if no files are sent to or from the current Tivoli Storage Manager server. The estimate function also does not take into account whether or not files are excluded from backup. The assumption made by the estimation algorithm is that all the files selected will be sent to the server.

**Note:** Tivoli Storage Manager creates the `dsm.ini` file to record statistics from prior backup activity between the Tivoli Storage Manager client and a given Tivoli Storage Manager server. These statistics are used in subsequent estimation

calculations. If a prior backup has never been performed between this client node and a given server, then there is no historic data upon which an estimation can be made and the estimation function will simply indicate "N/A".

- To perform subsequent incremental backups, from the Tivoli Storage Manager main window, open the **Actions** menu and select **Backup Domain**.

### Specifying drives in your domain

When you start Tivoli Storage Manager, it sets your default domain to the drives you specify with the *domain* option in the dsm.opt file (see "Domain" on page 239). If you do not set the *domain* option, the default domain is all local fixed drives (the drives on your workstation).

You can exclude any domain (including the systemobject domain) in your default domain from backup processing using the **Backup** tab in the Preferences editor. You can also exclude drives or the systemobject domain by specifying the dash (-) operator before the drive or the systemobject domain. For example, in the following option Tivoli Storage Manager will process all local drives except for the c: drive and systemobject domain:

```
domain ALL-LOCAL -c: -systemobject
```

Using the backup-archive command line client, you can specify drives to include in addition to your default domain. For example, if your default domain contains drives c: and d:, and you want to back up those drives as well as the diskette in drive a:, enter:

```
dsmc incremental -domain="a:"
```

You can also select **Actions** → **Backup Domain** from the Tivoli Storage Manager GUI to perform these backup functions.

## Backing up data using the command line

You can use the **incremental** or **selective** commands to perform backups. Table 20 shows examples of using commands to perform different tasks. For information on the commands to use for backing up system objects, see Chapter 10, "Using commands," on page 423.

Table 20. 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 475 for more information about the <b>incremental</b> command. See "Full and partial incremental backup" on page 65 for detailed information about incremental backups.
Back up the g: and h: drives in addition to the c:, d:, and e: drives defined in your client domain.	dsmc incremental -domain="g: h:"	See "Domain" on page 239 for more information about the <i>domain</i> option.
Back up all local volumes defined in your client domain <i>except</i> for the c: drive and systemobject domain.	dsmc incremental -domain="all-local -c: -systemobject"	You cannot use the (-) operator in front of the domain keyword all-local. See "Domain" on page 239 for more information. For Windows Server 2003 and Windows Vista you can also exclude the systemstate domain from backup processing in this way.

Table 20. Command line backup examples (continued)

Task	Command	Considerations
Back up all local volumes defined in your client domain <i>except</i> for the c: drive and systemstate domain.	<code>dsmc incremental -domain="all-local -c: -systemstate"</code>	You cannot use the (-) operator in front of the domain keyword all-local. See "Domain" on page 239 for more information.
Back up <i>only</i> the g: and h: drives.	<code>dsmc incremental g: h:</code>	None
Back up all files in the c:\Accounting directory and all its subdirectories.	<code>dsmc incremental c:\Accounting\* -sub=yes</code>	See "Subdir" on page 381 for more information about the <i>subdir</i> option.
Assuming that you initiated a snapshot of the C: drive and mounted the snapshot as the logical volume \\florence\c\$\snapshots\snapshot.0, 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 C:.	<code>dsmc incremental c: -snapshot= \\florence\c\$\snapshots\ snapshot.0</code>	See "Snapshotroot" on page 375 for more information.
<i>Incremental-by-date backup</i>		
Perform an incremental-by-date backup of your default client domain.	<code>dsmc incremental -incrbydate</code>	Use the <i>incrbydate</i> option with the <b>incremental</b> 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 286 for more information about the <i>incrbydate</i> option.
<i>Selective backups</i>		
Back up all files in the d:\proj directory.	<code>dsmc selective d:\proj\</code>	Use the <b>selective</b> command to back up specific files, a group of files with similar names, or empty directories and their attributes regardless of whether those files or directories were backed up during your last incremental backup and without affecting the backup server's last incremental backup count. You can use wildcards to back up multiple files at once. See "Selective" on page 559 for more information about the <b>selective</b> command.
Back up the d:\proj directory and all its subdirectories.	<code>dsmc selective d:\proj\ -subdir=yes</code>	See "Subdir" on page 381 for more information about the <i>subdir</i> option.



Table 20. Command line backup examples (continued)

Task	Command	Considerations
Back up the d:\h1.doc and d:\test.doc files.	<code>dsmc selective d:\h1.doc d:\test.doc</code>	You can specify as many file specifications as available resources or other operating system limits permit. Separate file specifications with a space. You can also use the <i>filelist</i> 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 262 for more information.
Back up a list of files in the c: drive.	<code>dsmc selective -filelist=c:\filelist.txt</code>	Use the <i>filelist</i> option to process a list of files. See “Filelist” on page 262 for more information.
Assuming that you initiated a snapshot of the C: drive and mounted the snapshot as the logical volume \\florence\c\$\snapshots\snapshot.0, run a selective backup of the c:\dir1\sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name C:.	<code>dsmc selective c:\dir1\sub1* -subdir=yes snapshot=\\florence\c\$\snapshots\snapshot.0</code>	See “Snapshotroot” on page 375 for more information.

See “Backup: Additional considerations” on page 97 for additional information to consider when performing a backup.

## 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. For example, you might need to delete sensitive data that was backed up (intentionally or unintentionally), and now needs to be removed from the server. Or you might need to delete files that were backed up, but were later found to contain viruses. To determine if you have the authority to delete individual backup copies from the Tivoli Storage Manager server without deleting the entire file space, 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. If you want to estimate the amount of time it takes to process your files and directories, click the **Estimate** button.
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 465 for more information.

---

## 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.

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 280 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 122 for more information about backup sets.

You can perform a full or differential backup using the *mode* option. See “Backup Group” on page 440 and “Mode” on page 307 for more information.

For example, to perform a full backup of all the files in the `c:\dir1\filelist1` file to the virtual file space `\virtfs`, containing the group leader `c:\group1` file, enter the following command:

```
dsmc backup group -filelist=c:\dir1\filelist1 -groupname=group1 -virtualfsname=\virtfs -mode=full
```

---

## Backup with client-node proxy support

Backups of multiple nodes that share storage can be consolidated to a common target node name on the Tivoli Storage Manager server. This is useful when the 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:

1. Install the backup-archive client on all nodes in a shared data environment.
2. Register each node with the Tivoli Storage Manager server, if it does not exist. Register the common target node name to be shared by each of the agent nodes used in your shared data environment.
3. Register each of the nodes in the shared data environment with the Tivoli Storage Manager server. This is the agent node name that is used for authentication purposes. Data will not be stored using the node name when the *asnodename* option is used.
4. Grant proxy authority to all nodes in the shared environment to access the target node name on the Tivoli Storage Manager server, using the **grant proxynode** command (Tivoli Storage Manager administrator).
5. Use the **query proxynode** administrative client command to display the authorized user's client nodes, granted by the **grant proxynode** command.

See "Asnodename" on page 194 for more information.

#### **Suggestions:**

- All agent nodes in the multiple node environment should be of the same platform type.
- Do not use target nodes as traditional nodes, especially if you encrypt your files before backing them up to the server.

## **Enabling multiple node operation from the GUI**

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.

Perform one of the following steps to verify that your client node is now accessing the server as the target node:

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

To return to single node operation, delete the **As Node Name** from the **Accessing As Node** field in the **General** → **Preferences** tab. **Restrictions enforced within a proxied session:**

- You cannot perform a system object or a system state object backup or restore.
- If a backupset contains system object or system state, it will not be restored 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.

## Setting up encryption

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.

### Important:

- If you change the encryption key, you must repeat the previous steps.
- Use the same encryption key for all files backed up in the shared node environment.

## Scheduling 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 by either configuring a Tivoli Storage Manager schedule service, or by 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.

Perform the following steps to enable scheduling of multiple nodes:

1. Ensure that all agent nodes must have proxy authority over the common target node
2. Ensure that all agent nodes must have a schedule defined on the server:

```
def sched domain_name sched_name options='-asnode=target'
```
3. Ensure that each agent node must have its schedule associated with a node:

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

---

## 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 375 for more information.

---

## Backing up system objects (Windows XP)

You can back up Windows XP system objects together or individually. Microsoft recommends that all system objects be backed up together to maintain a consistent system state. The following are valid system objects:

- Active Directory (domain controller only)
- Certificate server database
- Cluster Database (cluster node only)
- COM+ database
- Event logs (system, security and application)
- Windows Registry
- System and boot files
- System Volume
- Removable Storage Management Database (RSM)
- Replicated file systems (FRS)
- Windows Management Instrumentation (WMI)

You can exclude system objects from backup processing using the `exclude.systemobject` option. See “Exclude options” on page 257 for more information.

To back up system objects using the GUI:

1. Click **Backup** from the GUI main window. The Backup window appears.
2. Expand the directory tree by clicking the plus sign **+**. To display files in a folder, click the folder icon.

3. Locate the System Objects node in the directory tree and expand it.
4. Click the selection box for the system object(s) you want to back up.
5. Click **Backup**. The Backup Task List window displays the backup processing status. When processing completes, the Backup Report window displays processing details.

On the command line, use the **backup systemobject** command to back up all valid system objects. To back up system objects individually, see corresponding commands in Chapter 10, "Using commands," on page 423.

See "Backup: Additional considerations" on page 97 for additional information to consider when performing a backup. You do not need to understand this information to perform basic backups.

---

## Backing up Windows system state

Tivoli Storage Manager supports the Microsoft Volume Shadowcopy Service (VSS) on Windows Server 2003 and Windows Vista. Tivoli Storage Manager uses VSS to back up all system state components as a single object, to provide a consistent point-in-time snapshot of the system state. System state consists of all bootable system state and system services components.

**Note:** Allow additional time to back up system state components on a Windows Vista operating system, because of the number of files involved.

Bootable system state components include the following:

- Active Directory (domain controller only)
- System Volume
- Certificate Server Database
- COM+ database
- Windows Registry
- System and boot files

System services components include the following:

- Background Intelligent Transfer Service (BITS)
- Event logs
- Removable Storage Management Database (RSM)
- Cluster Database (cluster node only)
- Remote Storage Service
- Terminal Server Licensing
- Windows Management Instrumentation (WMI)
- Internet Information Services (IIS) metabase
- Dynamic Host Configuration Protocol (DHCP)
- Windows Internet Name Service (WINS)

The list of bootable system state and system services components are dynamic and can change depending on service pack and operating system features installed. Tivoli Storage Manager allows for the dynamic discovery and back up of these components.

You must have administrative authority to back up system state information.

To back up a system state object using the GUI:

1. Click **Backup** from the GUI main window. The Backup window appears.

2. Expand the directory tree by clicking the plus sign (+). To display files in a folder, click the folder icon.
3. Locate the system state node in the directory tree. You can expand the system state node to display the components.
4. Click the selection box next to the system state node to back up the entire system state object. You can back up the system state node only as a single entity because of dependencies among the system state components. By default, all components are selected; you cannot back up individual system state components.
5. Click **Backup**. The Backup Task List window displays the backup processing status. When processing completes, the Backup Report window displays processing details.

On the command line, use the **backup systemstate** command to back up all system state or system services components as a single object. See “Backup Systemstate” on page 453 for more information. Use the **query systemstate** command to display information about a backup of the system state on the Tivoli Storage Manager server. See “Query Systemstate” on page 517 for more information.

**Notes:**

1. Your Windows client must be connected to a Tivoli Storage Manager Version 5.2.0 or higher server.
2. System and boot files are backed up as a group only if one of the members of the group (one of the files) changes. If the files have not changed since the last backup, the system and boot files are not redundantly backed up.
3. You can use the *include.systemstate* option to assign management classes for back up of the system state. The default is to bind the system state object to the default management class. See “Include options” on page 280 for more information.
4. You can use the *exclude.systemservice* option in your client options file (dsm.opt) to exclude individual system services components from backup. Input can be the keyword or component name to be excluded from backup. See “Exclude options” on page 257 for more information.
5. You can use the *domain* option to exclude the entire system state from domain incremental backup processing. See “Domain” on page 239 for more information.
6. The system dllcache directory is now included in the boot partition backup of Windows 2003 and Windows Vista systems.

When the dllcache files are not available for restore of Windows 2003 or Windows Vista, system recovery might require availability of the operating system installation media. By backing up the dllcache directory, you can avoid the need for installation media during system restores.

If you do not want the dllcache directory included in the backup of your boot partition, and you understand the limitations of not backing up the dllcache directory, then you can use an exclude.dir statement to suppress backup of those files. For example:

```
exclude.dir c:\windows\system32\dllcache
```

See “Restoring Windows system state” on page 111 for instructions to restore a backup of the system state.

---

## Backing up Automated System Recovery (ASR) files (Windows XP, Windows Server 2003)

ASR is a restore feature of Windows XP Professional and Windows Server 2003 that provides a framework for saving and recovering the Windows XP or Windows Server 2003 operating state, in the event of a catastrophic system or hardware failure. Tivoli Storage Manager creates the files required for ASR recovery and stores them on the Tivoli Storage Manager server.

You must have administrative authority to back up ASR files. To generate and back up the ASR files using the Tivoli Storage Manager GUI:

1. Click **Backup** from the GUI main window. The Backup window appears.
2. Expand the directory tree by clicking the plus sign +. To display files in a folder, click the folder icon.
3. Click the selection box next to the Automated System Recovery node. You can also click the ASR selection box from the file list.
4. Click **Backup**. The Backup Task List window displays the backup processing status. When processing completes, the Backup Report window displays processing details.

Tivoli Storage Manager will generate the ASR files in the <system drive>:\adsm.sys\ASR staging directory on your local machine and store these files in the ASR file space on the Tivoli Storage Manager server.

You can restore these files to a diskette for use during ASR recovery. See “ASR preparation procedure” on page 116 for more information.

You can also use the **backup asr** command to back up ASR files to the Tivoli Storage Manager server. See “Backup ASR” on page 436 for more information.

---

## Performing an image backup

From your local workstation, you can back up a logical volume as a single object (image backup) on your system. The traditional *offline* image backup prevents write access to the volume by other system applications during the operation. These volumes can be formatted FAT, FAT32, NTFS, or unformatted RAW volumes. If a volume is NTFS-formatted, only those blocks used by the file system or smaller than the *imagegapsize* parameter will be backed up.

Normally you cannot restore an image backup of the system drive over itself since an exclusive lock of the system drive is not possible. However, when using WinPE, an image restore of the system drive is possible. See *Tivoli Storage Manager Recovery Techniques Using Windows Preinstallation Environment (Windows PE)* for more information.

You cannot restore an image backup to the volume from which the client is currently running. Consider installing the Tivoli Storage Manager client on the system drive.

Image backup does not guarantee consistency of system objects, such as the Active Directory. System objects can be spread out across multiple volumes, and should be backed up using the corresponding BACKUP commands.

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 *offline* image backup prevents write access to the volume by other system applications during the operation. When backing up an image using *snapshotproviderimage=none*, after a restore always run the fsck utility.

To restore an image backup of a volume, the Tivoli Storage Manager client must be able to obtain an exclusive lock on the volume being restored.

If the Tivoli Storage Manager online image support is configured, Tivoli Storage Manager performs an online image backup, during which the volume is available to other system applications. The snapshot provider, as specified by the *snapshotproviderimage* option, maintains a consistent image of a volume during online image backup. To install or configure online image support, see “Configuring online image backup support” on page 28.

To enable online image backups, see “Configuring online image backup support” on page 28. If you use VSS, you do not need to install LVSA.

You can use the *snapshotproviderimage* option with the **backup image** command or the *include.image* option to specify whether to perform an offline or online image backup. See “Snapshotproviderimage” on page 374 for more information.

## Before you perform an image backup

Before you perform an image backup, consider the following:

- *To perform an offline or online image backup you must have administrative authority on the system.*
- You do not need more than one drive to perform an image backup.
- Ensure that no other application is using the volume when you run an offline image backup. To ensure a consistent image during backup processing, the client will lock the volume, so that no other applications can write to it. If the volume is in use when the client attempts to lock the volume, the backup will fail. If the client cannot lock a volume because it is in use, you can perform an online image backup.
- 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 280 for more information. See Chapter 8, “Understanding storage management policies,” on page 153 for more information about management classes.

**Note:** If the *snapshotproviderimage* option is set to *none*, then the copy serialization parameters set by the management class will be used. See “Snapshotproviderimage” on page 374 for more information.

- You can exclude a volume from image backup using the *exclude.image* option. See “Exclude options” on page 257 for more information.

- You must use the mount point or drive letter for the volume on which you want to perform an image backup. Tivoli Storage Manager will not back up a volume without the use of a drive letter or mount point.
- Do not include the system drive in an image backup because the client cannot have an exclusive lock of the system drive during the restore and the system drive image cannot be restored to the same location. Image backup does not guarantee consistency of system objects, such as the Active Directory. System objects can be spread out across multiple volumes, and should be backed up using the corresponding backup commands. Because you cannot restore an image backup to the volume from which the client is currently running (or any volume for which an exclusive lock cannot be obtained) you should install your client program on the system drive.

**Note:** When using WinPE, an image restore of the system drive is possible. See *Tivoli Storage Manager Recovery Techniques Using Windows Preinstallation Environment (Windows PE)* for more information.

- If bad disk sectors are detected on the source drive during a LAN-free or LAN-based image backup, data corruption can occur. In this case, bad sectors are skipped when sending image data to the Tivoli Storage Manager server. If bad disk sectors are detected during the image backup, a warning message is issued after the image backup completes.

## 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 GUI” on page 71 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 86 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 121 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 an 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 86 for instructions.
2. Perform an incremental-by-date image backup of the file system (See "Performing an image backup using the GUI" on page 86 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" on page 86 for instructions).
4. Restore your volume by performing an incremental restore (See "Performing an image restore using the GUI" on page 121 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 of the incremental backups performed since that date.

**Note:** You should perform full image backups periodically in the following cases:

- When a file system changes substantially (more than 40%), as indicated in step 4 of method 1 and step 3 of method 2. On restore, this would provide a file system image close to what existed at the time of the last incremental-by-date image backup and it also improves restore time.
- As appropriate for your environment.

This 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 21 on page 86 is a comparison of methods 1 and 2.

Table 21. 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

If the Tivoli Storage Manager online image feature is configured, Tivoli Storage Manager performs a snapshot-based image backup, during which the real volume is available to other system applications. A consistent image of the volume is maintained during the online image backup.

When you perform an image backup using the client GUI **Image Backup** option, Tivoli Storage Manager honors the *snapshotproviderimage* setting in your client options file (dsm.opt). If the Tivoli Storage Manager online image support is configured, Tivoli Storage Manager performs an online image backup, during which the volume is available to other system applications. See “Snapshotproviderimage” on page 374 for more information.

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 an offline image backup, select **Image Backup** from the drop-down list.
  - To perform an online image backup, select **Snapshot Image Backup** from the drop-down list.
  - 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*.
- Because image backup allows you to back up only used blocks in a file system, the stored image size on the Tivoli Storage Manager server could be smaller than the volume size. For online image backups, the stored image can be larger

than the file system based on the size of the cache files. To determine the actual stored image size, select **View** → **File Details**. The actual stored image size is noted in the Stored Size field.

## 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 442 and “Restore Image” on page 536 for more information.

You can use the *snapshotproviderimage* option with the **backup image** command or the *include.image* option in your `dsm.opt` file or on the command line to specify whether to perform an offline or online image backup. See “Snapshotproviderimage” on page 374 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 307 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 8 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

## Support for CDP Persistent Storage Manager (PSM)

PSM is the snapshot technology that is included with a number of Microsoft Server Appliance Kit based NAS boxes including the IBM TotalStorage® NAS 200, 300, and 300G. You can use the Tivoli Storage Manager backup-archive client to back up the persistent images (PI) of a volume produced by PSM. You must first ensure that the volume has a label. You can then use PSM to schedule or create a persistent image with a specific image name, such as `snapshot.daily`, and set the number of images to save to 1. PSM will overwrite the PI as needed and you can use Tivoli Storage Manager to incrementally back up the PI. In this case, Tivoli Storage Manager will only back up the files that changed between snapshots. One advantage of backing up a PSM PI rather than the actual volume, is that there are no open files in the PI.

Considerations:

- By default, a PSM schedule uses a variable name (`snapshot.%i`) and keeps a number of images. It is not recommended that you use Tivoli Storage Manager with PSM in this way because Tivoli Storage Manager will see each image as unique and will make a complete copy of each.
- Tivoli Storage Manager requires that the volume used to make the PI has a label. If the volume does not have a label, Tivoli Storage Manager will not back up its PI.
- You use the Tivoli Storage Manager image backup function to back up the original volume used to create the PI, but you cannot use the backup image function to back up the PI.
- To avoid backing up unnecessary files when backing up PSM, include the following entries in your client option file (`dsm.opt`):

```
exclude.dir "Persistent Storage Manager State"
exclude.file "*.psm"
exclude.file "*.otm"
```

## 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 20.

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 client options file (dsm.opt) to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup. See “Toc” on page 399 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.

6. (Optional) To monitor processing of an operation from the GUI main window, open the **Actions** menu and select **TSM Activities**. During a backup, the status bar indicates processing status. A percentage estimate is not displayed for differential backups.

**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 126.

## Backing up NAS file systems using the command line

Table 22 lists the commands and options you can use to back up NAS file system images from the command line.

Table 22. 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.	242
<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.	257
<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 client options file (dsm.opt). See “Toc” on page 399 for more information. This option is valid for all Windows clients <i>only</i> .	280
<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.	507
<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.	447
<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.	399
<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.	485
<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.	460
<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.	491
<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.	499
<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.	468



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 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}. Do not use brace delimiters in the option file.

**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.

## Backup and restore of NAS file servers using CIFS

Tivoli Storage Manager can process NAS file server data that is accessed using the Common Internet File System (CIFS), but there are limitations to this approach:

1. There is no guarantee that the security information associated with the files and directories will be processed.
2. There will be performance degradation due to the fact that data is being accessed remotely. The Tivoli Storage Manager client reads the data over the network from the remote NAS file server, then sends the data over the network to the Tivoli Storage Manager server.
3. The mapped drives are presented to Tivoli Storage Manager as NTFS file systems, but in fact might not have full NTFS functionality and might show anomalies. For example, there have been instances where a file's encryption attribute is set, but when the client attempts to back up the file, it fails because the volume-level encryption setting indicates that the volume does not support encryption.

Tivoli Storage Manager will support backup and restore of NAS file servers with CIFS, but will not support problems related to these restrictions.

If the NAS file server is supported by Tivoli Storage Manager NDMP support, it is strongly recommended to use that support, rather than trying to back up and restore the volumes using remote mapped drives.

Tivoli Storage Manager supports the following methods for backup and recovery of data on NAS devices:

- One approach is to use a Tivoli Storage Manager backup-archive client to back up and restore data, by using CIFS or NFS to access files from the backup-archive client. Data can be stored on the Tivoli Storage Manager server with file-level granularity using progressive-incremental backup. The data is stored in the Tivoli Storage Manager storage hierarchy and can be migrated, reclaimed, and backed up to a copy storage pool. However, this approach introduces overhead for accessing individual files by the Tivoli Storage Manager client, requires data flow over a network and through the Tivoli Storage Manager client, requires data flow through the Tivoli Storage Manager server unless a LAN-free configuration is used, and might not preserve access-control lists because of limitations in the methods used for accessing files.

- A variation on the above approach is to use a Tivoli Storage Manager backup-archive client running on the NAS device, provided that the NAS operating system allows external programs. This method avoids the overhead of CIFS or NFS. Data can be stored on the Tivoli Storage Manager server with file-level granularity using progressive-incremental backup. The data is stored in the Tivoli Storage Manager storage hierarchy and can be migrated, reclaimed, and backed up to a copy storage pool. However, this approach requires data flow through the Tivoli Storage Manager client. This method also requires data flow over a network and through the Tivoli Storage Manager server unless a LAN-free configuration is used.
- Another method is to use the support for NDMP provided in Tivoli Storage Manager. File systems are backed up as full images (all files) or differential images (all files that have changed since the last full backup). Backed up images are stored on a tape device that is accessed by the NAS file server. This method provides high performance and scalability because there is no data flow over a network or through a Tivoli Storage Manager client or server. Data backed up to Tivoli Storage Manager using NDMP cannot be migrated, reclaimed, or backed up to a copy storage pool.

---

## Using VMware Consolidated Backup

VMware Consolidated Backup is a VMware backup solution for the ESX 3.0 server in a SAN environment. Consolidated backup allows the backup of multiple virtual machines to be offloaded to a dedicated physical host. By off-loading the backup away from virtual machines and the ESX server, it also allows backup in a LAN-free environment.

Before you can start using VMware Consolidated Backup, you need to prepare your systems by following these steps:

1. Implement your ESX server farm with SAN-attached storage, so that all virtual disks for all virtual machines are stored on the SAN.
2. Plan for a dedicated Windows 2003 server which has multiple SAN connections to act as the VMware Consolidated Backup backup proxy.  
**Recommendation:** If you want to use LAN-free, you need an additional SAN connection so that tape traffic can be separated from disk traffic.
3. Zone your SAN and configure your disk subsystem host mappings so that all ESX servers and the backup proxy can access the same disk volumes.
4. Follow all of the prerequisite and configuration VMware Consolidated Backup setup requirements and limitations that are documented in the VMware *Virtual Machine Backup Guide*, at: [http://www.vmware.com/pdf/vi3\\_vm\\_backup.pdf](http://www.vmware.com/pdf/vi3_vm_backup.pdf). Pay specific attention to the software and hardware requirements, and the "Before you Begin" sections.
5. Download and install, on your backup proxy machine, the VMware Consolidated Backup framework (*VMware Consolidated Backup for ESX Server*) from VMware, at: [http://www.vmware.com/download/vi/drivers\\_tools.html](http://www.vmware.com/download/vi/drivers_tools.html). The integration module (Consolidated Backup Integration Module for Tivoli Storage Manager) is not required.
6. Install the Tivoli Storage Manager client and run the setup wizard to establish basic communication with the Tivoli Storage Manager server on the backup proxy and each virtual machine. Pay attention to the Tivoli Storage Manager node names that are used on each backup proxy and virtual machine.  
**Important:** Your Tivoli Storage Manager administrator might need to register the nodes for you.

7. Grant proxy between the node for the VMware Consolidated Backup proxy, and all of the virtual machine nodes on the Tivoli Storage Manager Server. For example, you have agent=vcbproxynode and targets of vm1, vm2, through vmN. Here are the commands that you need to run on the Tivoli Storage Manager server:

```
grant proxy target=vm1 agent=vcbproxynode
grant proxy target=vm2 agent=vcbproxynode
...
grant proxy target=vmN agent=vcbproxynode
```

8. Update the options file on the VMware Consolidated Backup proxy to include the following options:
- VMCHost (see “Vmchost” on page 412)
  - VMCPW (see “Vmcpw” on page 413)
  - VMCUSER (see “Vmcuser” on page 414)
  - VMLIST (see “Vmlist” on page 415)

**Restriction:** Check the VMware documentation regarding the use and setup of the VirtualCenter. If you cannot use a VirtualCenter server and you need to perform backups of multiple machines on multiple ESX servers, do not specify these options in the options file, but specify the options with the command so that it can be varied for each ESX server.

9. Ensure that the VMware tools are installed on each virtual machine, from the VMware virtual infrastructure client.
10. (Optional) For each virtual machine, configure the VMware pre-freeze and post-thaw scripts that are run within each virtual machine (for most operating systems). These files can be created on the virtual machines: C:\Windows\pre-freeze-script.bat and C:\Windows\post-thaw-script.bat. Among other functions, this can be used to achieve application-consistent backup in Windows virtual machines. In addition to the pre-freeze and post-thaw scripts, VMware will quiesce NTFS and FAT file systems (only for Microsoft Windows virtual machines). This ensures that no file system writes are pending at the time the snapshot is taken, allowing the creation of file-system consistent backups.

11. (Optional) For each virtual machine, configure the Tivoli Storage Manager exclude list on the backup proxy. Define the Tivoli Storage Manager exclude list in the dsm.opt file of the virtual machine and in the dsm.opt file of the Backup Proxy machine for files that you do not want Tivoli Storage Manager to back up. The exclude list on the Backup Proxy machine can be set for all virtual machines or for a specific virtual machine. **The exclude list on the virtual machine must match the exclude list on the Backup Proxy if incremental backups are to be done from within the virtual machine.** In most cases the incremental backups will be run from the Backup Proxy machine.

Here is a sample exclude statement for a backup proxy:

```
exclude "\\mutalist\c$\temp\...\*"
exclude "C:\temp\...\*" - all virtual machines
```

Here is a sample exclude statement for a virtual machine:

```
exclude "\\mutalist\c$\...\*"
```

12. After you’ve completed the previous steps, test your systems by running some test backups on the proxy machine using the **backup VM** command (see “Backup VM” on page 457).

Back up a single virtual machine using the following command:

```
dsmc backup vm -vmlist=vm1
```

Back up all virtual machines in a single backup operation. Add each virtual machine to the *VMLIST* option in the *dsm.opt* file:

```
dsmc backup vm
```

Back up virtual machines in an environment where DNS name resolution is not available:

```
dsmc backup vm -vmlist=9.100.45.101[vm1],9.100.45.102[vm2]  
-vmchost=9.100.45.100 -vmcuser=virctr\admin -vmcpw=xxxxx
```

Back up virtual machines from several ESX servers in an environment where a virtual center server is not available:

```
dsmc backup vm -vmlist=vm1,vm2,vm3  
-vmchost=esxserver1 -vmcuser=root -vmcpw=xxxxx
```

```
dsmc backup vm -vmlist=vm4,vm5,vm6  
-vmchost=esxserver2 -vmcuser=root -vmcpw=xxxxx
```

```
dsmc backup vm -vmlist=vm7,vm8,vm9  
-vmchost=esxserver3 -vmcuser=root -vmcpw=xxxxx
```

**Note:** Only VMware virtual machines installed with Microsoft Windows operating systems are allowed with the **backup VM** command.

13. Perform some file level restores of these backups using the Tivoli Storage Manager Backup-Archive Client running within the virtual machine. For each virtual machine, ensure that the drive has been backed up and that the incremental data has been sent.
14. After you have prepared and run a test backup and restore of your systems, define the schedule on the Tivoli Storage Manager Server, and set up a daily schedule to run the **backup VM** command to process all virtual machines defined in the *VMLIST* option one at a time. Create the schedules as MACRO type schedule, with the action parameter of 'macro' and the object parameter of the macro file (for example: SCHED\_MACRO\_BackupVM.MACRO). The VMware options of *VMCHost*, *VMCUSER*, *VMCPW*, and *VMLIST* must be set in the client *dsm.opt* file of the Backup Proxy, or on the Tivoli Storage Manager Server schedule definition.

Here are some example schedule commands which are run from the Tivoli Storage Manager Server command line admin client:

```
define schedule STANDARD VM_DAILY_BACKUP action=macro  
objects="SCHED_MACRO_BackupVM.MACRO" period=24  
perunits=hours duration=4 durunits=hours
```

```
define association STANDARD VM_DAILY_BACKUP vcbproxynode
```

15. On the backup proxy machine (nodename vcbproxynode), use the GUI setup wizard to set up the client for Tivoli Storage Manager Scheduler service.
16. Create a file in the Tivoli Storage Manager client directory C:\Program Files\Tivoli\TSM\baclient\SCHED\_MACRO\_BackupVM.MACRO with one line, "backup vm", inside the Tivoli Storage Manager macro file.

This will process all virtual machines defined in the *VMLIST* option, one at a time, using the Tivoli Storage Manager client. The *resourceutilization* option can be used to control whether multiple backup sessions are used to process more than one virtual drive at a time using multiple sessions per virtual machine. Optionally, virtual machines list (*VMLIST*) can be processed concurrently by defining additional schedules and configuring multiple schedule services on the Backup Proxy machine to process separate *VMLIST* options.

The daily backup schedules should run from the Backup Proxy. Restore should be performed from the clients running inside the virtual machines.

---

## Backing up NetApp CIFS share definitions

Previous versions of the Tivoli Storage Manager Windows client (prior to Version 5.3.3) did not allow backing up and restoring NetApp CIFS share definitions. Note that the share definition includes share permissions set on the filer. Beginning with Tivoli Storage Manager Version 5.3.3, the Windows client now backs up the CIFS share definition under the root directory, the mapped CIFS share, or the UNC name. This support requires that the NetApp filer is running DATA ONTAP software which presents CIFS shares to remote clients as ordinary remote NTFS shares.

The root directory of a CIFS share is backed up with a full progressive incremental backup of the mapped drive/UNC name. See the following two examples:

```
net use x: \\NetAppFiler\CifsShareName
dsmc incr x:
dsmc incr \\NetAppFiler\CifsShareName
```

The following command line client output is displayed when the root directory (and share definition) is backed up:

```
Directory--> 0 \\NetAppFiler\CifsShare\ [Sent]
```

Refer to “Restoring NetApp CIFS shares” on page 130 for information about restoring NetApp CIFS shares.

---

## Backing up 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 (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 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 23. Client command line informational messages

Informational message	Meaning
Directory-->	Indicates the directory that you back up.
Updating-->	Indicates that only the file meta data is sent, not the data itself.
Expiring-->	Indicates an object (file or directory) on the server that no longer exists on the client is expired and made inactive on the server.
Subfile-->	Indicates that a file is processed using adaptive subfile backup. The file size indicated in this message is the size of the entire file, not the amount of data actually sent.
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 161 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 65 for more information.
Total number of objects failed:	Objects can fail for several reasons. Check the <code>dsmerror.log</code> for details.
Total number of subfile objects:	The number of files processed using adaptive subfile backup.
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.
Aggregate data transfer rate:	The average rate at which Tivoli Storage Manager and the network transfer data between the client and the server. This is calculated by dividing the total number of bytes transferred by the time that elapses from the beginning to the end of the process. Both Tivoli Storage Manager processing and network time are included in the aggregate transfer rate. Therefore, the aggregate transfer rate is lower than the network transfer rate.

**Note:** On occasion, the aggregate data transfer rate might be higher than the network data transfer rate. This is because the backup-archive client can have multiple simultaneous sessions with the backup server. If you set the *resourceutilization* option, Tivoli Storage Manager attempts to improve performance and load balancing by using multiple sessions when it backs up a volume 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.

Table 23. Client command line informational messages (continued)

Informational message	Meaning
Objects compressed by:	Specifies the percentage of data sent over the network divided by the original size of the file on disk. For example, if the net data-bytes are 10K and the file is 100K, then Objects compressed by: $== (1 - (10240/102400)) \times 100 == 90\%$ .
Subfile objects reduced by:	Expresses the savings achieved by adaptive subfile backup as a percentage of actual bytes sent versus actual file sizes. For example, if the actual sizes of the files backed up is 100 MB, but with adaptive subfile backup, only 25 MB were sent, then the reduction would be: $(1 - (25 \text{ MB}/100 \text{ MB})) \times 100 == 75\%$ . <b>Note:</b> The calculation considers all files that were processed, not only those files that were processed by adaptive subfile backup. The calculation does not include the effect of compression.
Elapsed processing time:	The active processing time to complete a command. This is calculated by subtracting the starting time of a command process from the ending time of the completed command process.
Total number of bytes transferred:	As indicated.
LanFree bytes transferred:	The total number of data bytes transferred during a lan-free operation. If the <i>enablelanfree</i> option is set to <i>no</i> , this line will not appear.

## Backup: Additional considerations

This section discusses additional information to consider when performing a backup. You do not need to understand this information to perform basic backups.

### Backing up open files

Some files on your system might be in use when you try to back them up. These are called *open files* because they are locked by an application for its exclusive use. It is not very common for files to be opened in this 'locked' mode. An application can open a file in this way to avoid other applications or users from reading or accessing the file, but it has the side effect of preventing backup products from reading the file for backup.

It is not always desirable to use the Open file support feature to backup open or 'locked' files. There might be cases where an application opens a file or group of files in this 'locked' mode to prevent the access of these files in an inconsistent state. In these cases the following points should be considered. These considerations could also be used on platforms where the Open file support feature is not available or is not in use, to avoid the overhead of creating volume snapshot for each backup.

- If the file is unimportant or can be easily rebuilt (a temporary file for example), you might not care if the file is backed up, and might choose to exclude it.
- If the file is important:
  - Ensure the file is closed before backing it up. If backups are run according to a schedule, use the *preschedulecmd* option to enter a command that closes the file. For example, if the open file is a database, issue a command to close the database. You 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, 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 changes 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 a management class with **dynamic** or **shared dynamic** serialization. See “Selecting a management class for files” on page 159 for information about assigning management classes, and see “Displaying information about management classes and copy groups” on page 155 for information to determine which management classes are available to you.

**Note:** If open file support is not configured: While Tivoli Storage Manager attempts to back up open files, this is not always possible. Some files are open exclusively for the application that opened them. If Tivoli Storage Manager encounters such a file, it cannot read it for backup purposes. If you are aware of such file types in your environment, you should exclude them from backup to avoid seeing error messages in the log file.

## 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 that you assign to the file using an *include* option in the include-exclude options list. The selected management class must contain a backup copy group for the file to be backed up.

Select **Utilities** → **View Policy Information** from the backup-archive client or Web client GUI to view the backup policies defined by the Tivoli Storage Manager server for your client node. See Chapter 7, “Automating tasks,” on page 145 and Chapter 8, “Understanding storage management policies,” on page 153 for more information on management classes, how an Authorized User assigns management classes to files, and storage management policies.

## Understanding how deleted file systems are handled

When a file system or drive has been deleted, or it is no longer backed up by Tivoli Storage Manager, the existing backup versions for each file are managed according to the following policy attributes:

- Number of days to keep inactive backup versions
- Number of days to keep the last backup version (if there is no active version)

If you do nothing else, active backup versions remain indefinitely. If you do not need to keep the active versions indefinitely, use the **expire** command to inactive the active versions.

You can also use the **delete backup** command to delete individual backup versions, or the **delete filespace** command to delete the entire file space. Your Tivoli Storage Manager server administrator must give you “delete backup” authority to use these commands. If the file space also contains archive versions, you must also have delete archive authority to use **delete filespace**. Use the **query session** command to determine whether you have delete backup and delete archive authority. Alternatively, you can ask your Tivoli Storage Manager server administrator to delete the file space for you.

Note that deletion of a file system has no effect on existing archive versions. However, if you no longer require the archive versions, you can use the **delete archive** or **delete filespace** commands to delete unneeded archives.



See Chapter 8, “Understanding storage management policies,” on page 153 for further information about policy management.

## Backing up removable media

Tivoli Storage Manager backs up your removable media (such as tapes, cartridges or diskettes) based on the drive label, not the drive letter. If a drive has no label, the backup does not occur. This use of drive labels permits you to perform such tasks as backing up different diskettes from the a: drive.

For a restore or retrieve, Tivoli Storage Manager maintains a separate file space for each drive label. These labels become the file space names on the server. If you change the label of a drive you already backed up, Tivoli Storage Manager views it as a new drive and does not relate it to your previous drive.

Because Tivoli Storage Manager uses the labels to manage backups and archives of your removable media, you occasionally need to use those labels to locate data when using commands. For example, if you try to restore a file on diskette or CD-ROM using `d:\projx\file.exe` as a file name, Tivoli Storage Manager substitutes the current label of your d: drive for the d:. If the d: drive label is d-disk, `d:\projx\file.exe` becomes `{d-disk}\projx\file.exe`, and the label is enclosed in braces.

If the label of the d: drive does not match a file space name on the server, Tivoli Storage Manager cannot locate your files using the current d: drive's label. However, Tivoli Storage Manager can locate your files if you use the file space name based on the original drive label. A mismatch between a label and a file space name might occur if you label your drives again, or if you access Tivoli Storage Manager from a different workstation than the one from which you backed up the files. If you have not relabeled the drive, and you are at the same workstation where the file was backed up, then you can use the drive letter as a shorthand version of the file space name (drive label).

## Backing up fixed drives

Tivoli Storage Manager can back up your fixed drives even if they do not have a label, including drive aliases created with the DOS `subst` command. This applies to both the drive alias and the underlying physical drive, because the alias name and the physical drive label are the same.

## Backing up NTFS file spaces

When you back up files on NTFS partitions, Tivoli Storage Manager also backs up file security information and these file descriptors:

- Owner security information (SID)
- Primary group SID
- Discretionary access control list (permissions)
- System access control list (auditing information)

You must specify a mixed or lowercase NTFS file space name enclosed in quotes and braces. For example, `{"NTFSDrive"}`. Single quotes or double quotes are valid in loop mode. For example: `{"NTFSDrive"}` and `{'NTFSDrive'}` are both valid. In batch mode, only single quotes are valid. The single quotes requirement is a restriction of the operating system.

## Using Universal Naming Convention names

A Universal Naming Convention (UNC) name is a network resource name for a share point on a workstation. The resource name includes the workstation name assigned to the workstation and a name you assign to a drive or directory so that it can be shared. The name you assign is also called a *share point name*.

### Using UNC names in domain lists

The following rules apply when using UNC names to specify a domain list. You must specify:

- A drive letter for removable media
- Drive letters or UNC name for local fixed drives
- Drive letters or UNC names for remote mapped drives
- UNC names for remote unmapped drives

Example 1: To specify drive a: containing removable media, enter

```
domain a: \\local\c$
```

Example 2: To specify fixed drive c:, enter

```
domain c: \\remote\share1 \\remote\c$
```

### Backing up specific files using the Universal Naming Convention

You can back up shared files in a network through the use of a UNC name. A UNC name is a network resource name for a share point on a workstation. The resource name includes the workstation name assigned to the workstation and a name you assign to a drive or directory so that it can be shared. The name you assign is also called a share point name.

Using a UNC name permits you to back up specific shared directories to a separate file space. This is useful if, for example, you or an administrator want to back up a small portion of data that you would otherwise be unable to access. Drives are not backed up to a separate file space.

Every local drive is accessible using a UNC name except for drives containing removable media (such as tapes, cartridges or diskettes). Access these drives by using a predefined administrative share name consisting of the workstation name and the local drive letter, followed by \$. For example, to specify a UNC name on the c: drive for workstation ocean, enter:

```
\\ocean\c$
```

The \$ sign *must* be included with the drive letter.

To enter a UNC name for workstation ocean and share point wave, enter:

```
\\ocean\wave
```

When accessing files, you do not need to enter the letter of the drive except for drives containing removable media.

See Table 24 on page 101 for examples showing selective backup of files using UNC names. In these examples, assume that:

- The workstation running **dsmc** is major.
- Share names betarc and testdir from workstation alpha1 are mapped to drives **r** and **t**, respectively.

Table 24. UNC examples

Example	Comment
dsmc sel \\alpha1\c\$\	name of remote file space is \\alpha1\c\$
dsmc sel \\major\c\$\	name of local, fixed file space is \\major\c\$
dsmc sel a:\	name of local, removable file space is volume label of a:
dsmc sel \\alpha1\betarc\	name of remote file space is \\alpha1\betarc
dsmc sel \\alpha1\testdir\	name of remote file space is \\alpha1\testdir
dsmc sel d:\	name of local, fixed file space is \\major\d\$
dsmc sel c:\	file space name is \\major\c\$
dsmc sel r:\	file space name is \\alpha1\betarc

You can also specify UNC names for files in your include-exclude and domain lists; see “Creating an include-exclude list (optional)” on page 31 and “Domain” on page 239 for more information.

## Backing up Microsoft Dfs files

### Notes:

1. This feature applies to Windows Server 2003.
2. See the product README file for current limitations of this feature.

You can use the following recommended methods for protecting the data in your Microsoft Dfs environment:

1. Back up the Dfs link metadata and the actual data at the share target of each link from the machine hosting the Dfs root. This method simplifies back up and restore by consolidating all of the Tivoli Storage Manager activities on a single machine. This method has the disadvantage of requiring an additional network transfer during backup to access the data stored at link targets.
2. Back up only the Dfs link metadata that is local to the machine hosting the Dfs root. Back up the data at the target of each link from the machine(s) which the data is local too. This method increases back up and restore performance by eliminating the extra network transfer, but requires Tivoli Storage Manager back up and restores to be coordinated among several machines.

Files contained on a Dfs server component are accessed using a standard UNC name, for example:

```
\\servername\dfsroot\
```

where *servername* is the name of the host computer and *dfsroot* is the name of the Dfs root.

If you set the *dfsbackupmntpnt* option to *yes* (the default), an incremental backup of a Dfs root does not traverse the Dfs junctions. Only the junction metadata is backed up. This is the recommended setting so that Tivoli Storage Manager can be used to restore Dfs links.

You can use the *dfsbackupmntpnt* option to specify whether Tivoli Storage Manager sees a Dfs mount point as a Microsoft Dfs junction or as a directory. For more information on using this option, see “Dfsbackupmntpnt” on page 233.

**Recommendation:** Restore the Dfs junction metadata first. This will recreate the links. Then restore each junction and the data at each junction separately. If you do not restore the junction metadata first, Tivoli Storage Manager creates a directory under the Dfs root using the same name as the junction point and restores the data in that directory.

The following example relates to method 1 above and illustrates how to use Tivoli Storage Manager to back up and restore a Microsoft Dfs environment. Assume the existence of a domain Dfs environment hosted by the machine itanium4:

**Dfs root**

```
\\itanium4\snj64test
```

**Dfs link1**

```
\\itanium4\snj64test\tools
```

**Dfs link2**

```
\\itanium4\snj64test\trees
```

Backup procedure:

1. Set the *dfsbackupmntpnt* option to *yes* in your client options file (dsm.opt).
2. Enter the following command to back up link junction information:

```
dsmc inc \\itanium4\snj64test
```

3. Enter the following command to back up data at the tools link:

```
dsmc inc \\itanium4\snj64test\tools
```

4. Enter the following command to back up data at the trees link:

```
dsmc inc \\itanium4\snj64test\trees
```

**Note:** DFS™ Replication uses staging folders to act as caches for new and changed files to be replicated from sending members to receiving members. If you do not want to backup these files, you can exclude them from your backup using the *exclude.dir* option.

```
exclude.dir x:\...\Dfsrprivate
```

Restore procedure:

1. Manually recreate shares at target machines only if they no longer exist.
2. Manually recreate the Dfs root using the exact name as it existed at the time of back up.

3. Enter the following command to recover data from the tools link. This step is not necessary if the data still exists at the link target:

```
dsmc restore \\itanium4\snj64test\tools\* -sub=yes
```

4. Enter the following command to recover data from the trees link. This step is not necessary if the data still exists at the link target:

```
dsmc restore \\itanium4\snj64test\trees\* -sub=yes
```

5. Use the Distributed File System management console snap-in to reestablish replication for each link, if necessary.

Tivoli Storage Manager limitations:

- Tivoli Storage Manager does not restore root of Dfs. To recreate the Dfs tree, manually create the Dfs root first, then start restore to recreate the links.
- Tivoli Storage Manager can back up the Dfs tree (both domain based Dfs and stand alone Dfs) hosted on local machine only. You cannot back up Dfs if the Dfs host server is not your local machine.

- Tivoli Storage Manager cannot recreate shared folders on restore. For example, if you delete the junction and the shared folder the junction points to, restoring the Dfs root will recreate the Dfs junction, but restoring a junction creates a local folder instead of creating the original backed up shared network folder.
- If a Dfs link is created with replica and the replica share is on a different server, then Tivoli Storage Manager does not display the replica data.
- If a Dfs root is added or modified, Tivoli Storage Manager will not back it up. You must specify the Dfs root in the *domain* option in the client options file (dsm.opt) regardless of whether DOMAIN ALL-LOCAL is specified.



---

## Chapter 5. Restoring your data

Use the backup-archive client to store backup versions of your files on the Tivoli Storage Manager server. You can restore these backup versions if the original files are lost or damaged. This chapter discusses various ways you can restore your data.

All client backup and 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 51 for information on starting the Web client.

*Table 25. Restore: Primary tasks*

Task	Page
Restoring files and directories	107
Restoring system objects	111
Restoring Windows system state	111
Restoring Automated System Recovery (ASR) files (Windows XP, Windows Server 2003)	113
Restoring Microsoft Dfs trees and files	120
Restoring an image	120
Restoring data from a backupset	122
Performing point-in-time restores	125
Restoring NAS file systems	126
Restoring the WebSphere Application Server	131
Authorizing another user to restore or retrieve your files	132
Restoring or retrieving files from another client node	133
Restoring or retrieving your files to another workstation	134
Restoring or retrieving files to another type of workstation	134
Restoring a disk in the event of disk loss	134
Deleting file spaces	135

---

### Duplicate file names

If you attempt to restore or retrieve a file whose name is the same as an existing file's short name, a file name collision will occur (existence of duplicate file names). An example is when the file *abcdefghijk.doc* has a short name of *abcdef~1.doc*, and you attempt to restore or retrieve a file explicitly named *abcdef~1.doc* into the same directory. In this case, a collision will occur because the name of the file you are restoring conflicts with the short name for *abcdefghijk.doc*.

A collision can occur even if the files are restored or retrieved to an empty directory. For example, files *abcdef~1.doc* and *abcdefghijk.doc* might originally have existed in the directory as *abcdefghijk.doc* and *abcdef~2.doc*. During the restore, if

*abcdefghijkl.doc* is restored first, it is assigned a short name of *abcdef~1.doc* by the Windows operating system. When you restore *abcdef~1.doc*, the duplicate file name situation occurs.

Tivoli Storage Manager handles these situations based on the value of the *replace* option. Use the **replace** option to specify whether to overwrite an existing file, or to prompt you for your selection when you restore or retrieve files. See “Replace” on page 340 for more information.

If a file name collision occurs, you can do any of the following:

- Restore or retrieve the file with the short file name to a different location.
- Stop the restore or retrieve and change the name of the existing file.
- Disable short file name support on Windows.
- Do not use file names, such as *abcdef~1.doc*, that would conflict with the short file naming convention.

---

## Restoring files using Universal Naming Convention names

Using a UNC name permits you to restore specific shared files to a separate file space. This is useful if, for example, you or an administrator want to restore a portion of data that you would otherwise be unable to access.

Except for drives with removable media, every local drive letter is accessible using a local UNC name that includes the workstation name and a designation for the drive letter. For example, to enter a UNC name on drive c: for workstation ocean, enter:

```
\\ocean\c$
```

The \$ sign *must* be included with the drive letter.

To enter a UNC name for workstation ocean and share point wave, enter:

```
\\ocean\wave
```

When accessing files, you do not need to enter the letter of the drive *except* for drives with removable media.

---

## 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 both an active and inactive version of a file at the same 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 277 for more information.



---

## Restoring files and directories

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

**Note:** When restoring a directory, its modification date and time is set to the date and time of the restore, not to the date and time the directory had when it was backed up. This is because Tivoli Storage Manager restores the directories first, then adds the files to the directories.

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

### Restoring data using the GUI

To restore files and directories using the client GUI, perform the following steps:

1. Click **Restore** on the main window. The Restore window appears.
2. Expand the directory tree by clicking the plus (+) sign or the folder icon next to an object in the tree. Select the object that you want to restore. To search or filter files, click the **Search** icon from the tool bar.
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 appropriate information.
5. Click **Restore**. The Restore Task List window displays the processing status.

### Restoring data using the command line

You can use the **restore** command to restore files. See “Restore” on page 522 for more information about the **restore** command. For information about restoring system objects using commands, see Chapter 10, “Using commands,” on page 423.

Table 26 shows examples of using the **restore** command to restore objects from Tivoli Storage Manager server storage. See “Restore” on page 522 for additional examples.

Table 26. Command line restore examples

Task	Command	Considerations
Restore the most recent backup version of the c:\doc\h1.doc file, even if the backup is inactive.	<code>dsmc restore c:\doc\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 298 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 c:\project\* -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 322 and “Inactive” on page 277 for more information.

Table 26. Command line restore examples (continued)

Task	Command	Considerations
Restore all files with a file extension of .c from the c:\devel\projecta directory.	dsmc restore c:\devel\projecta\*.c	If you do not specify a destination, the files are restored to their original location.
Restore the c:\project\doc\h1.doc file to its original directory.	dsmc restore c:\project\doc\h1.doc	If you do not specify a destination, the files are restored to their original location.
Restore the c:\project\doc\h1.doc file under a new name and directory.	dsmc restore c:\project\doc\h1.doc c:\project\newdoc\h2.doc	None
Restore the files in the e: drive and all of its subdirectories.	dsmc restore e:\ -subdir=yes	You must use the <i>subdir</i> option to restore directory attributes/permissions. See "Subdir" on page 381 for more information about the <i>subdir</i> option.
Restore all files in the c:\mydir directory to their state as of 1:00 PM on August 17, 2002.	dsmc restore -pitd=8/17/2002 -pitt=13:00:00 c:\mydir\	See "Pitdate" on page 323 and "Pittime" on page 324 for more information about the <i>pitdate</i> and <i>pittime</i> options.
Restore the c:\doc\h2.doc file to its original directory on the workstation, named <i>star</i> .	dsmc restore c:\doc\h2.doc \\star\c\$\  To restore the file to <i>star</i> which has been renamed <i>meteor</i> , enter: dsmc restore \\star\c\$\doc\h2.doc \\meteor\c\$\  You could also enter: dsmc restore \\star\c\$\doc\h2.doc c:\  This example is valid because if the workstation name is not included in the specification, the local workstation is assumed ( <i>meteor</i> , in this case).	For the purposes of this manual, the workstation name is part of the file name. Therefore, if you back up files on one workstation and you want to restore them to another workstation, you must specify a destination. This is true even if you are restoring to the same physical workstation, but the workstation has a new name.
Restore a file that was originally backed up from the diskette <b>workathome</b> in the a: drive, and restore it to a diskette in the a: drive labeled <b>extra</b> .	dsmc restore {workathome}\doc\h2.doc a:\doc\h2.doc	If you are restoring a file to a disk with a different label than the disk from which the file was backed up, you must use the file space name (label) of the backup disk instead of the drive letter.
Restore files specified in the c:\filelist.txt file to the d:\dir directory.	dsmc restore -filelist=c:\filelist.txt d:\dir\	See "Filelist" on page 262 for more information about restoring a list of files.
Restore all members of the virtfs\group1 group backup stored on the Tivoli Storage Manager server.	dsmc restore group {virtfs}\group1	See "Restore Group" on page 534 for more information.

## Restoring large amounts of data

If you need to restore a large number of files, you will get faster performance using the command line interface rather than the GUI interface. In addition, you will improve performance if you enter multiple **restore** commands at one time. For example, to restore all the files in your c: file space, enter:

```
dsmc restore c:\* -subdir=yes -replace=all -tapeprompt=no
```

However, if you enter multiple commands for the root directories in your c: file space, you can restore the files faster. For example, enter these commands:

```
dsmc restore c:\users\ -subdir=yes -replace=all -tapeprompt=no
dsmc restore c:\data1\ -subdir=yes -replace=all -tapeprompt=no
dsmc restore c:\data2\ -subdir=yes -replace=all -tapeprompt=no
```

Or, if you need to restore files for multiple drives, enter these commands:

```
dsmc restore c:\* -subdir=yes -replace=all -tapeprompt=no
dsmc restore d:\* -subdir=yes -replace=all -tapeprompt=no
dsmc restore e:\* -subdir=yes -replace=all -tapeprompt=no
```

You can also use the *quiet* option with the **restore** command 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 options file, it is not necessary to include these 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. Do not use any overlapping file specifications in the commands.

To display a list of the root directories in a file space, use the **query backup** command. For example:

```
dsmc query backup -dirsonly -subdir=no c:\
```

As a general rule, you can enter two to four **restore** commands at one time. The maximum number you can run at one time without degrading performance depends on factors such as network utilization and how much memory you have. For example, if \users and \data1 are on the same tape, the restore for \data1 must wait until the restore for \users is complete. However, if \data2 is on a different tape, and there are at least two tape drives available, the restore for \data2 can begin at the same time as the restore for \users.

The speed at which you can restore the files also depends upon how many tape drives are available and whether your administrator is using collocation to keep file spaces assigned to as few volumes as possible. If your administrator is using collocation, the number of sequential access media mounts required for restore operations is also reduced.

### **Standard query restore, no query restore, and restartable restore**

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

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

1. The client queries the server for a list of files backed up for the client file space you want to restore.
2. The server sends a list of backed up files that match the restore criteria. If you want to restore both active and inactive files, the server sends information about all backed up files to the client.
3. The list of files returned from the server is sorted in client memory to determine the file restore order and to minimize tape mounts required to perform the restore.
4. The client tells the server to restore file data and directory objects.

5. The directories and files you want to restore are sent from the server to the client.

**No query restore process:** The no query restore process is outlined below.

1. The client tells the server that a no query restore is going to be performed and provides the server with details about file spaces, directories, and files.
2. The server sorts the data using an internal sort table which minimizes tape mounts.
3. The data to be restored is sent to the client. File and directory objects stored on disk are sent immediately since sorting for such data is not required before restoring it.
4. You can use multiple sessions to restore the data. If the data resides on multiple tapes, there are multiple mount points available at the server. The combination of using the *resourceutilization* option and MAXNUMMP allows multiple sessions. See “Resourceutilization” on page 344 for more information.

When you enter an unrestricted wildcard source file specification on the **restore** command and do not specify any of the options: *inactive*, *latest*, *pick*, *fromdate*, *todate*, or *volinformation*, the client uses a *no query restore* method for restoring files and directories from the server. This method is called *no query restore* because instead of querying the server for each object to be restored, a single restore request is sent to the server. In this case, the server returns the files and directories to the client without further action by the client. The client merely accepts the data coming from the server and restores it to the destination named on the **restore** command.

Using the command-line client, an example of an unrestricted wildcard command would be:

```
c:\mydocs\2004\*
```

An example of a restricted wildcard file specification would be:

```
c:\mydocs\2004\sales.*
```

**Restartable restore process:** If the restore process stops because of a power outage or network failure, the server records the point at which this occurred. This record is known to the client as a *restartable restore*. It is possible to have more than one restartable restore session. Use the **query restore** command 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 522. To perform restartable restores using the GUI, follow these steps:

1. Select **Actions** → **Restartable restores** from the main panel.
2. Select the restartable restore session you want to complete.

3. Click the **Restart** button at the bottom of the panel.

## Restoring migrated files

When you restore migrated files, you also need temporary space for the resident version of these files. If you do not specify enough space for the migrated files and the resident versions of the files combined, the operation might fail. Refer to *IBM Tivoli Storage Manager for HSM for Windows Administration Guide* for details about the *restoremigstate* option.

---

## Restoring system objects

Certain Windows XP system objects must be restored together in order to create a consistent system state.

The following system objects should be restored together:

- Active Directory (domain controller only)
- Certificate server database
- Cluster database (cluster node only)
- COM+ database
- Windows Registry
- System and boot files
- System volume

Other system objects can be restored individually. These include the following:

- Event logs (system, security and application)
- Removable Storage Management Database (RSM)
- Replicated file systems (FRS)
- Windows Management Instrumentation (WMI) repository

To restore files and directories using the client GUI, perform the following steps:

1. Click **Restore** on the main window. The Restore window appears.
2. Expand the directory tree by clicking the plus (+) sign or the folder icon next to an object in the tree.
3. Locate the System Objects node in the directory tree and expand it.
4. Click the selection box next to the system object(s) that you want to restore.
5. To modify specific restore options, click the **Options** button. Any options you change are effective during the current session *only*.
6. Click **Restore**. The Restore Task List window displays the processing status.

On the command line, use the **restore systemobject** command to restore all valid system objects. To restore system objects individually, see corresponding commands in Chapter 10, "Using commands," on page 423.

---

## Restoring Windows system state

Tivoli Storage Manager supports the Microsoft Volume Shadowcopy Service (VSS) on Windows Server 2003 and Windows Vista. Tivoli Storage Manager uses VSS to restore the system state. By default, all system state components are restored together. However, you can restore just the bootable system state components or individual system services components. Restoring an individual system services component will restore only a specific system service and not necessarily your Windows operating system.

During system state restore, all data is restored to their original locations except for files that are protected by System File Protection (SFP). SFP files are restored to a temporary location. During reboot, the original SFP files are deleted, and the restored SFP files are renamed to the original file names.

Extra disk space is required to restore SFP files because they are restored to a temporary location on your disk. On Windows Server 2003, most of the system state data is SFP files, so a lot of free space is required. If there is absolutely no extra space on the system drive, contact your service representative.

Bootable system state components can include the following:

- Active Directory (domain controller only)
- System volume (domain controller only)
- Certificate Server Database
- COM+ database
- Windows Registry
- System and boot files

System services components can include the following:

- Background Intelligent Transfer Service (BITS)
- Event logs
- Removable Storage Management Database (RSM)
- Cluster Database (cluster node only)
- Remote Storage Service
- Terminal Server Licensing
- Windows Management Instrumentation (WMI)
- Internet Information Services (IIS) metabase
- DHCP database
- Wins database

**Attention:** Restoring system state in a situation other than system recovery is not recommended.

You must have administrative authority to restore Windows system state information using the GUI:

1. Click **Restore** from the GUI main window. The Restore window appears.
2. Expand the directory tree by clicking the plus sign **+**. To display files in a folder, click the folder icon.
3. Locate the system state node in the directory tree. You can expand the system state node to display the components.
4. Click the selection box next to the system state node to restore the entire system state. You can restore the bootable system state node only as a single entity because of dependencies among the bootable system state components. By default, all components are selected; you can restore individual system services components.
5. Click **Restore**. The Task List window displays the restore processing status.

On the command line, use the **restore systemstate** command to restore a backup of system state. See “Restore Systemstate” on page 547 for more information.

Considerations:

- You can restore system state data to an alternate machine.

- If you are upgrading from a Windows 2000 machine to a Windows 2003 machine, you cannot restore the Windows 2000 system objects that were backed up to the server.
- If you are upgrading from a Windows XP machine to a Windows Vista machine, you cannot restore the Windows XP system objects that were backed up to the server.
- Your Windows client must be connected to a Tivoli Storage Manager Version 5.2.0 or higher server.
- If Active Directory is installed, you must be in Active Directory restore mode.
- See “Performing a Windows XP or Windows Server 2003 system recovery” on page 114 for procedures on how to perform the following tasks:
  - Your operating system is still functioning, but a complete system restore is required.
  - A complete recovery is required, including an operating system re-installation.

---

## Restoring Automated System Recovery (ASR) files (Windows XP, Windows Server 2003)

ASR is a restore feature of Windows XP Professional and Windows Server 2003 that provides a framework for saving and recovering the Windows XP or Windows Server 2003 operating state, in the event of a catastrophic system or hardware failure. Tivoli Storage Manager creates the files required for ASR recovery and stores them on the Tivoli Storage Manager server.

You must have administrative authority to restore ASR files. To restore ASR files using the GUI:

1. Click the **Restore** button in the GUI main windows. The Restore window appears.
2. Expand the directory tree by clicking the plus sign + or the folder icon next to an object in the tree.
3. Click the selection box next to the Automated System Recovery node or click on the ASR selection box from the file list.
4. Click **Restore**. The Restore Destination window appears. Enter the appropriate information. You can choose to restore the ASR files to their original destination (<system drive>:\adsm.sys\ASR) or to a different location, such as a diskette (see “ASR preparation procedure” on page 116.)
5. Click **Restore**. The Restore Task List window displays the processing status.

### Notes:

1. When you restore inactive ASR files, you must exercise care and ensure that the restore of system state and files to the same time is performed. This requires manual modification of the tsmasr.cmd file because the default behavior is to restore the system to the last active backup. The restore commands in tsmasr.cmd do not specify inactive point-in-time parameters.
2. When you select Automated System Recovery from the directory tree, the results are displayed in the file list. Any inactive version of a particular file is indicated by the icon marked with an X next to the selection box.

On the Tivoli Storage Manager command line, use the **restore asr** command to restore the Automated System Recovery (ASR) files from the Tivoli Storage Manager server to a specified location. By default, the ASR files are restored to the adsm.sys\ASR staging directory. See “Restore ASR” on page 528 for more information.

---

## Performing a Windows XP or Windows Server 2003 system recovery

This section discusses how you can use the Tivoli Storage Manager client in conjunction with Automated System Recovery (ASR) on Windows XP or Windows Server 2003 to recover your Windows system in the event of a catastrophic system or hardware failure. If your operating system is still functioning, but a complete system restore is required, see “Operating system functions but requires complete system restore.”

The scope of this publication does not include disaster recovery planning. See Appendix E, “ASR supplemental information,” on page 625 for references to additional resources.

The functions described in this section are also documented in the Tivoli field guide, *Using Microsoft Windows Automated System Recovery (ASR) to Recover Windows XP and Windows 2003 Systems with the IBM Tivoli Storage Manager Backup-Archive Client for Windows*, available at the following Web site:

[http://www.ibm.com/software/sysmgmt/products/support/Field\\_Guides.html](http://www.ibm.com/software/sysmgmt/products/support/Field_Guides.html)

### Operating system functions but requires complete system restore

If your operating system is still functioning, but a complete system restore is required, use the following procedure. If Active Directory is installed, you must be in Active Directory restore mode. When performing a system recovery including the system state, the following restore order is recommended. Do not restart the system between each step, even though you are prompted to do so:

1. Restore the system drive. For example:  

```
dsmc restore c:\* -sub=yes -rep=all
```
2. Restore the system state as follows:  

```
dsmc restore systemstate
```

### Complete recovery required, including operating system re-installation

This section describes ASR hardware and software requirements and how to perform a complete recovery, including an operating system re-installation.

#### Windows ASR hardware requirements

To perform a successful system recovery:

1. The hardware configuration of the target system must be identical to that of the original system, with the exception of the hard disks, video cards, and network interface cards.  
**Note:** If there is a change in the video cards or networking cards, they must be manually configured.
2. The target system must have the same number of disks as the original system.
3. The size of each disk on the target system must be greater than or equal to the size of the corresponding disk on the original system. Note that disks that appear to be the same size might not have the same number of sectors due to geometry differences. Choose a disk on the target system which has identical disk geometry as the original system or choose a larger disk on the target system.
4. All disks must have 512 byte sectors.



5. All disks (and only such disks) that are available to ASR on which to perform a recovery must be accessible.
6. The ASR files must be accessible from a local floppy drive.
7. The floppy and CD drives cannot be external PC-card drives.
8. The original system must be an x86 or an Itanium machine.
9. The disks on the original system must have 512 byte sectors.
10. All basic volumes in the original system must be simple. Non-simple volumes are not supported. If a volume is mirrored, it must be a dynamic volume, not a basic volume.

### **Windows ASR software requirements**

1. Recovering a Windows XP SP1 system requires a Windows XP SP1 integrated (*slipstream*) installation CD. You cannot recover a Windows XP SP1 system using a base Windows XP operating system installation CD. SP1 ASR references the `asr_pfu.exe` file, which is not present on the base Windows XP installation. You can use the base level Windows XP installation CD to recover pre-SP1 Windows XP machines.
2. The Tivoli Storage Manager client installation package used on the TSMCLI CD you create for ASR recovery must be at the same or higher level than the level of client used to create the backups for the system being recovered. See “ASR preparation procedure” on page 116 for information about how to create the TSMCLI CD.
3. ASR requires a Tivoli Storage Manager Version 5.2.0 or higher client. Additionally, ASR for Windows Server 2003 requires Tivoli Storage Manager Version 5.2.0 or higher server.
4. Recovering a Windows XP 64-Bit Edition or Windows Server 2003 64-bit Edition system requires the hotfix or circumvention described in Microsoft Knowledgebase article 817708. This hotfix is integrated into service pack 1 (SP1) for Windows 2003.

### **ASR preparation overview**

To perform a successful ASR restore, you must prepare the following items:

1. Generate and back up the ASR files. See “Backing up Automated System Recovery (ASR) files (Windows XP, Windows Server 2003)” on page 82 for more information.
2. Create an Automatic System Recovery diskette. See “ASR preparation procedure” on page 116 for more information.
3. Perform a complete incremental backup of your system and boot drives using Tivoli Storage Manager.
4. Back up the system state on Windows Server 2003 or back up the System Objects on Windows XP. See “Backing up Windows system state” on page 80 or “Backing up system objects (Windows XP)” on page 79 for more information.
5. The operating system installation CD. For Windows XP, recovering a Windows XP SP1 system requires a Windows XP SP1 integrated (*slipstream*) installation CD. You cannot recover a Windows XP SP1 system using a base Windows XP operating system installation CD. SP1 ASR references the `asr_pfu.exe` file, which is not present on the base Windows XP installation. You can use the base level Windows XP installation CD to recover pre-SP1 Windows XP machines.
6. The Tivoli Storage Manager Windows client package in package-for-the-web format burned onto a CD. This is the single self-extracting installation package that IBM makes available with FTP.

7. A network connection supporting DHCP is required for ASR recovery from Tivoli Storage Manager server storage. Alternatively, if your network does not support DHCP, you can use local backup sets generated by Tivoli Storage Manager containing the backup of your system and boot drives; system state (Windows Server 2003) or system objects (Windows XP).
8. The Tivoli Storage Manager node password.
9. (Optional) One important facility provided by Microsoft is the Windows Recovery Console. You can use the Windows Recovery Console to facilitate easier problem determination in the event of recovery errors. You should consider using security policy administrative tools to enable the recovery console to access all paths and removable media. You must set these permissions prior to an ASR recovery situation.

## ASR preparation procedure

1. Create an ASR recovery diskette. You should create a recovery diskette any time there is a hardware configuration change or service pack update to your machine. Tivoli Storage Manager generates and backs up the required ASR files during system object (Windows XP) and system state (Windows 2003) backup. The ASR recovery diskette creation process will place the latest backup of these files on the diskette. You must have administrative authority to create an ASR diskette.

To create an ASR diskette using the GUI client:

- a. Generate and back up the ASR files. See “Backing up Automated System Recovery (ASR) files (Windows XP, Windows Server 2003)” on page 82 for more information.
- b. Select **Utilities** → **Create ASR Diskette** from the GUI main window. The Create Automated System Recovery Diskette window appears.
- c. Specify the floppy drive for creating the ASR diskette.
- d. Insert a blank formatted floppy disk in the designated drive.
- e. Select the designated drive letter in the Create Automated System Recovery Diskette window and click the **Finish** button. Tivoli Storage Manager copies the required ASR files to the diskette and labels the diskette volume TSMASR. If you forget to back up the ASR files prior to creating the diskette, the following error is shown:

No ASR filespace was found on server. Make sure you have backed up ASR prior to performing this operation.

You can also create the ASR recovery diskette for another node if you know the password for the other node. To create the diskette for another node:

- a. Set the *virtualnodename* option in your client options file (dsm.opt) to the node name of the workstation for which you are creating the ASR recovery diskette.
- b. Select **Utilities** → **Create ASR Diskette** from the GUI main window. The Create Automated System Recovery Diskette window appears.
- c. Specify the floppy drive for creating the ASR diskette.
- d. Insert a blank formatted floppy disk in the designated drive.
- e. Select the designated drive letter in the Create Automated System Recovery Diskette window and click the **Finish** button. Tivoli Storage Manager copies the required ASR files to the diskette and labels the diskette volume TSMASR.

### Notes:

- a. It is recommended that you always create an ASR diskette after you back up ASR files.

- b. Store the ASR recovery diskette in a safe location for future use.

To create an ASR diskette using the command line client:

- a. Back up the ASR files by issuing the following Tivoli Storage Manager command:

```
dsmc backup asr
```

The Windows Automated System Recovery window is displayed by the Windows API that creates the ASR files. See “Backup ASR” on page 436 for more information about the **backup asr** command.

- b. Insert a blank formatted diskette into the floppy drive. Assuming that your floppy drive is a:, issue the following command:

```
dsmc restore asr a:\
```

You can also create the ASR recovery diskette for another node if you know the password for the other node. To create the diskette for another node, use the *virtualnodename* option with the **restore asr** command to specify the node name of the workstation for which you are creating the ASR recovery diskette:

```
dsmc restore asr a:\ -virtualnodename=dagordon
```

See “Restore ASR” on page 528 for more information about the **restore asr** command.

- c. Use the Windows **label** command from a command line prompt to make the diskette label TSMASR. *The volume label for the diskette must be TSMASR for recovery to succeed.*

```
label a: tsmasr
```

**Note:** You must back up the ASR files using the **backup asr** command prior to performing an incremental backup of the system and boot drives. Backing up the ASR files generates the `ntdll.ASR` and `smss.ASR` files and places them in the `windows\repair` directory. These files must be present in the incremental backup of your system and boot drives in order for ASR recovery to succeed.

To verify that these files were backed up, use the **query backup** command. For example, assuming that your Windows system is installed on drive c:, enter the following command:

```
query backup c:\windows\repair\*.asr
```

2. Perform a complete incremental backup of your system and boot drives. The system drive is the partition containing the boot files such as `ntldr` and `boot.ini`. The boot drive is the partition containing your Windows system directory. Assuming your system and boot files are on drive c:, enter the following command:

```
dsmc incremental c:
```

3. Back up the system state or system objects. To back up the system state, enter the following command:

```
dsmc backup systemstate
```

To back up system objects, enter the following command:

```
dsmc backup systemobject
```

To verify that you backed up the system state, enter the following command:

```
dsmc query systemstate
```

**Note:** You can specify `-showmembers=yes` with these commands to display file level detail.

To verify that you backed up the system objects, enter the following command:

```
dsmc query backup systemobject
```

4. Locate your operating system installation CD. ASR boots the machine from the operating system CD inserted in the locally attached CD drive during recovery. Network booting of the operating system is not supported for ASR recovery.  
**Note:** For Windows XP, you must use installation media that has integrated (*slipstream*) service pack 1 (SP1). Integrated installation media for XP SP1 is currently available from Microsoft in the English and German languages. See “Reference information” on page 633 for instructions to create your own slipstream CD.
5. Prepare the Tivoli Storage Manager client package in package-for-the-web format burned onto a CD. This is the single self-extracting installation package of Tivoli Storage Manager that IBM makes available with FTP. The volume label of the CD must be TSMCLI and the client package must be renamed to TSMCLI.EXE and placed in the root directory of the CD.
  - a. Download the Tivoli Storage Manager client package.
  - b. Rename the package to TSMCLI.EXE. From a command prompt, enter the following command:

```
RENAM IPxxxxx.EXE TSMCLI.EXE
```
  - c. Using the method described by your CD burning software, specify TSMCLI as the volume label.
  - d. Ensure that the TSMCLI.EXE file path is designated as the root directory ( \ ) in your CD burning software.
  - e. Burn the CD using the method described by your CD burning software.
6. A network connection supporting DHCP is required for ASR recovery from Tivoli Storage Manager server storage. If your network does not support DHCP, you can use local backup sets containing the backup of your system and boot drives and system state or system objects.  
**Note:** For best performance, you should generate the system state backup in a backup set separate from the boot and system drives.
7. You must know your Tivoli Storage Manager node password. Your Tivoli Storage Manager node password is requested during ASR recovery, even if you set the *passwordaccess* option to *generate*. During ASR recovery, the generated Tivoli Storage Manager password is not available in the Windows Registry of the temporary operating system image installed by Windows. If you do not know your Tivoli Storage Manager node password, ask your Tivoli Storage Manager administrator to update the node password to a known value for you. After ASR recovery is completed, you can use the **set password** command to change the password to another value.
8. (Optional) Prior to an ASR recovery situation, use security policy administrative tools to enable the Windows Recovery Console to access all paths and removable media. This will facilitate problem determination in the event of recovery errors.

## ASR restore procedure

During a system recovery, perform the following steps:

1. Insert the Windows XP or Windows Server 2003 operating system installation CD into the CD drive.
2. Restart your computer. Press the appropriate key to boot from CD. You might need to configure your computer BIOS to enable boot from CD.
3. Press **F6** if you need to install any third party SCSI or RAID drivers.
4. Press **F2** to enter the ASR recovery mode during text mode section of Setup.

5. Insert the ASR recovery diskette (TSMASR) into the floppy drive when prompted for the Windows ASR recovery diskette. The following tasks are performed automatically:
  - Windows reads the `asrpnpsif` file from the diskette.
  - Windows reformats the boot volume and might reformat other partitions on the disk if the partition layout is different than the original system. This partitioning and format process can cause several reboots.
6. Insert the Tivoli Storage Manager installation CD (TSMCLI) into the CD drive when prompted. The Tivoli Storage Manager client package (TSMCLI.EXE) is copied to the `%SystemRoot%\temp` directory.
7. Insert the ASR recovery diskette (TSMASR) again when prompted. The `tmasr.opt`, `tmasr.cmd`, and `waitforevent.exe` files are copied to the `%SystemRoot%\temp` directory.
8. Remove the diskette when prompted, prior to system reboot. The system reboots into GUI-mode setup.
9. Insert the operating system CD into the drive when prompted. The prompt message is unclear; it reads  
 Setup was unable [to] access files needed to continue. This could be caused by an error on the CD Media or the CD is no longer present in the drive.
10. After setup is completed, a command window is opened, and the `tmasr.cmd` file is run in this window.
11. The Tivoli Storage Manager client is silently installed.
12. You are prompted to choose whether you want to restore from a network connected Tivoli Storage Manager server (press 1), a local backup set from file (press 2), or a local backup set from tape (press 3).
13. Tivoli Storage Manager prompts for your ID and password if you are restoring from the Tivoli Storage Manager server. If you are restoring from the local backup set, you are prompted for the path to the local backup set, for example `r:\backupsetfile1.ost`. You will be prompted for subsequent volumes if your backup set occupies more than one volume.
14. Tivoli Storage Manager commands are issued to restore the system drive.
15. Depending on your operating system, issue the **restore systemobject** or **restore systemstate** command.
16. The command file exits and the machine reboots. Remove the ASR recovery diskette (TSMASR) from the drive. The operating system comes up in fully recovered state.

**Notes:**

1. After your system restarts, you can use Tivoli Storage Manager to restore your data files from the server. Note that during recovery ASR brings your operating system back to the point of the last successful backup. When ASR is complete, you will need to recover any applications and data that were not resident on the operating system critical volumes.
2. Store the ASR recovery diskette in a safe location for future use.
3. See Appendix E, "ASR supplemental information," on page 625 for information for ASR questions and answers, diagnostic tips, and an in-depth discussion of ASR.

---

## Restoring Microsoft Dfs trees and files

To restore Dfs junctions and the data for each junction, restore the Dfs junction metadata first and then restore each junction separately. If the junction metadata is not restored, Tivoli Storage Manager creates a directory under the Dfs root using the same name as that of the junction point and restores the data in that directory. See “Backing up Microsoft Dfs files” on page 101 for more information.

---

## Restoring an image

Before you perform an image restore, consider the following:

- To perform an image restore (offline or online) you must have administrative authority on the system.
- 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.
- The image restore operation will overwrite the volume label on the destination volume with the one that existed on the source volume.
- 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 does not have to be the same type as the original. The volume does not even have to be formatted. The image restore process will create the appropriately formatted file system for you.
- 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 created an image of the system drive, you cannot restore the image to the same location because the client cannot have an exclusive lock of the system drive. Also, because of different system component configurations, the system image might not be consistent across components (such as Active Directory). Some of these components can be configured to use different volumes where parts are installed on the system drive and others to non-system volumes.
- 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 should run **chkdsk** to check for and repair any bad sectors (unless the restored volume is RAW).

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 408 for more information.

If bad sectors 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 276 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 drive letter or 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. Because image backup allows you to back up only used blocks in a file system, the stored image size on the Tivoli Storage Manager server could be smaller than the volume size. For online image backups, the stored image can be larger than the file system based on the size of the cache files.
  - 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 536 for more information.

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 408 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 276 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, DVD, and Iomega JAZ or ZIP drives. Current device support information is available at the following Web site:

- [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)

You can restore backupsets from the following locations:

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

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

### Instant archive

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

### Rapid recovery

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



**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 Windows standalone workstation, without requiring a server connection, use the *localbackupset* option. See “Localbackupset” on page 299 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 node name or viewing the backup set in its raw format. Encryption or physical protection of the media are the only methods to ensure that the data is protected.

If you restore backupset data from the server, individual files, directories or entire backupset data can be restored in a single operation from the GUI or the command line. When you restore backupset data locally, the GUI can only display and restore an entire backupset. The command line can be used to restore individual files or directories stored in a backupset locally. Refer to the *backupsetname* option on the **restore** and **query** commands for more information.

## Restoring backupsets using the 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.

You can select files from within a backupset if that backupset is located at the server and it was generated with a table of contents.
- Click **Restore**. The Restore Destination window appears. Enter the appropriate information.
- Click **Restore**. The Task List window displays the restore processing status.

**Notes:**

1. If the object you want to restore is part of a backupset generated on a node, and the node name is changed on the server, any backupset objects that were generated prior to the name change will not match the new node name. Ensure that the node name is the same as the node for which the backupset was generated.
2. The client can be used to restore a backupset on an attached device with or without a server connection. If the server connection fails, a prompt will appear to continue for purposes of local backupset restore. Also, the `localbackupset` option can be used to tell the client not to attempt the connection to the server.
3. Certain local devices 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. Restoring individual system objects or system state components.
  - c. 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.
  - d. 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. To restore individual files, folders, and images from a backupset, use the `restore` command with the `backupsetname` option instead. Refer to the `backupsetname` option on the `restore backupset`, `restore`, and `query` commands for more information.

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

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

---

## 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 system object, 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 442.

Perform incremental backups to support a point-in-time restore. During an incremental backup, the client notifies the server when files are deleted from a client file space or directory. Selective and incremental-by-date backups do not notify the server about deleted files. Run incremental backups at a frequency consistent with possible restore requirements.

If you request a point-in-time restore with a date and time that is prior to the oldest version maintained by the Tivoli Storage Manager server, the object is not restored to your system. Files which were deleted from your workstation prior to the point-in-time specified will not be restored.

### Notes:

1. Your administrator must define copy group settings that maintain enough inactive versions of a file to guarantee that you can restore that file to a specific date and time. If enough versions are not maintained, Tivoli Storage Manager might not be able to restore all objects to the point-in-time you specify.
2. If you delete a file or directory, the next time you run an incremental backup, the active backup version becomes inactive and the oldest versions that exceed the number specified by the *versions data deleted* attribute of the management class are deleted. See Chapter 8, “Understanding storage management policies,” on page 153 for more information about the *versions data deleted* attribute.

When performing a point-in-time restore, consider the following:

- Tivoli Storage Manager restores file versions from the most recent backup before the specified point-in-time date. Ensure the point-in-time that you specify is not the same as the date and time this backup was performed.
- If the date and time you specify for the object you are trying to restore is earlier than the oldest version that exists on the server, Tivoli Storage Manager cannot restore that object.
- Point-in-time restore will restore files deleted from the client workstation after the point-in-time date but not files deleted before this date.

- Tivoli Storage Manager cannot restore a file created after the point-in-time date and time. When a point-in-time restore runs, files that were created on the client after the point-in-time date are not deleted.

To perform a point-in-time restore using the client GUI, use the following steps:

1. Click the **Restore** button in the main window. The Restore window appears.
2. Click the **Point-in-Time** button from the Restore window. The Point in Time Restore window appears.
3. Select the **Use a Point-in-Time Date** selection box. Select the date and time and click **OK**. The point in time that you specified appears in the Point in Time display field in the Restore window.
4. Display the objects you want to restore. You can search for an object by name, filter the directory tree, or work with the directories in the directory tree.
5. Click the selection boxes next to the objects you want to restore.
6. Click the **Restore** button. The Restore Destination window is displayed. Enter the appropriate information.
7. Click the **Restore** button to start the restore. The Restore Task List window displays the restore processing status.

**Note:** If there are no backup versions of a directory for the point-in-time you specify, files within that directory are not restorable from the GUI. However, you can restore these files from the command line.

You can start point-in-time restore from the command-line client using the *pitdate* and *pittime* options with the **query backup** and **restore** commands. For example, when you use the *pitdate* and *pittime* options with the **query backup** command, you establish the point-in-time for which file information is returned. When you use *pitdate* and *pittime* with the **restore** command, the date and time values you specify establish the point-in-time for which files are returned. If you specify *pitdate* without a *pittime* value, *pittime* defaults to 23:59:59. If you specify *pittime* without a *pitdate* value, it is ignored.

---

## Restoring 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 20.

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 20. To restore NAS file systems using the Web client GUI:

1. Click the **Restore** button from the main window. The Restore window appears.
2. Expand the directory tree if necessary. To expand a node in the tree, click the plus sign (+) next to an object in the tree.

**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 files and directories using the Web client

You can use the *toc* option with the *include.fs.nas* option in your client options file to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup. See "Toc" on page 399 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. 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.

To restore NAS files and directories:

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 appear 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 display the **File Level** object.
  4. Expand the **File Level** object to display the volumes, directories, and files that were last backed up. When you expand the volume object, and complete TOC information is available on the server for the latest backup, the Load Table of Contents dialog appears. If complete TOC information is not available for the latest backup, no objects will appear below the volume object. The next step explains how to display objects from backups other than the latest backup. Complete TOC information is provided if you performed either of the following operations:
    - A differential image backup with TOC information and its corresponding full image backup with TOC information
    - A full image backup with TOC information
  5. Click the selection boxes next to the directories or files that you want to restore.
    - a. If you want to restore files from a NAS image that was backed up on a particular date or display files from several older versions, highlight the volume you want to restore and click the **Point In Time** button.
    - b. If you select **Use a Point in Time Date** in the Point in Time Restore windows, files from the image backed up on that date, and if it is a differential image, files from its corresponding full image appear under the **File Level** object.
    - c. If you click **Use Selected Images** in the Point in Time Restore window, the Selected Images window appears for you to select images. The contents of the selected images appear in the **File Level** object.
  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.
  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, remote (NAS), and Websphere Application Server 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 (either workstation, NAS, or Websphere Application Server).
- 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**. You can delete both workstation and remote objects.

## Restoring NAS file systems using the command line

Table 27 lists the commands and options you can use to restore NAS file system images from the command line.

Table 27. 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.	507
<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.	491
<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.	499
<b>restore nas</b>	Restores the image of a file system belonging to a Network Attached Storage (NAS) file server.	540
<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.	485
<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.	460
<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.	468

A NAS file system specification uses the following conventions:

- 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}.

**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 NetApp CIFS shares

For information about backing up NetApp CIFS shares, see “Backing up NetApp CIFS share definitions” on page 95.

Restoring the share definition requires restoring the root directory of the share file space, which under most circumstances can be done as follows:

```
dsmc rest \\NetAppFiler\CifsShareName\ -dirsonly
```

The following command line client output indicates that the root directory (and share definition has been restored):

```
Restoring          0 \\NetAppFiler\CifsShareName\ [Done]
```

If the CIFS share definition is deleted on the NetApp filer, the client is unable to directly restore the share definition because the share is no longer accessible.

The share definition can be restored indirectly by creating a temporary local share and restoring the share definition to the temporary share as follows:

```
md c:\tempdir net share tempshare=c:\tempdir
/remark:"Temporary Share for Restoring Deleted CIFS Share"
net use z: \\LocalMachineName\tempshare
dsmc res \\NetAppFiler\CifsShareName\ z:\ -dirsonly
```

This will restore the original share definition (including permissions) on the filer.

Older versions of the Tivoli Storage Manager server might have a problem which prevents restoring the root directory and the CIFS share definition. If this problem occurs, it can be circumvented by using by one of the following methods:

1. Use the DISABLENQR testflag to restore the root directory as follows:

```
dsmc res \\NetAppFiler\CifsShareName\ -test=disablenqr -dirsonly
```
2. Use the command line client `-pick` option with a restore command and select the root directory:

```
dsmc res \\NetAppFiler\CifsShareName\ -dirsonly -pick
```

---

## Restoring VMware Consolidated Backup data

VMware Consolidated Backup restores must be done from within the virtual machine, using the normal command line client and the GUI client.

**Note:** You can perform restores from any machine, including the backup proxy machine, by using the *virtualnodename* option. Refer to “Virtualnodename” on page 410 for information on using this option.

Refer to “Using VMware Consolidated Backup” on page 92 for information on preparing your system, installing, configuring, and testing your VMware Consolidated Backup setup.



---

## Restoring the WebSphere Application Server

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 Websphere Application Server, 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:

- **Recommendation:** 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 Websphere Application Server 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 Websphere Application Server was backed up, the WebSphere Application Server tree displays only the WAS\_<NODENAME> and WAS\_ND\_<NODENAME> file spaces.
  - c. If multiple instances of Websphere Application Server 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 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 inactive objects.
  - b. If you want to see all active and inactive Websphere Application Server backups before you select them, select **View** → **Display active/inactive files** from the menu bar. Inactive Websphere Application Server 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 Websphere Application Server 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 Websphere Application Server object (from the Restore window), select a Websphere Application Server object, click **View** → **File Details**.
- If you need to recover a Websphere Application Server 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 Websphere Application Server file space. See “Restore WAS” on page 549 for more information.
- To delete Websphere Application Server file spaces, select **Utilities** → **Delete Filespaces** from the Restore window.

To restore Websphere Application Server objects from the Tivoli Storage Manager command line, use the **restore was** command. See “Restore WAS” on page 549 for more information.

---

## Restore: Additional considerations

This section discusses additional 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 a user on another node to restore your backup versions or retrieve your archive copies. This lets you share files with other people or with other workstations that you use with a different node name.

You can also authorize other nodes to access the ASR file space. In the event of a workstation disaster where the workstation’s ASR diskette is not available, another node can be used to create the ASR diskette so that the workstation can be recovered using ASR and Tivoli Storage Manager. See “ASR preparation procedure” on page 116 for more information.

To authorize another node 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. Type the node name of the user in the Grant Access to Node field. Type the node name of the user’s host machine in the Grant Access to Node field.
5. Type the name of a user on the host machine in the User field.
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. 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.

For example, to give the node user2 access to all backup files and subdirectories under the d:\user1 directory, create a rule with the following values:

```
Permit Access to: Backed up Objects
Grant Access to Node: user2
Filespace and Directory: d:\user1
Filename: *
Include subdirectories: Selected
```

The node you are authorizing must be registered with your Tivoli Storage Manager server.

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 the following sections:

- “Set Access” on page 563
- “Query Access” on page 487
- “Delete Access” on page 462

## 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, by following these steps:

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.
4. Click the **Set** button.

If you are using commands, use the *fromnode* option to indicate the node. You must also use the file space name, rather than the drive letter, to select the restore-retrieve drive you want to access. Include the file space name in braces and use it like a drive letter. For example, to restore cougar’s files from the \projx directory on the d-disk file space to your own \projx directory, enter:

```
dsmc restore -fromnode=cougar \\cougar\d$\projx\* d:\projx\
```

Use the **query filesystem** command to display a list of file spaces. For example, to display a list of cougar’s file spaces, enter:

```
dsmc query filesystem -fromnode=cougar
```

For more information about using the *fromnode* option with the **restore** command, see “Restore” on page 522. See “Retrieve” on page 553 for the **retrieve** command. For more information about the *fromnode* option, see “Fromnode” on page 268.

## Restoring or retrieving your files to another workstation

When you are using a different workstation, you can restore or retrieve files you backed up from your own workstation. Your backup versions and archive copies are stored according to your node, not your specific workstation. Your Tivoli Storage Manager password protects your data.

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. You can use the *virtualnodename* option when starting Tivoli Storage Manager or place the option in your client options file, *dsm.opt*, on the workstation. If you are using a workstation other than your own, use the *virtualnodename* option with the **dsm** command. For example, if your node name is cougar, enter:

```
start dsm -virtualnodename=cougar
```

You can then restore or retrieve files as if you were working on your original workstation.

You can use *virtualnodename* on commands also. For example, if you want to restore your *\projx* files to your local *c:\myfiles* directory, enter:

```
dsmc restore -virtualnodename=cougar \\cougar\d$\projx\*. * c:\myfiles\
```

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

## Restoring or retrieving files to another type of workstation

Because the file formats used on Windows are compatible, you can restore or retrieve files from one system type to another. This is called *cross-client restore*.

**Attention:** You must have the appropriate permissions to access the file spaces of the other workstation, as explained in “Authorizing another user to restore or retrieve your files” on page 132 and “Restoring or retrieving files from another client node” on page 133.

NTFS drives permit file and directory names that are longer than those permitted on FAT drives. If you are recovering files to a FAT drive with long file names, you should always specify a destination file specification for each file. This ensures that you get the name you want.

When you use the Windows client to recover files with long names to an NTFS disk, the long names are preserved, even if you are recovering the file to a different type of drive than the source drive.

The considerations for retrieving files are the same as for restoring them.

## Restoring a disk in the event of disk loss

The Tivoli Storage Manager client must be installed in order to recover files. If the disk that contains the client is lost, you must reinstall the client before you can recover your files.

If your administrator has created a backup set of your files on portable media, you can use that set to recover the files without the use of a server. See “Restoring data from a backupset” on page 122 for information on backup sets.

If you lose the disk containing the operating system, you must recover the operating system before you can connect to the server.

For Windows XP and Windows Server 2003: If you generated and backed up Automated System Recovery (ASR) files, you can automatically save the ASR files to diskettes. You can use the ASR diskettes in case you need to recover your system due to a catastrophic system or hardware failure. See “Restoring Automated System Recovery (ASR) files (Windows XP, Windows Server 2003)” on page 113 for more information.

- An operating system that permits you to perform basic functions. You might also want to be able to boot DOS on your computer in the event that it does not start.

Windows provides a *safe mode* start-up that allows you to boot to a command line prompt.

- A client with an appropriately customized options file. The command line client is sufficient for this.

Consult your operating system and communication software manuals for help in setting up these diskettes.

If you have large volumes of data that might need to be recovered, your administrator might want to create a backup set of those files on portable media. You can then use that set to recover the files without using a server. See “Restoring data from a backupset” on page 122 for information on backup sets.

## Deleting file spaces

If your Tivoli Storage Manager administrator grants you authority, you can delete entire file spaces from the server. You cannot delete individual backup versions that are kept on the server. When you delete a file space, you delete all the files, both backup versions and archive copies, that are contained within the file space. For example, if you delete the file space for your c: drive, you are deleting every backup for every file on that disk and every file you archived from that disk.

**Carefully consider what you are doing before you attempt to delete a file space.**

You can delete file spaces using the 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 client, perform the following steps:

1. Select **Utilities** → **Delete Filespaces** from the main window.
2. Selection 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 468 for more information. Use the *class* option with the **delete filesystem** command to delete NAS file spaces. See “Class” on page 208 for more information.



---

## Chapter 6. Archiving and retrieving your data

Archiving and retrieving files is similar to backing up and restoring files. This chapter discusses the primary archive and retrieve tasks. See “When to back up and when to archive files” on page 58 for a discussion of the difference between backups and archives.

An **Estimate** function is provided on the Archive and Retrieve windows, which are accessible from the main Tivoli Storage Manager window. Use this function if you want to estimate the amount of time it takes to process your files and directories. The estimated transfer is a rough calculation of the time it takes Tivoli Storage Manager to transfer your data. It is based on previous transfers of data between your workstation and the current server. The actual transfer time could be longer or shorter than the estimate due to factors such as network traffic, system load on your workstation, or system load on the server.

Unless otherwise specified, references to Windows refer to all supported Windows operating systems.

All client archive and retrieve procedures in this chapter also apply to the Web client, except the following:

- Estimate
- Preferences editor
- Setup wizard

See “Starting a Web client session” on page 51 for information on starting the Web client.

Table 28 identifies tasks described in this chapter:

*Table 28. Archive and retrieve: Primary tasks*

Task	Page
Archiving data using the GUI	138
Archiving data using the command line	139
Deleting archive data	142
Retrieving archives using the GUI	143
Retrieving archives using the command line	143

---

### Archiving files

To archive files, select the files you want to archive. You can select the files by name or description, or select them from a directory tree. Your administrator might have set up schedules to automatically archive certain files on your workstation. See Chapter 7, “Automating tasks,” on page 145 for information about checking and running the schedules available to you. The following sections discuss how to archive files without using a schedule.

Tivoli Storage Manager requires that you assign an archive description for all archived files. An archive description identifies data through a meaningful description that you can use later to identify files and directories. You can enter as

many as 254 characters to describe your archived data. If you do not enter a description, Tivoli Storage Manager assigns the following default archive description:

Archive Date: mm/dd/yyyy

where mm/dd/yyyy is the current date.

When you select the archive function from the backup-archive GUI, Tivoli Storage Manager displays a list of all previously used archive descriptions. You can assign these archive descriptions to future archives.

Incremental backup might recall migrated files, while selective backup and archive will always recall migrated files, if you do not use the *skipmigrated* option. Refer to *IBM Tivoli Storage Manager for HSM for Windows Administration Guide* for details about using the *skipmigrated* option.

## Archiving files using open file support

If open file support has been configured (see “Configuring Open File Support (OFS)” on page 29), Tivoli Storage Manager performs a snapshot backup or archive of files that are locked (or “in use”) by other applications. The snapshot allows the archive to be taken from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to the file system are not included in the archive. You can set the *snapshotproviderfs* parameter of the *include.fs* option to *none* to specify which drives do not use open file support.

To control an open file support operation with LVSA, you can specify these additional options in your dsm.opt file or as values of the *include.fs* option: *snapshotcachelocation*, *snapshotcachesize*, *snapshotfsidleretries*, *snapshotfsidlewait*, *snapshotproviderfs*, *presnapshotcmd*, *postsnapshotcmd*. See Chapter 9, “Using processing options,” on page 165 for information about using these options.

### Notes:

1. You can use the *include.fs* option to set snapshot options on a per file system basis.
2. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with FAT, FAT32 or NTFS file systems. This support includes SAN-attached volumes that meet these requirements.
3. If the client is unable to create a snapshot, failover to non-OFS backup occurs; the same backup support that would be done if the OFS feature was not installed.
4. To enable open file support in a cluster environment all machines in the cluster should have the OFS feature configured.

For information about Tivoli Storage Manager Open File Support restrictions and issues, search for the **TSM Client v5.2 Open File Support** document under the **Storage Management** product category at the following Web site:

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

## Archiving data using the GUI

You can archive specific files or entire directories from a directory tree. You can also assign a unique description for each group of files you archive (archive package).



To archive your files:

1. Click the **Archive** button in the GUI main window. The Archive window appears.
2. Expand the directory tree by clicking the plus sign (+) or the folder icon next to an object in the tree. To search or filter files, click the **Search** icon from the tool bar.
3. Enter a description, accept the default description, or select an existing description for your archive package in the **Description** pull down box.
4. To modify specific archive options, click the **Options** button. Any options you change are effective during the current session *only*. To estimate the transfer time for your archive selections click the **Estimate** button.
5. Click the **Archive** button. The Archive Status window displays the progress of the archive.

## Archiving data using the command line

You request archive services when you want to preserve copies of files in their current state, either for later use or for historical or legal purposes. You can archive a single file, a group of files, or all the files in a directory or subdirectory. After you archive a file, you can choose to delete the original file from your workstation. Use the **archive** command to archive files. See “Archive” on page 433 for more information about the **archive** command.

### Open file support for archive operations

If open file support has been configured (see “Configuring Open File Support (OFS)” on page 29), Tivoli Storage Manager performs a snapshot archive of files that are locked (or “in use”) by other applications. Severe open file support errors are reported in the Windows event log. See “Archiving files using open file support” on page 138 for more information.

### 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 375 for more information.

Table 29 shows examples of using the **archive** command to archive objects. See “Archive” on page 433 for additional examples.

Table 29. Command line archive examples

Task	Command	Considerations
Archive all files in the c:\plan\proj1 directory with a file extension of .txt.	<code>dsmc archive c:\plan\proj1\*.txt</code>	Use wildcards to archive more than one file at a time.
Archive all files in the c:\small\testdir directory and delete the files on your workstation.	<code>dsmc archive c:\small\testdir\* -deletefiles</code>	Retrieve the archived files to your workstation whenever you need them again. See “Deletefiles” on page 229 for more information about the <i>deletefiles</i> option.

Table 29. Command line archive examples (continued)

Task	Command	Considerations
Archive the c:\proj1\h1.doc file and the c:\proj2\h2.doc file	<code>dsmc archive c:\proj1\h1.doc c:\proj2\h2.doc</code>	You can specify as many file specifications as available resources or other operating system limits permit. Separate file specifications with a space. You can also use the <i>filelist</i> 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 262 for more information.
Archive a list of files in the c:\filelist.txt file.	<code>dsmc archive -filelist=c:\filelist.txt</code>	Use the <i>filelist</i> option to process a list of files. See “Filelist” on page 262 for more information.
Archive the a:\ch1.doc file and assign a description to the archive.	<code>dsmc archive a:\ch1.doc -description="Chapter 1, first version"</code>	If you do not specify a description with the <b>archive</b> command, the default is Archive Date:x, where x is the current system date. See “Description” on page 230 for more information about the <i>description</i> option.
Archive all of the files in the d:\proj directory and its subdirectories.	<code>dsmc archive d:\proj\ -subdir=yes</code>	See “Subdir” on page 381 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 c:\relx\dir1 directory.	<code>dsmc archive c:\relx\dir1\ -v2archive</code>	Tivoli Storage Manager archives only files in the c:\relx\dir1 directory. Directories that exist in the path are not processed. See “V2archive” on page 406 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 c:\plan\proj1\ budget.jan\*</code>	See “Archmc” on page 193 for more information about the <i>archmc</i> option. See Chapter 8, “Understanding storage management policies,” on page 153 for more information about management classes.
Assuming that you initiated a snapshot of the C: drive and mounted the snapshot as the logical volume \\florence\c\$\snapshots\ snapshot.0, archive the c:\dir1\sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name C:	<code>dsmc archive c:\dir1\sub1\* -subdir=yes -snapshotroot=\\florence\c\$\snapshots\snapshot.0</code>	See “Snapshotroot” on page 375 for more information.

## Archiving data with client node proxy

Archives of multiple nodes that share storage can be consolidated to a common target node name on the Tivoli Storage Manager server. This is useful when the 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 that 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 node name 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 node name that is used for authentication purposes. Data will not be stored using the node name 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 or system state 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 194 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 463 for more information.

---

## Retrieving archives

Select the **retrieve** function to recover an archive copy of a file or a directory.

**Note:** When retrieving a directory, its modification date and time is set to the date and time of the retrieve, not to the date and time the directory had when it was archived. This is because Tivoli Storage Manager retrieves the directories first, then adds the files to the directories.

You can also retrieve archive copies from the directory tree, filter the directory tree, and retrieve archive copies of files owned by someone else. To do any of these, click the **Retrieve** button on the Tivoli Storage Manager main window and follow the directions provided in the task help of the GUI.

For information about how to resolve conflicts between duplicate file names, see “Duplicate file names” on page 105.

**Important:** When you retrieve a file without any specifications, and more than one version of the archive copy exists on the server, Tivoli Storage Manager retrieves all of the copies. After the first copy is retrieved, the second copy will be retrieved. If there is an existing copy on your client machine, you are prompted to replace, skip, or cancel.

## Retrieving archives using the GUI

To retrieve archives using the Tivoli Storage Manager GUI, perform the following steps:

1. Click the **Retrieve** button on the 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 located to the right of the **Estimate** button. Any options you change are effective during the current session *only*. To estimate the transfer time for your archived selections click the **Estimate** button.
  4. Click **Retrieve**. The Retrieve Destination window appears. You can retrieve files to a directory or drive other than the one from where they were originally archived. You can also select how much of the parent directory structure will be recreated at the retrieve location.
  5. Click **Retrieve**. The Retrieve Status window displays the processing status.

## Retrieving archives 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. See “Retrieve” on page 553 for more information about the **retrieve** command. Table 30 shows examples of using the **retrieve** command. See “Retrieve” on page 553 for additional examples.

Table 30. Command line examples of retrieving archives

Task	Command	Considerations
Retrieve the c:\doc\h2.doc file to its original directory.	dsmc retrieve c:\doc\h2.doc	If you do not specify a destination, the files are retrieved to their original location.
Retrieve the c:\doc\h2.doc file under a new name and directory.	dsmc retrieve c:\doc\h2.doc c:\proj2\h3.doc	None

Table 30. Command line examples of retrieving archives (continued)

Task	Command	Considerations
Retrieve all files archived with a specific description to a directory named retr1 at a new location	<code>dsmc retrieve c:\* d:\retr1\ -sub=yes -desc="My first archive"</code>	None
Retrieve all files from the c:\projecta directory that end with the characters .bak to the c:\projectn directory.	<code>dsmc retrieve c:\projecta\*.bak c:\projectn</code>	None
Use the <i>pick</i> option display a list of archives from which you can select files to retrieve.	<code>dsmc retrieve c:\project\* -pick</code>	See "Pick" on page 322 for more information about the <i>pick</i> option.
Retrieve a file originally archived from the diskette, <b>workathome</b> on the a: drive, to a diskette in the a: drive labeled <b>extra</b> .	<code>dsmc retrieve {workathome}\doc\ h2.doc a:\doc\h2.doc</code>	If you are retrieving a file to a disk that has a different label other than the disk from which the file was archived, use the file space name (label) of the archive disk rather than the drive letter.
Retrieve the c:\doc\h2.doc file to its original directory on the workstation, named <i>star</i> .	<code>dsmc retrieve c:\doc\h2.doc \\star\c\$\</code>  To retrieve the file to <i>star</i> which has been renamed <i>meteor</i> , enter: <code>dsmc retrieve \\star\c\$\ doc\h2.doc \\meteor\c\$\</code>  You could also enter: <code>dsmc retrieve \\star\c\$\ doc\h2.doc c:\</code>  This example is valid because if the workstation name is not included in the specification, the local workstation is assumed ( <i>meteor</i> , in this case).	For the purposes of this manual, the workstation name is part of the file name. Therefore, if you archive files on one workstation and you want to retrieve them to another workstation, you must specify a destination. This is true even if you are retrieving to the same physical workstation, but the workstation has a new name.

---

## Chapter 7. Automating tasks

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

With client scheduling, you can also:

- Display information about available schedules.
- Display information about work that the schedule has completed.
- Modify scheduling options in the client options file (dsm.opt). See “Scheduling options” on page 179 for more information.
- The scheduler can be installed as a service and can be set to start automatically when the system is rebooted. The scheduler can be setup to run as a service using the Setup wizard.
- When setup to run as a service, the scheduler can be configured to run in normal mode, or managed by the Tivoli Storage Manager client acceptor daemon. The client acceptor daemon will automatically stop and start the schedule service when necessary to save system resources when it is not in use. You can configure this using the Setup wizard.

### 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 options file (dsm.opt) 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 179.

---

## 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 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: e: f:
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 on the e: and f: drives.

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 147 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 as a service, output from scheduled commands displays in the application event log. Output is also directed to the `dsmsched.log` file in the current directory unless you change the path and file name using the *schelogname* option. The amount of detail is determined by whether *verbose* or *quiet* is set in the `dsm.opt` file. The scheduler service also posts messages to the Windows event log.

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:

## Event logging

The scheduler service also logs information into the application event log and provides an event identification (event ID) number for each event in the log. The following section shows examples of events that are logged to the application event log.

### Event log examples

#### Scheduler Service

##### 4097: Informational message

###### Example 1:

Event Type: Information  
Event Source: AdsmClientService  
Event Category: None  
Event ID: 4097  
Date: 10/31/2002  
Time: 8:29:57 AM  
User: DILE\Administrator  
Computer: MIKEDILE  
Description:  
TSM 515 Scheduler halted.

###### Example 2:

Event Type: Information  
Event Source: AdsmClientService  
Event Category: None  
Event ID: 4097  
Date: 10/31/2002  
Time: 8:29:57 AM  
User: DILE\Administrator  
Computer: MIKEDILE  
Description:  
Scheduler Terminated, service ending.

###### Example 3:

Event Type: Information  
Event Source: AdsmClientService  
Event Category: None  
Event ID: 4097  
Date: 10/31/2002  
Time: 8:29:56 AM  
User: DILE\Administrator  
Computer: MIKEDILE  
Description:  
TSM Client Scheduler 'TSM 515 Scheduler'  
Started.

###### Example 4:

Event Type: Information  
Event Source: AdsmClientService  
Event Category: None  
Event ID: 4097  
Date: 10/31/2002  
Time: 8:29:56 AM  
User: DILE\Administrator  
Computer: MIKEDILE  
Description:  
Starting Scheduler.

###### Example 5:

Event Type: Information  
Event Source: AdsmClientService  
Event Category: None  
Event ID: 4097  
Date: 10/30/2002  
Time: 8:06:09 PM  
User: DILE\Administrator  
Computer: MIKEDILE  
Description:  
Incremental backup of volume '\\MIKEDILE\C\$'

##### 4098: Warning message

###### Example 1:

Event Type: Warning  
Event Source: AdsmClientService  
Event Category: None  
Event ID: 4098  
Date: 10/31/2002  
Time: 8:29:56 AM  
User: DILE\Administrator  
Computer: MIKEDILE  
Description:  
Error Initializing TSM Api, unable to verify  
Registry Password, see dsierror.log.

###### Example 2:

Event Type: Warning  
Event Source: AdsmClientService  
Event Category: None  
Event ID: 4098  
Date: 9/20/2002  
Time: 6:20:10 PM  
User: DILE\Administrator  
Computer: MIKEDILE  
Description:  
ANS1802E Incremental backup of '\\mikedile\  
c\$' finished with 3 failure

##### 4099: Error message

###### Example 1:

Event Type: Error  
Event Source: AdsmClientService  
Event Category: None  
Event ID: 4099  
Date: 9/17/2002  
Time: 6:53:13 PM  
User: DILE\Administrator  
Computer: MIKEDILE  
Description:  
Scheduler exited with a result code of 4.

## Example 2:

Event Type: Error  
Event Source: AdsmClientService  
Event Category: None  
Event ID: 4099  
Date: 9/17/2002  
Time: 6:27:19 PM  
User: DILE\Administrator  
Computer: MIKEDILE  
Description:  
ANS4987E Error processing '\\mikedile\e\$\  
tsm520c\client\winnt\mak \dsmwin32.ncb':  
the object is in use by another process

### 4100: Scheduler command message

Event Type: Information  
Event Source: AdsmClientService  
Event Category: None  
Event ID: 4100  
Date: 10/31/2002  
Time: 8:29:56 AM  
User: DILE\Administrator  
Computer: MIKEDILE  
Description:  
Next Scheduled Event Obtained from Server  
SNJEDS1 (MVS):  
-----  
Schedule Name: NIGHTLY\_BACKUP  
Action: Incremental  
Objects: (none)  
Options: (none)  
Server Window Start: 19:00:00 on 10/31/2002

### 4101: Backup/Archive statistics for scheduler

Displays backup and archive statistics, which might be useful in determining the success or failure of a command.

Event Type: Information  
Event Source: AdsmClientService  
Event Category: None  
Event ID: 4101  
Date: 10/30/2002  
Time: 8:29:21 PM  
User: DILE\Administrator  
Computer: MIKEDILE  
Description:  
Backup/Archive Statistics for Schedule Backup  
NIGHTLY\_BACKUP :  
-----  
Total number of objects inspected: 158,688  
Total number of objects backed up: 2,486  
Total number of objects updated: 0  
Total number of objects rebound: 0  
Total number of objects deleted: 0  
Total number of objects expired: 12  
Total number of objects failed: 0

Total number of bytes transferred: 1.15 GB  
Data transfer time: 104.35 sec  
Network data transfer rate: 11,564.84 KB/sec  
Aggregate data transfer rate: 866.99 KB/sec  
Objects compressed by: 100%  
Elapsed processing time: 00:23:11

### 4102: Restore/Retrieve statistics for scheduler

Similar to 4101, except that it is for scheduled restore and retrieve commands.

### 4103: Backup-Archive Client Service Startup Parameters

Event Type: Information  
Event Source: AdsmClientService  
Event Category: None  
Event ID: 4103  
Date: 10/31/2002  
Time: 8:29:56 AM  
User: DILE\Administrator  
Computer: MIKEDILE  
Description:  
Backup/Archive Client Service Startup  
Parameters:  
-----  
Service Name : TSM 515 Scheduler  
Last Update : Oct 14 2002  
Client PTF Level : 5.1.5.2  
Service Directory : D:\Program Files\  
Tivoli\TSM515\baclient  
Client Options File : E:\users\mikedile\  
logfiles\dsm.opt  
Client Node : MIKEDILE  
Comm Method : (default or obtained from  
client options file)  
Server : (default or obtained from client  
options file)  
Port : (default or obtained from client  
options file)  
Schedule Log : E:\users\mikedile\logfiles\  
dsmsched.log  
Error Log : E:\users\mikedile\logfiles\  
dsmerror.log  
MS Cluster Mode : (default or obtained  
from client options file)

### Journal Based Backup Service:

#### 4097: Informational message

4098: Warning message  
4099: Error message  
4100: Journal Based Backup service file  
monitor parameters  
4101: Journal Based Backup service database  
parameters  
4102: Journal Based Backup Service  
configuration parameters

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 145 for more information.

---

## Enabling or disabling scheduled commands

You can use the *schedcmddisabled* option to disable the scheduling of commands by the server. Commands are scheduled by using the *action=command* option on the **define schedule** server command.

The *schedcmddisabled* option does not disable the *preschedulecmd* and *postschedulecmd* commands. However, you can specify *preschedulecmd* or *postschedulecmd* with a blank or a null string to disable the scheduling of these commands. See “Schedcmddisabled” on page 350 for more information.

You can use the *schedrestretrdisabled* option to prevent the Tivoli Storage Manager Server administrator from executing restore or retrieve schedule operations. See “Schedrestretrdisabled” on page 357 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 378 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 379 for more information.

---

## Accessing Windows network drives during a scheduled backup

The Tivoli Storage Manager scheduler service logs into a user account when running. By default, the service will use the *System* account which does not have the ability to access network drives. In order for the scheduler service to access network drives, the service properties must be changed to log in as a different user which has the ability to log on as a service, to log in interactively, and to access the drives to map. The user account must log in once initially to map the required network drives persistently so that they are automatically mapped every time this account logs in.

To access network drives during a scheduled backup, you must specify a Windows domain user ID with authority to access the network drives you want to back up. Perform the following steps to establish this access.

1. Click the **Start** menu.
2. Select **Control Panel** → **Administrative Tools** → **Services** .
3. Select the scheduler service from the list of Windows services.
4. Select the **Log On** tab.
5. Ensure the radio button next to *This Account* is enabled in the Login As section.
6. Enter the domain account, or use the **Browse** button to locate the domain account.
7. Enter the password for the domain account.
8. Click **OK** and then click **Start**.

The scheduled backup will run using the security credentials of the Windows domain user you specified.

The Tivoli Storage Manager scheduler service is also available with the Tivoli Storage Manager Web client agent.

---

## Mobile dial-up support

Tivoli Storage Manager supports remote network connections to a server. With a remote network connection, mobile users no longer need to dial-in to their company network when a backup is scheduled to run. Tivoli Storage Manager automatically establishes a connection before the scheduled backup occurs. If the connection fails, Tivoli Storage Manager reestablishes the connection before attempting the backup.

Consult your operating system or Microsoft Internet Explorer help menus for current instructions to do the following:

*Configure a network or dial-up connection*

Options are available for the type of connection you want to establish.

*Configure autodial feature*

The schedule client connects to the internet and displays a dialup box if the autodial feature is not enabled.

*Add and configure a phone book*

Available phone numbers are placed in the operating system phone book. Consult your operating system or Microsoft Internet Explorer help menus for instructions on how to work with phone book entries.

---

## 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 31.

*Table 31. Client return codes and their meanings*

Code	Explanation
0	All operations completed successfully.

Table 31. Client return codes and their meanings (continued)

Code	Explanation
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 157.</li> </ul>
8	The operation completed with at least one warning message. For scheduled events, the status will be <i>Completed</i> . Review dsmerror.log (and dsmsched.log for scheduled events) to determine what warning messages were issued and to assess their impact on the operation.
12	The operation completed with at least one error message (except for error messages for skipped files). For scheduled events, the status will be <i>Failed</i> . Review the dsmerror.log file (and dsmsched.log file for scheduled events) to determine what error messages were issued and to assess their impact on the operation. As a general rule, this return code means that the error was severe enough to prevent the successful completion of the operation. For example, an error that prevents an entire drive from being processed yields return code 12. When a file is not found the operation yields return code 12.
<i>other</i>	For scheduled operations where the scheduled action is <i>COMMAND</i> , 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> .  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.

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 c:\MyTools\* -subdir=yes
incremental c:\MyPrograms\TestODBCDriver\* -subdir=yes
archive e:\TSM510C\* -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 574.

---

## Chapter 8. Understanding storage management policies

Storage management policies are rules your administrator defines in order to manage your backups and archives on the server. Your data is associated (or bound) to these policies; then when the data is backed up or archived, it is managed according to policy criteria. Policy criteria include a policy domain, a policy set, a management class, and a copy group .

Policies determine:

- Whether a file is eligible for backup or archive services.
- How many backup versions to keep.
- How long to keep inactive backup versions and archive copies.
- Where to place the copies in storage.
- For incremental backup, policies also determine:
  - How frequently a file can be backed up.
  - Whether a file must change before it is backed up again.

This chapter explains:

- Policy criteria (policy domains, policy sets, copy groups, and management classes).
- How to display policies.
- How Tivoli Storage Manager associates your data with policies.

---

### Using policy domains and policy sets

A *policy domain* is a group of clients with similar requirements for backing up and archiving data. Policy domains contain one or more policy sets. An administrator uses policy domains to manage a group of client nodes in a logical way. For example, a policy domain might include:

- A department, such as Accounting.
- A physical location, such as a particular building or floor.
- A local area network, such as all clients associated with a particular file server.

Tivoli Storage Manager includes a default policy domain named *Standard*. At first, your client node might be associated with the default policy domain. However, your administrator can define additional policy domains if there are groups of users with unique backup and archive requirements.

A *policy set* is a group of one or more management classes. Each policy domain can hold many policy sets. The administrator uses a policy set to implement different management classes based on business and user needs. Only one of these policy sets can be active at a time. This is called the *active policy set*. Each policy set contains a *default management class* and any number of additional management classes.

---

### Using management classes and copy groups

A *management class* is a collection of backup and archive copy groups that establishes and contains specific storage management requirements for backing up and archiving data. An administrator can establish separate management classes to meet the backup and archive requirements for different kinds of data, such as:

- System data that is critical for the business.

- Application data that changes frequently.
- Report data that Management reviews monthly.
- Legal information that must be retained indefinitely, requiring a large amount of disk space.

Most of the work you do with storage management policies is with management classes. Each file and directory that you back up and each file that you archive is associated with (or *bound* to) a management class, as follows:

- If your data is not associated (or bound) to a file with a management class, Tivoli Storage Manager uses the default management class in the active policy set.
- For backing up directories, you can specify a management class with an *include* statement or the *dirmc* option. If you do not specify a management class, Tivoli Storage Manager uses the management class in the active policy set specifying the longest retention period.
- For archiving directories, you can specify a management class with an *include.archive* statement or the *archmc* option. If you do not specify a management class, the server assigns the default management class to the archived directory. If the default management class has no archive copy group, the server assigns the management class that currently has the archive copy group with the shortest retention time.

You can use *include* statements in your include-exclude list to associate files with management classes. See “Selecting a management class for files” on page 159 for more information. In your client options file (*dsm.opt*), you can associate directories with a management class, using the *dirmc* option. See “Selecting a management class for directories” on page 160 for more information.

Within a management class, the specific backup and archive requirements are in *copy groups*. Copy groups define the specific storage management attributes that describe how the server manages backed up or archived data. Copy groups include both *backup copy groups* and *archive copy groups*. A management class can have one backup copy group, one archive copy group, both, or neither.

A *backup copy group* contains attributes that are used during the backup process to determine:

- How many days must elapse before a file is backed up again.
- How a file is processed during a backup if it is in use.

It also contains attributes to manage the backup versions of your files on the server. These attributes control:

- On which media type the server stores backup versions of your files and directories.
- How many backup versions the server keeps of your files and directories.
- How long the server keeps backup versions of your files and directories.
- How long the server keeps inactive backup versions.
- How long the last remaining inactive version of a file is kept.

An *archive copy group* contains attributes that control:

- Whether a file is archived if it is in use
- On which media type the server stores archived copies of your files
- How long the server keeps archived copies of your files

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 162.

---

## 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 32 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 32. Default values in the standard management class*

Attribute	Backup default	Archive default
Copy group name	Standard	Standard
Copy type	Backup	Archive
Copy frequency	0 days	CMD (Command)
Versions data exists	Two versions	Does not apply
Versions data deleted	One version	Does not apply
Retain extra versions	30 days	Does not apply
Retain only version	60 days	Does not apply
Copy serialization	Shared static	Shared static
Copy mode	Modified	Absolute
Copy destination	Backuppool	Archivepool
Retain versions	Does not apply	365 days

## Copy group name

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

## Copy type

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

## Copy frequency

*Copy frequency* is the minimum number of days that must elapse between successive incremental backups. Use this attribute during a full incremental backup. 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. Journal-based incremental backup differs from the traditional full incremental backup in that Tivoli Storage Manager does not enforce non-default copy frequencies (other than 0).

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 drive. 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 drive. This parameter is ignored as long as the file or directory remains on your drive.

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 system. Any subsequent updates to this parameter will not affect files that are already inactive. For example: If this parameter is set to 10 days when a file is inactivated during an incremental backup, the file will be deleted from the server in 10 days.

## Copy serialization

The *Copy Serialization* attribute determines whether a file can be in use during a backup or archive, and what to do if it is. The value for this attribute can be one of the following:

- **Static.** A file or 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:** For all supported Windows 32-bit versions: 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 *snapshotproviderimage* option. See “Snapshotproviderimage” on page 374 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 client options file. If the object is changed during every backup or archive attempt, it is not backed up or archived.
- **Dynamic.** A file or directory is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.

**Note:** For all supported Windows 32-bit versions: 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 *snapshotproviderimage* option. See “Snapshotproviderimage” on page 374 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 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.

For Windows 2003 and XP *only*: If open file support is configured (see “Configuring Open File Support (OFS)” on page 29), Tivoli Storage Manager performs a *snapshot* backup of files that are open (or locked) by other applications. The snapshot allows the backup to be taken from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to the file system are not included in the backup. See “Open file support for backup operations” on page 69 for more information.

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.
  - The file attributes, with the exception of archive, are different. If only the file meta-data changes (such as access permissions), but the file data does not change, Tivoli Storage Manager might back up only the meta-data.
- **Absolute.** The file is considered for incremental backup regardless of whether it changed since the last backup. For archive copy groups, the mode is always *absolute*, indicating that a file is archived regardless of whether it changed since the last archive request.

## Copy destination

Names the destination where backups or archives are stored. The destination can be either a storage pool of disk devices or a storage pool of devices that support removable media, such as tape.

## Retain versions

Specifies the number of days an archived file remains in storage. When the specified number of days elapse for an archived copy of a file, it is deleted from server storage.

---

## Selecting a management class for files

If the default management class meets the backup and archive requirements for all the files on your workstation, it is not necessary to take any action to associate your files with that management class. This is done automatically when you back up or archive your files.

When selecting a different management class for your files, consider these questions:

- Does the management class contain a backup copy group?  
If you attempt to back up a file associated with a management class that does not contain a backup copy group, the file is not backed up.
- Does the management class contain an archive copy group?  
You cannot archive a file associated with a management class that does not contain an archive copy group.
- Does the backup copy group contain attributes that back up your files often enough?  
Mode and frequency work together to control how often a file is backed up when you use incremental backup. Tivoli Storage Manager does not check those attributes for selective backup.
- Do the copy groups specify either static or shared static for serialization?  
If serialization is shared dynamic or dynamic, you might get fuzzy backups or archive copies. Verify that this is acceptable. For example, you might want to use shared dynamic or dynamic serialization for a file to which log records are continuously added. If you used static or shared static serialization, the file might never be backed up because it is constantly in use. With shared dynamic or dynamic serialization, the file is backed up, but the backup version of the file might contain inconsistent data. Do not use shared dynamic or dynamic serialization for a file if it is very important that the backup version or archive copy contain all changes.
- Does the backup copy group specify an adequate number of backup versions to keep, along with an adequate length of time to keep them?
- Does the archive copy group specify an adequate length of time to keep archived copies of files?

---

## Assigning a management class to files

A management class defines when your files are included in a backup, how long they are kept on the server, and how many versions of the file the server should keep. The server administrator selects a default management class. You can specify your own management class to override the default management class.

To assign a management class other than the default to directories, use the *dirmc* option in your client options file (*dsm.opt*). See “Dirmc” on page 234 for more information.

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

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

```
include c:\adsm\proj2\costs\* budget
```

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

```
include ?:\...\* managall
```

The example below demonstrates how to use a management class:

```
exclude ?:\...\*.sno
include c:\winter\...\*.ice mcweekly
include c:\winter\december\*.ice mcdaily
include c:\winter\january\*.ice mcmonthly
include c:\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 `c:\winter\january` directory is assigned to the management class, **mcmonthly**.
3. Any file with an extension of `ice` in the `c:\winter\december` directory is assigned to the management class, **mcdaily**.
4. Any other files with an extension of `ice` in any directory under `\winter` are assigned to the management class, **mcweekly**.
5. Any file with an extension of `sno` (except `c:\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 280 for more information about the *include* option.

---

## Overriding the management class for archived files

When you archive a file, you can override the assigned management class using the graphical user interface (GUI), or by using the *archmc* option on the **archive** command. Overriding the management class using the GUI is equivalent to using the *archmc* option on the **archive** command. To use the GUI, press the **Options** button on the archive tree to override the management class and select a different management class.

On the command line, to associate the file `budget.jan` with the management class **ret2yrs**, you would enter:

```
dsmc archive -archmc=ret2yrs c:\plan\proj1\budget.jan
```

---

## Selecting a management class for directories

If the management class in your active policy set containing the longest “Retain only version” (REONLY) 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 REONLY 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 566 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 183 for more information.

---

### Overview of processing options

Tivoli Storage Manager uses *processing options* that you specify in your client system 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
- 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 15 for information on how to create and modify your client 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 47 on page 185.

**Note:** Some of the processing options used by the Tivoli Storage Manager central scheduler are defined in the Windows registry when the schedule services are configured. These options can also be specified in the client options file. When the scheduler runs as a service, processing options specified in the registry override the same options specified in the client options file. For more information, see “Changing processing options used by the scheduler service” on page 577.

---

## Communication options

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

For all Windows clients, use one of the following protocols:

- TCP/IP
- Named pipes
- Shared memory

Use the *commmethod* option to specify the communication protocol. For more information, see “Commmethod” on page 217.

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 client options file. The other TCP/IP options have default values that you can modify if you want to change the default value.

Table 33. TCP/IP options

Option	Description	Page
<i>httpport</i>	Specifies a TCP/IP port address for the Tivoli Storage Manager Web client.	273
<i>lanfreetcpport</i>	Specifies the TCP/IP port number where the Tivoli Storage Manager storage agent is listening.	294
<i>lanfreetcpserveraddress</i>	Specifies the TCP/IP address for the Tivoli Storage Manager storage agent.	295
<i>tcpbuffsize</i>	Specifies the size, in kilobytes, of the Tivoli Storage Manager internal TCP/IP communication buffer.	389
<i>tcpnodelay</i>	Specifies whether the server or client disables the delay of sending successive small packets on the network.	393
<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.	388
<i>tcpcadaddress</i>	Specifies a TCP/IP address for dsmcad.	390
<i>tcpport</i>	Specifies the TCP/IP port address for a Tivoli Storage Manager server.	394
<i>tcpserveraddress</i>	Specifies the TCP/IP address for a Tivoli Storage Manager server.	395
<i>tcpwindowsize</i>	Specifies the size, in kilobytes, of the TCP/IP sliding window for your client node.	396
<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 service and the Web Client Agent service for communications with the Web GUI.	422

## Named Pipes option

The communication option for Named Pipes is:

Table 34. Named Pipes communication option

Option	Description	Page
<i>namedpipename</i>	Specifies the name of a named pipe to use for communications between a Tivoli Storage Manager client and server on the same Windows server domain.	310

## Shared memory options

Table 35. 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.	293
<i>shmport</i>	Specifies the unique number that is used by the client and the server to identify shared memory area used for communications.	293

---

## Node options

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

Table 36. node options

Option	Description	Page
<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.	194
<i>clusterdisksonly</i>	Specifies whether the Tivoli Storage Manager Backup-Archive client allows the backup of only clustered disks, when running in a Microsoft Cluster Server (MSCS) or VERITAS Cluster Server (VCS) environment and using the <i>clusternode=yes</i> processing option.	210
<i>clusternode</i>	Specifies whether Tivoli Storage Manager manages cluster drives in a Microsoft Cluster Server (MSCS) or Veritas Cluster Server (VCS). environment.	214
<i>nasnodename</i>	Specifies the node name for the NAS file server when processing NAS file systems.	311
<i>nodename</i>	Use the <i>nodename</i> option in your client options file <i>dsm.opt</i> to identify your workstation to the server to establish communications.	312
<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.	410

## Backup and archive processing options

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

Table 37. 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.	193
<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.	194
<i>autofsrename</i>	Specifies whether to rename an existing file space on a Unicode-enabled server so a Unicode-enabled file space can be created for the current operation.	202
<i>changingretries</i>	Specifies the number of retries when attempting to back up or archive a file that is in use.	207
<i>class</i>	Specifies whether to display a list of NAS, client, or WebSphere® Application Server objects during a <b>query backup</b> , <b>query filesystem</b> , or <b>delete filesystem</b> operation.	208
<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.	221
<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.	222
<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.	229
<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.	230
<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.	232
<i>dfsbackupmntpnt</i>	Specifies whether Tivoli Storage Manager sees a Dfs (NTFS or FAT) as a junction or a directory. This option is valid for Windows 2003 server <i>only</i> .	233

Table 37. Backup and archive processing options (continued)

Option	Description	Page
<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.	234
<i>dirsonly</i>	Backs up, restores, archives, retrieves, or queries directories <i>only</i> .	235
<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.	238
<i>domain</i>	Specifies the drives to include in your default client domain for an incremental backup.	239
<i>domain.image</i>	Specifies the file systems and raw logical volumes that you want to include in your client domain for an image backup.	241
<i>domain.nas</i>	Specifies the volumes to include in your default domain for NAS image backups.	242
<i>enablearchiveretentionprotection</i>	Allows the client to connect to a data retention server.	247
<i>enablelanfree</i>	Specifies whether to enable an available LAN-free path to a storage area network (SAN) attached storage device.	249
<i>exclude</i> <i>exclude.backup</i> <i>exclude.file</i> <i>exclude.file.backup</i>	Use these options to exclude a file or group of files from backup services.	257
<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.	250
<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.	251
<i>exclude.archive</i>	Excludes a file or a group of files that match the pattern from archive services <i>only</i> .	257
<i>exclude.compression</i>	Excludes files from compression processing if you set the <i>compression</i> option to <i>yes</i> . This option applies to backups and archives.	257
<i>exclude.dir</i>	Excludes a directory, its files, and all its subdirectories and their files from backup processing.	257
<i>exclude.encrypt</i>	Excludes specified files from encryption processing.	257
<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.	257

Table 37. Backup and archive processing options (continued)

Option	Description	Page
<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> .	257
<i>exclude.subfile</i>	Excludes files from adaptive subfile backup processing. This option does not apply to archive processing.	257
<i>exclude.systemobject</i>	Excludes individual system objects from backup services. Excluded system object types that you backed up previously are not expired during subsequent backups. This option only excludes the system object types that you specify from subsequent backups.	257
<i>exclude.systemservice</i>	Excludes individual system services components from backup services. Input can be the keyword or component name to be excluded from backup.	257
<i>encryptiontype</i>	Specifies whether to use AES 128-bit or DES 56-bit data encryption during backup and archive operations. AES 128-bit encryption is the default.	250
<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.	262
<i>filesonly</i>	Backs up, restores, retrieves, or queries files <i>only</i> .	266
<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.	271
<i>guitreeviewafterbackup</i>	Specifies whether the GUI client is returned to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.	272
<i>imagegapsize</i>	Specifies the minimum size of empty regions on a volume that you want to skip during backup. This option is valid for Windows 32-bit clients.	275
<i>incl excl</i>	Specifies the path and file name of an include-exclude options file.	278
<i>include</i> <i>include.backup</i> <i>include.file</i>	Use these options to include files or assign management classes for backup processing.	280
<i>include.archive</i>	Includes files or assigns management classes for archive processing.	280
<i>include.compression</i>	Includes files for compression processing if you set the <i>compression</i> option to <i>yes</i> . This option applies to backups and archives.	280
<i>include.encrypt</i>	Includes the specified files for encryption processing. By default, Tivoli Storage Manager does not perform encryption processing.	280



Table 37. Backup and archive processing options (continued)

Option	Description	Page
<i>include.fs</i>	Use the <i>include.fs</i> option to specify processing options for a given file system. Use the <i>include.fs</i> option to specify which drives use open file support and to control how full file space incremental backups are processed.	280
<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 fileclient options file (dsm.opt). See “Toc” on page 399 for more information.	280
<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. Use the <i>include.fs</i> option to specify which drives use open file support and to control how full file space incremental backups are processed.	280
<i>include.subfile</i>	Includes files for adaptive subfile backup processing. This option does not apply to archive processing.	280
<i>include.systemobject</i>	Assigns management classes for backup of Windows XP system objects. By default, Tivoli Storage Manager binds all system objects to the default management class. You cannot use this option to bind individual system object components to a different management class. You cannot use this option to include or exclude a system object from processing.	280
<i>include.systemstate</i>	Assigns management classes for back up of the Windows Server 2003 and Windows Vista system state. The default is to bind the system object to the default management class.	280
<i>incrbydate</i>	Use with the <b>incremental</b> command to request an incremental backup by date.	286
<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. <i>only</i> .	287
<i>incrthreshold</i>	The <i>incrthreshold</i> option specifies the threshold value for the number of directories in any journaled file space that might have active objects on the server, but no equivalent object on the workstation.	288
<i>memoryefficientbackup</i>	Specifies a memory-saving backup algorithm for incremental backups when used with the <b>incremental</b> command.	305

Table 37. Backup and archive processing options (continued)

Option	Description	Page
<i>mode</i>	Use the <i>mode</i> option with these commands, as follows:  <b>backup image</b> To specify whether to perform a selective or incremental image backup of client file systems.  <b>backup nas</b> To specify whether to perform a full or differential image backup of NAS file systems.  <b>backup was</b> To specify whether to perform a full or differential backup of the WebSphere Application Server Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server to the Tivoli Storage Manager server.  <b>backup group</b> To specify whether to perform a full or differential group backup containing a list of files from one or more file space origins.	307
<i>monitor</i>	Specifies whether you want to monitor an image backup of file systems belonging to a Network Attached Storage (NAS) file server.	309
<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.	315
<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.	314
<i>optfile</i>	Specifies the client options file you want to use when you start a Tivoli Storage Manager session.	318
<i>postsnapshotcmd</i>	During an online image backup or open file support operation, this option allows you to manually bring up an application after the snapshot provider starts a snapshot. This option is only valid if the OFS or online image support is enabled.	327
<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.	331

Table 37. Backup and archive processing options (continued)

Option	Description	Page
<i>presnapshotcmd</i>	During an online image backup or open file support operation, this option allows you to manually quiesce an application before the snapshot provider starts a snapshot. This option is only valid if the OFS or online image support is enabled.	336
<i>resetarchiveattribute</i>	Specifies whether Tivoli Storage Manager resets the Windows archive attribute on files that are successfully backed up to a Tivoli Storage Manager server. This option is valid for all Windows clients.	342
<i>skipntpermissions</i>	Specifies whether to back up, archive, retrieve or restore Windows security information.	365
<i>skipntsecuritycrc</i>	Specifies whether to compute the security CRC for permission comparison during subsequent backups. Use this option on all Windows clients.	366
<i>snapshotcachelocation</i>	During an online image backup or open file support operation use this option to specify the location where the LVSA places the cache files. This option is only valid if the LVSA is installed and configured for open file or online image support.	367
<i>snapshotcachesize</i>	Use this option to determine the size of the cache file containing the original data for blocks that change during the operation. 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 file system 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 <i>snapshotcachelocation</i> option. This option is only valid if the LVSA is installed and configured for open file or online image support.	368
<i>snapshotfsidleretries</i>	During an online image backup or open file support operation, use this option to specify the number of additional times the LVSA should try to achieve the snapshot file system idle wait time before the online image backup or open file support operation fails. This option is only valid if the LVSA is installed and configured for open file or online image support.	369
<i>snapshotfsidlewait</i>	During an online image backup or open file support operation, use this option to specify the amount of time that must pass in which there is no write activity (read activity is ignored) on a volume before a snapshot can be taken. This option is only valid if the LVSA is installed and configured for open file or online image support.	371

Table 37. Backup and archive processing options (continued)

Option	Description	Page
<i>snapshotproviderfs</i>	Use the <i>snapshotproviderfs</i> option to enable snapshot-based file backup and archive operations, and to specify a snapshot provider.	373
<i>snapshotproviderimage</i>	Use the <i>snapshotproviderimage</i> option to enable snapshot-based online image backup, and to specify a snapshot provider.	374
<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.	375
<i>subdir</i>	Specifies whether to include subdirectories of a named directory.	381
<i>subfilebackup</i>	Specifies whether Tivoli Storage Manager uses adaptive subfile backup.	383
<i>subfilecachepath</i>	Specifies the path where the client cache resides for adaptive subfile backup processing.	384
<i>subfilecachesize</i>	Specifies the client cache size for adaptive subfile backup.	385
<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.	386
<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.	399
<i>type</i>	Use the <i>type</i> option with the <b>query node</b> command to specify the type of node to query.	404
<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.	406
<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.	409

Table 37. Backup and archive processing options (continued)

Option	Description	Page
<i>vmchost</i>	Used with the <b>backup VM</b> or <b>query VM</b> commands to specify the host name of the VMware VirtualCenter or ESX server where the VMware Consolidated Backup commands are directed.	412
<i>vmcpw</i>	Used with the <b>backup VM</b> or <b>query VM</b> commands to specify the password of the VirtualCenter or ESX user specified with the <b>VMUSER</b> option.	413
<i>vmcuser</i>	Used with the <b>backup VM</b> or <b>query VM</b> commands to specify the user name of the VMware VirtualCenter or ESX server where the VMware Consolidated Backup commands are directed.	414
<i>vmlist</i>	Used with the <b>backup VM</b> or <b>query VM</b> commands to specify the host name or list of host names of the virtual machine to back up.	415
<i>wasnode</i>	If Websphere Application Server 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 Websphere Application Server on your system.	419
<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.	416
<i>washome</i>	Use the <i>washome</i> option in your client options file (dsm.opt) to specify an override base install path for the Application Server	417
<i>wasndhome</i>	Use the <i>wasndhome</i> option in your client options file (dsm.opt) to specify an override base install path for the Network Deployment Manager.	418
<i>wastype</i>	Use the <i>wastype</i> option with the <b>backup was</b> command to back up the Websphere Application Server Network Deployment Manager or the Application Server. If Websphere Application Server security is enabled, use the <i>wastype</i> option with the <b>set waspassword</b> command to specify the Websphere Application Server Network Deployment Manager or Application Server when setting the user name and password for each installation of Websphere Application Server on your system.	420
<i>wasuser</i>	If Websphere Application Server 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 Websphere Application Server on your system.	421

The following options are backup-archive client options that apply only to HSM for Windows migrated files. Refer to *IBM Tivoli Storage Manager for HSM for Windows Administration Guide* for details about these options:

- *Restorecheckstubaccess*
- *Restoremigstate*
- *Skipmigrated*

## Restore and retrieve processing options

The following options relate to restore and retrieve processing.

Table 38. Restore and retrieve processing options

Option	Description	Page
<i>activatekey</i>	Specifies whether to activate the Windows Registry key to update the Windows Registry after restoring files.	192
<i>asrmode</i>	Use this option with the <b>restore</b> , <b>restore systemobject</b> , and <b>restore systemstate</b> commands to specify whether to perform a restore operation in system ASR recovery mode. This option is used in the context of restore commands generated in the asr.sif file by the <b>backup asr</b> command only. This option should not be used outside the context of ASR recovery mode.	196
<i>backupsetname</i>	The <i>backupsetname</i> 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 <i>location</i> option.	204
<i>computername</i>	Use the <i>computername</i> option to specify the name of the computer for which you are performing system state recovery in Automated System Recovery (ASR) mode. This option is used in the context of restore commands generated in the asr.sif file by the <b>backup asr</b> command only. This option should not be used outside the context of ASR recovery mode.	224
<i>dirsonly</i>	Qualifies the operation (backup, archive, restore, retrieve) to process directories alone.	235
<i>disablenqr</i>	Specifies whether the Tivoli Storage Manager Backup-Archive Client can use the no-query restore method for restoring files and directories from the server.	236
<i>filelist</i>	Specifies a file containing a list of files to be processed by the specified command.	262
<i>filesonly</i>	Qualifies the operation (backup, archive, restore, retrieve) to process files alone.	266
<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.	267

Table 38. Restore and retrieve processing options (continued)

Option	Description	Page
<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.	268
<i>frontime</i>	Use the <i>frontime</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.	269
<i>frsprimaryrestore</i>	Specifies whether Tivoli Storage Manager allows the primary restoration of the SYSVOL system object if all replication partners have been lost. For example, in a disaster recovery situation when all systems containing the SYSVOL system object are lost, it is necessary to restore the first SYSVOL in this manner.	270
<i>guitreeviewafterbackup</i>	Specifies whether the GUI client is returned to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.	272
<i>ifnewer</i>	Replaces an existing file with the latest backup version only if the backup version is newer than the existing file.	274
<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.	276
<i>inactive</i>	Displays a list of active and inactive files when used with the <i>pick</i> option.	277
<i>latest</i>	Restores the most recent backup version of a file whether it is active or inactive.	298
<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.	299
<i>location</i>	Specifies where Tivoli Storage Manager searches for the backup set during a query or restore operation.	300
<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.	309
<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.	315
<i>optfile</i>	Specifies the client options file you want to use when you start a Tivoli Storage Manager session.	318

Table 38. Restore and retrieve processing options (continued)

Option	Description	Page
<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.	322
<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.	323
<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.	324
<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.	333
<i>replace</i>	Specifies whether to overwrite an existing file, or to prompt you for your selection when you restore or retrieve files.	340
<i>showmembers</i>	Displays all members of a group.	364
<i>subdir</i>	Specifies whether you want to include subdirectories of a named directory.	381
<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.	386
<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.	401
<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.	401
<i>type</i>	Use the <i>type</i> option with the <b>query node</b> command to specify the type of node to query.	404
<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. .	408
<i>washome</i>	Use the <i>washome</i> option in your client options file (dsm.opt) to specify an override base install path for the Application Server	417
<i>wasndhome</i>	Use the <i>wasndhome</i> option in your client options file (dsm.opt) to specify an override base install path for the Network Deployment Manager.	418



Table 38. Restore and retrieve processing options (continued)

Option	Description	Page
<i>wastype</i>	Use the <i>wastype</i> option with the <b>restore was</b> command to restore the Websphere Application Server Network Deployment Manager or Application Server.	420

The following options are backup-archive client options that apply only to HSM for Windows migrated files. Refer to *IBM Tivoli Storage Manager for HSM for Windows Administration Guide* for details about these options:

- *Restorecheckstubaccess*
- *Restoremigstate*
- *Skipmigrated*

## 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 39. Scheduling options

Option	Description	Page
<i>managedservices</i>	Specifies whether the Storage Manager Client Acceptor service manages the Web client, the scheduler, or both.	302
<i>maxcmdretries</i>	Specifies the maximum number of times the client scheduler attempts to process a scheduled command that fails.	304
<i>postschedulecmd, postnschedulecmd</i>	Specifies a command to process after running a schedule.	325
<i>preschedulecmd, prenschedulecmd</i>	Specifies a command to process before running a schedule.	329
<i>queryschedperiod</i>	Specifies the number of hours the client scheduler waits between attempts to contact the server for scheduled work.	338
<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.	347
<i>runasservice</i>	Forces the client command process to continue running, even if the account that started the client logs off. Use this option on all Windows clients.	349
<i>schedcmddisabled</i>	Specifies whether to disable the scheduling of generic commands specified by your Tivoli Storage Manager administrator.	350
<i>schedlogmax</i>	Specifies the maximum size of the scheduler log, in megabytes.	351
<i>schedlogname</i>	Specifies the path and file name where you want to store schedule log information.	352
<i>schedlogretention</i>	Specifies the number of days to keep log file entries in the schedule log, and whether to save pruned entries.	353

Table 39. Scheduling options (continued)

Option	Description	Page
<i>schedmode</i>	Specifies which schedule mode to use, <i>polling</i> or <i>prompted</i> .	355
<i>schedrestretrdisabled</i>	Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing restore or retrieve schedule operations.	357
<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.	361
<i>srprepuestoscheddisabled</i>	Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing pre-schedule and post-schedule commands when performing scheduled operations.	378
<i>srprepuestosnapdisabled</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.	379
<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 355) for details.	391
<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 355) for details.	392

## Format and language options

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

Table 40. Format and language options

Option	Description	Page
<i>dateformat</i>	Specifies the format for displaying dates.	227
<i>language</i>	Specifies the language used for messages.	296
<i>numberformat</i>	Specifies the format for displaying numbers.	316
<i>timeformat</i>	Specifies the format for displaying time.	397

---

## Command processing options

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

Table 41. 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.	244
<i>quiet</i>	Limits the number of messages that are displayed on your screen during processing. This option can be overridden by the server.	339
<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> .	358
<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.	359
<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.	407

---

## Authorization options

These options control access to a Tivoli Storage Manager server.

Table 42. Authorization options

Option	Description	Page
<i>password</i>	Specifies a Tivoli Storage Manager password.	319
<i>passwordaccess</i>	Specifies whether you want to use a generated password or be prompted for a password each time you start the client.	320
<i>revokeremoteaccess</i>	Restricts an administrator with client access privileges from accessing your workstation through the Web client.	348

---

## 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 43. Error processing options

Option	Description	Page
<i>errorlogmax</i>	Specifies the maximum size of the error log, in megabytes.	253
<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.	254
<i>errorlogretention</i>	Specifies how many days to maintain error log entries before pruning, and whether to save the pruned entries.	255

---

## Transaction processing options

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

Table 44. 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).	215
<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.	219
<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.	220
<i>diskbuffsize</i>	Specifies the maximum disk I/O buffer size (in kilobytes) that the client can use when reading files.	237
<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.  <b>Recommendation:</b> Discontinue the use of <i>largecommbuffers</i> because future releases of Tivoli Storage Manager might not accept this option.	
<i>resourceutilization</i>	Use the <i>resourceutilization</i> option in your fileclient options file <i>dsm.opt</i> to regulate the level of resources the Tivoli Storage Manager server and client can use during processing.	344
<i>txnbytelimit</i>	Specifies the number of kilobytes the client program buffers before it sends a transaction to the server.	403
<i>usedirectory</i>	Provides a convenient way to simplify client communication configuration by overriding <i>commmethod</i> parameters set in the client options file and instead querying the Active Directory for the communication method and server with which to connect.	405

---

## Web client options

The following are options for the Tivoli Storage Manager Web Client.

Table 45. Web client options

Option	Description	Page
<i>httpport</i>	Specifies a TCP/IP port address for the Web client.	273
<i>managedservices</i>	Specifies whether the Storage Manager Client Acceptor service manages the Web client, the scheduler, or both.	302
<i>revokeremoteaccess</i>	Restricts administrator access on a client workstation through the Web client.	348
<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 service and the Web Client Agent service for communications with the Web client.	422

---

## 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 514 for more information.

Table 46. 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.	225
<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.	264

---

## Using options with commands

You can override some of the options in your client 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 47 on page 185.

## 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="Project A" c:\devel\proj1\*
```

- For options that do not include parameters, enter a command, a dash (-) and the option name. For example,

```
dsmc incremental -quiet
```

**Note:** Use a leading dash (-) to indicate that the following text is the name of an option. If an object name begins with a dash, you must surround it in either single quotes (') or quotation marks ("). Most operating system command-line processors strip the quotes before submitting the command-line arguments to the Tivoli Storage Manager client application. In such cases, using escape characters or doubling the quotes allows the client to receive the quoted object name. In loop mode, surround such objects in either single quotes (') or quotation marks (").

- Enter either the option name, or an abbreviation for the option name. For example, to enter the *latest* option, enter either -lat or -latest. The capital letters in the syntax of each option indicate the minimum abbreviation for that option name. For information about how to read the syntax diagrams, see "Reading syntax diagrams" on page xii.
- Enter options before or after command parameters. For example, you can enter the option before or after a file specification:

```
dsmc selective -subdir=yes c:\devel\proj1\*  
dsmc selective c:\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" c:\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 bytes for a file name and file path combined is 8440. However, the file name itself cannot exceed 256 bytes and the path leading to the file cannot exceed 8184 bytes. Furthermore, directory names (including the directory delimiter) within a path are limited to 256 bytes. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary. Refer to Table 7 on page 18 for a description of file path names and limits.

Table 47 on page 185 lists client command options that you can enter only on the command line with specific commands.

Table 47. Client command options

Command option	Description	Commands	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.	<b>archive</b>	193
<i>asrmode</i>	Use this option with the <b>restore</b> , <b>restore systemobject</b> , and <b>restore systemstate</b> commands to specify whether to perform a restore operation in system ASR recovery mode. This option is used in the context of restore commands generated in the asr.sif file by the <b>backup asr</b> command only. This option should not be used outside the context of ASR recovery mode. This option is valid for the Windows XP and Windows 2003 .NET clients only.	<b>restore asr</b>	196
<i>computername</i>	Use the <i>computername</i> option to specify the name of the computer for which you are performing system state recovery in Automated System Recovery (ASR) mode. Use the <i>computername</i> option with restore commands generated in the asr.sif file by the <b>backup asr</b> command only. Do not use this option outside the context of ASR recovery mode. This option is valid for Windows XP and Windows Server 2003 clients only.	<b>restore systemobject</b> <b>restore systemstate</b> <b>restore backupset</b>	196
<i>class</i>	Specifies whether to display a list of NAS objects or client objects when using the following commands:	<b>query backup</b> <b>delete filespace</b> <b>query filespace</b>	208
<i>console</i>	Use the <i>console</i> option with the <b>query systeminfo</b> command to output system information to the console.	<b>query systeminfo</b>	225
<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 <b>restore image</b> 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.	<b>archive</b> <b>restore image</b>	229
<i>description</i>	Assigns or specifies a description for files when performing archive, delete, retrieve, or query archive operations.	<b>archive</b> <b>delete archive</b> <b>query archive</b> <b>query backupset</b> <b>retrieve</b>	230
<i>detail</i>	Displays management class, file space, backup, and archive information depending on the command with which it is used.	<b>delete filespace</b> <b>query archive</b> <b>query backup</b> <b>query filespace</b> <b>query mgmtclass</b>	232
<i>dirsonly</i>	Backs up, restores, archives, retrieves, or queries directories <i>only</i> .	<b>archive</b> <b>incremental</b> <b>query archive</b> <b>query backup</b> <b>restore</b> <b>restore backupset</b> <b>retrieve</b> <b>selective</b>	235

Table 47. Client command options (continued)

Command option	Description	Commands	Page
<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	262
<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.	query systeminfo	264
<i>filesonly</i>	Backs up, restores, retrieves, or queries files <i>only</i> .	archive incremental query archive query backup restore restore backupset retrieve selective	266
<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	267
<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.	query archive query backup query filespace query group query image query mgmtclass query was restore restore group restore image restore was retrieve	268
<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	269
<i>groupname</i>	Specifies the fully qualified name for a group.	backup group	271
<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	274



Table 47. Client command options (continued)

Command option	Description	Commands	Page
<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. .	<b>restore image</b>	276
<i>inactive</i>	Displays a list of active and inactive files when used with the <i>pick</i> option.	<b>delete group</b> <b>query backup</b> <b>query group</b> <b>query image</b> <b>query nas</b> <b>query systemobject</b> <b>query systemstate</b> <b>query was</b> <b>restore</b> <b>restore group</b> <b>restore image</b> <b>restore nas</b> <b>restore systemstate</b> <b>restore was</b>	277
<i>incrbydate</i>	Requests an incremental backup by date.	<b>incremental</b>	286
<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 <i>only</i> .	<b>restore image</b>	287
<i>latest</i>	Restores the most recent backup version of a file whether it is active or inactive.	<b>restore</b> <b>restore group</b> <b>restore was</b>	298
<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.	<b>query backupset</b> <b>restore backupset</b>	300
<i>mode</i>	Use the <i>mode</i> option with these commands, as follows:  <b>backup image</b> To specify whether to perform a selective or incremental image backup of client file systems.  <b>backup nas</b> To specify whether to perform a full or differential image backup of NAS file systems.  <b>backup was</b> To specify whether to perform a full or differential backup of the WebSphere Application Server Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server to the Tivoli Storage Manager server.  <b>backup group</b> To specify whether to perform a full or differential group backup containing a list of files from one or more file space origins.	<b>backup group</b> <b>backup nas</b> <b>backup image</b> <b>backup was</b> <b>restore nas</b>	307

Table 47. Client command options (continued)

Command option	Description	Commands	Page
<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.	<b>backup nas</b> <b>restore nas</b>	309
<i>nojournal</i>	Use this option with the 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.	<b>incremental</b>	314
<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.	<b>delete archive</b> <b>delete backup</b> <b>delete group</b> <b>expire</b> <b>restore image</b>	315
<i>optfile</i>	Specifies the client options file you want to use when you start a Tivoli Storage Manager session.	<b>dsmc.exe</b>	318
<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.	<b>delete archive</b> <b>delete group</b> <b>expire</b> <b>query nas</b> <b>restore</b> <b>restore asr</b> <b>restore group</b> <b>restore image</b> <b>restore nas</b> <b>restore was</b> <b>retrieve</b>	322
<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.	<b>query backup</b> <b>query group</b> <b>query image</b> <b>query nas</b> <b>query systemstate</b> <b>query was</b> <b>restore</b> <b>restore group</b> <b>restore image</b> <b>restore nas</b> <b>restore systemstate</b> <b>restore was</b> All query and restore system object commands	323
<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.	<b>query backup</b> <b>query image</b> <b>query nas</b> <b>query systemstate</b> <b>restore</b> <b>restore nas</b> <b>restore image</b> <b>restore systemstate</b> All query and restore system object commands	324

Table 47. Client command options (continued)

Command option	Description	Commands	Page
<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.	<b>restore</b> <b>restore backupset</b> <b>restore group</b> <b>retrieve</b> <b>restore was</b>	333
<i>runasservice</i>	Forces the client command process to continue running, even if the account that started the client logs off. Use this option on all Windows clients.	<b>schedule</b>	349
<i>showmembers</i>	Displays all members of a group.	<b>query group</b> <b>query systemstate</b> <b>query was</b> <b>restore group</b> <b>restore was</b>	364
<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.	<b>query archive</b> <b>query backup</b> <b>restore</b> <b>retrieve</b> <b>restore group</b> <b>restore was</b>	401
<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.	<b>query archive</b> <b>query backup</b> <b>restore</b> <b>retrieve</b> <b>restore group</b> <b>restore was</b>	402
<i>type</i>	Use the <i>type</i> option with the <b>query node</b> command to specify the type of node to query.	<b>query node</b>	404
<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.	<b>archive</b>	406
<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 Windows 32-bit platforms <i>only</i> .	<b>restore image</b>	408
<i>virtualfsname</i>	Specifies the name of the virtual file space for the group on which you want to perform the operation.	<b>backup group</b>	409
<i>wasnode</i>	Use the <i>wasnode</i> option with the <b>set waspassword</b> commands to specify the Websphere Application Server node name when performing the operation on the Websphere Application Server Network Deployment Manager or Application Server.	<b>set waspassword</b>	419
<i>wastype</i>	Use the <i>wastype</i> option with the <b>backup was</b> , <b>query was</b> , <b>restore was</b> , or <b>set waspassword</b> commands to perform the operation on the Websphere Application Server Network Deployment Manager or Application Server.	<b>backup was</b> <b>query was</b> <b>restore was</b> <b>set waspassword</b>	420

Table 47. Client command options (continued)

Command option	Description	Commands	Page
<i>wasuser</i>	If Websphere Application Server 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 Websphere Application Server on your system.	<b>set waspassword</b>	421

## 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 48. Options that are valid on the initial command line only

### Options valid on the initial command line

<i>asrmode</i>	<i>preschedulecmd</i> , <i>prenschedulecmd</i> (can be included in the schedule definition)
<i>backupregistry</i>	<i>presnapshotcmd</i>
<i>clusternode</i>	<i>queryschedperiod</i>
<i>commmethod</i>	<i>resourcutilization</i>
<i>computername</i>	<i>retryperiod</i>
<i>diskbuffsize</i>	<i>runasservice</i>
<i>editor</i>	<i>schedlogmax</i>
<i>enablelanfree</i>	<i>schedlogname</i>
<i>errorlogmax</i>	<i>schedlogretention</i>
<i>errorlogname</i>	<i>schedmode</i>
<i>errorlogretention</i>	<i>sessioninitiation</i>
<i>incrthreshold</i>	<i>subfilebackup</i>
<i>lanfreecommmethod</i>	<i>subfilecachepath</i>
<i>lanfreeshmport</i>	<i>subfilecachesize</i>
<i>lanfreetcpport</i>	<i>tcpbuffsize</i>
<i>language</i>	<i>tcpcadaddress</i>
<i>maxcmdretries</i>	<i>tcpclientaddress</i>
<i>namedpipename</i>	<i>tcpclientport</i>
<i>nodename</i>	<i>tcpport</i>
<i>optfile</i>	<i>tcpserveraddress</i>
<i>password</i>	<i>tcpwindowsize</i>
<i>postschedulecmd</i> , <i>postnschedulecmd</i> (can be included in the schedule definition)	<i>txnbytelimit</i>
<i>postsnapshotcmd</i>	<i>usedirectory</i>
	<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 xii 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**.

cr

---

## Activatekey

The **activatekey** option specifies whether to activate the Windows Registry key to update the Windows Registry after restoring files. Use this option with the **restore registry** command.

### Supported Clients

This option is valid for Windows XP 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). This option is not in the preference editor and must be added to the option file with a text editor.

### Syntax



### Parameters

- No* Specifies that Tivoli Storage Manager should not update the Windows Registry after restoring files.
- Yes* Specifies that Tivoli Storage Manager should update the Windows Registry after restoring files. This is the default.

### Examples

**Options file:**  
activate yes

**Command line:**  
-activate=yes

---

## 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 Windows 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 c:\plan\proj1\  
budget.jan\*
```

---

## Asnodename

### Authorized User

Use the *asnodename* option to allow agent nodes to back up or restore data on behalf of another node (the target node). This enables concurrent operations from multiple nodes to store data to the same target node and file space in parallel. Your client node must be granted access to the target node by the Tivoli Storage Manager server administrative client **grant proxynode** command.

When the Tivoli Storage Manager administrator grants a node proxy authority, and you use the *asnodename* option to become that node, you can query and restore all files as if you had root authority.

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

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

For example, you can use the following command to back up shared data for file space 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.
- You cannot use the *asnodename* option with system state components and system objects.

## Options File

Place this option in the dsm.opt file. 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

**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.

---

## Asrmode

Use the *asrmode* option with the **restore**, **restore systemobject**, and **restore systemstate** commands to specify whether to perform a restore operation in system ASR recovery mode. This option is used in the context of restore commands generated in the *asr.sif* file by the *backup asr* command only. This option should not be used outside the context of ASR recovery mode.

### Supported Clients

This option is valid for the Windows XP and Windows Server 2003 clients *only*.

### Syntax



### Parameters

- No* Specifies that Tivoli Storage Manager *does not* perform the restore operation in system ASR recovery mode. This is the default.
- Yes* Specifies that Tivoli Storage Manager performs the restore operation in system ASR recovery mode.

### Examples

#### Command line:

```
restore systemstate -asrmode=yes  
restore systemstate -asrmode=yes -inactive -pick
```

This option is valid for an interactive session, but cannot be changed by entering the option while running an interactive session.

---

## Auditlogging

Use the *auditlogging* option to generate an audit log which contains an entry for each file that is processed during an incremental, selective, archive, restore, or retrieve operation.

The audit log can be configured to capture either a basic level of information or a more inclusive (full) level of information.

The basic level of the audit logging feature captures the information that is in the schedule log and it records information that a file has been backed up, archived, updated, restored, retrieved, expired, deleted, skipped or failed during an incremental backup, selective backup, archive, restore or retrieve operation. In addition, the basic level of audit logging captures the input command for commands run through the Backup-Archive command line or scheduler clients.

The full level of audit logging records an action for each file that is processed by the Backup-Archive client. In addition to all of the events recorded by the basic level of audit logging, the full level of audit logging records information for a file that has been excluded or not sent during a progressive incremental backup operation because the file had not changed.

The following is an example of the messages that are issued when the audit log is configured to capture the basic level of information:

```
04/21/07 15:25:05 ANS1650I Command:
    sel c:\test\file.txt
04/21/07 15:25:05 ANS1651I Backed Up:
    \\spike\c$\test\file.txt
04/21/07 15:25:05 ANS1652I Archived:
    \\spike\c$\test\file.txt
04/21/07 15:25:05 ANS1653I Updated:
    \\spike\c$\test\file.txt
04/21/07 15:25:05 ANS1654E Failed:
    \\spike\c$\test\file.txt
04/21/07 15:25:05 ANS1655I Restored:
    \\spike\c$\test\file.txt
04/21/07 15:25:05 ANS1656I Retrieved:
    \\spike\c$\test\file.txt
04/21/07 15:25:05 ANS1657I Expired:
    \\spike\c$\test\file.txt
04/21/07 15:25:05 ANS1658I Deleted:
    \\spike\c$\test\file.txt
04/21/07 15:25:05 ANS1659I Skipped:
    \\spike\c$\test\file.txt
```

The following is an example of the messages that are issued when the audit log is configured to capture the full level of information (in addition to all messages issued for the basic level of audit logging):

```
04/21/07 15:25:05 ANS1660I Excluded:
    \\spike\c$\test\file.txt
04/21/07 15:25:05 ANS1661I Unchanged:
    \\spike\c$\test\file.txt
```

The audit log is not a substitute or a replacement for the standard error log (*dsmerror.log*) or for the schedule log (*dsm Sched.log*). If an error occurs that prevents a file from being processed, a message indicating that an error has occurred will be written to the audit log, but the message will not indicate the nature of the error. For problem diagnostics the standard error log must still be used.

The audit log entries only contain a time stamp and object name. There is no information to distinguish between files and directories or any information about the size of an object.

When using the Windows Backup-Archive Client, all object names will be written in the UNC format. Both the Windows and Macintosh Backup-Archive Clients will create the audit log as a Unicode file.

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

The **auditlogging** command is not supported with backup commands which interact with image-level objects such as **backup image** or **restore image**. The **auditlogging** command is supported with backup commands that interact with file-level objects such as **backup groups**, **backup was**, **backup systemobject**, and **backup systemstate**.

If you have enabled audit logging for an operation and there is a failure trying to write to the audit log (for example, the disk on which the audit log resides is out of space), the audit logging will be disabled for the rest of the operation and the return code for the operation will be set to 12, regardless of the outcome of the operation.

## Supported Clients

This option is valid for all clients.

## Options File

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

## Syntax



## Parameters

*off* Specifies that the audit logging facility is not engaged. This is the default.

*basic* Specifies that the audit log captures a basic level of information.

*full* Specifies that the audit log captures a more extensive level of information.

## Examples

Run an incremental backup with audit logging enabled.

### Command line:

```
dsmc i -auditlogging=basic
```

Back up a list of files using the maximum level of auditing, which enables a separate application, such as a Perl script, to verify the results.

```
| dsmd i -filelist=file.lst -auditlogging=full  
| -auditlogname="c:\program files\tivoli\tsm\baclient\  
| temp_audit001.log"  
  
|
```

---

## Auditlogname

The *auditlogname* option specifies the path and file name where you want to store audit log information. This option applies when audit logging is enabled. For more information about the audit logging facility refer to “Auditlogging” on page 197.

## Supported Clients

This option is valid for all clients.

## Options File

Place this option in the dsm.opt file.

## Syntax

▶—AUDITLOGName—*filespec*—▶

## Parameters

*filespec*

Specifies the path and file name where you want Tivoli Storage Manager to store audit log information.

If you specify a file name only, the file is stored in your current directory. The default is the installation directory with a file name of dsmaudit.log. The dsmaudit.log file *cannot* be a symbolic link.

## Examples

Run an incremental backup with audit logging enabled.

### Options file:

Store the audit log in a non-default path.

```
auditlogname c:\mypath\myaudit.log
```

### Command line:

Back up a list of files using the maximum level of auditing, which would enable a separate application, such as a Perl script, to verify the results:

```
dsmc i -filelist=file.lst -auditlogging=full  
-auditlogname="c:\program files\tivoli\tsm\baclient\  
temp_audit001.log"
```

### Sample output

The following is a sample execution and output file:

```
C:\Program Files\Tivoli\TSM\baclient>dsmc i  
c:\test\* -sub=yes -auditlogging=full  
IBM Tivoli Storage Manager  
Command Line Backup/Archive Client Interface  
Client Version 5, Release 5, Level 0.0  
Client date/time: 04/21/2007 15:52:25  
(c) Copyright by IBM Corporation and other(s)  
1990, 2007. All Rights Reserved.
```

```
Node Name: PATMOS  
Session established with server PATMOS_5331: Windows  
Server Version 5, Release 3, Level 3.1  
Server date/time: 04/21/2007 15:52:26  
Last access: 04/21/2007 15:52:06
```

```

Incremental backup of volume 'c:\test\*'
Normal File-->      1,048,576 \\patmos\c$\test
  \dir1\file1 [Sent]
Normal File-->      1,048,576 \\patmos\c$\test
  \dir1\file2 [Sent]
Normal File-->          1,024 \\patmos\c$\test
  \dir1\file3 [Sent]
Normal File-->      1,048,576 \\patmos\c$\test
  \dir2\file1 [Sent]
Normal File-->      1,048,576 \\patmos\c$\test
  \dir2\file2 [Sent]
Normal File-->          1,024 \\patmos\c$\test
  \dir2\file3 [Sent]
Successful incremental backup of '\\patmos\c$\test\*'

```

```

Total number of objects inspected:      12
Total number of objects backed up:      6
Total number of objects updated:        0
Total number of objects rebound:        0
Total number of objects deleted:         0
Total number of objects expired:         0
Total number of objects failed:          0
Total number of subfile objects:         0
Total number of bytes transferred:      400.85 KB
Data transfer time:                      0.00 sec
Network data transfer rate:              0.00 KB/sec
Aggregate data transfer rate:            382.85 KB/sec
Objects compressed by:                   91%
Subfile objects reduced by:              0%
Elapsed processing time:                 00:00:01
ANS1900I Return code is 0.
ANS1901I Highest return code was 0.

```

The following are the audit log contents:

```

04/21/2007 15:52:25 ANS1650I Command:
  i c:\test\*
04/21/2007 15:52:26 ANS1661I Unchanged:
  \\patmos\c$\test
04/21/2007 15:52:26 ANS1661I Unchanged:
  \\patmos\c$\test\dir1
04/21/2007 15:52:26 ANS1661I Unchanged:
  \\patmos\c$\test\dir2
04/21/2007 15:52:26 ANS1661I Unchanged:
  \\patmos\c$\test\file1
04/21/2007 15:52:26 ANS1661I Unchanged:
  \\patmos\c$\test\file2
04/21/2007 15:52:26 ANS1661I Unchanged:
  \\patmos\c$\test\file3
04/21/2007 15:52:26 ANS1651I Backed Up:
  \\patmos\c$\test\dir1\file1
04/21/2007 15:52:26 ANS1651I Backed Up:
  \\patmos\c$\test\dir1\file2
04/21/2007 15:52:26 ANS1651I Backed Up:
  \\patmos\c$\test\dir1\file3
04/21/2007 15:52:26 ANS1651I Backed Up:
  \\patmos\c$\test\dir2\file1
04/21/2007 15:52:26 ANS1651I Backed Up:
  \\patmos\c$\test\dir2\file2
04/21/2007 15:52:26 ANS1651I Backed Up:
  \\patmos\c$\test\dir2\file3

```

---

## Autofsrename

The *autofsrename* option renames an existing file space that is not Unicode-enabled on the Tivoli Storage Manager server so that a Unicode-enabled file space with the original name can be created for the current operation.

When you specify *autofsrename* yes in your client options file, and the server value of *autofsrename* is set to *client*, Tivoli Storage Manager generates a unique name by appending *\_OLD* to the file space name you specify in the current operation. For example, Tivoli Storage Manager renames the file space `\\your-node-name\h$` to `\\your-node-name\h$_OLD`. If the new file space name is too long, the suffix replaces the last characters of the file space name, as follows:

```
\\your-node-name_OLD
```

If the new file space name already exists on the server, Tivoli Storage Manager renames the new file space `\\your-node-name_OLDx`, where *x* is a unique number.

Tivoli Storage Manager creates new Unicode-enabled file spaces that contain only the data specified in the current operation. For example, to archive files from your H-disk named `\\your-node\h$`, issue the following archive command:

```
arc h:\logs\*.log
```

Before the archive takes place, the server renames the file space to `\\your-node\h$_OLD`. The archive places the data specified in the current operation into the Unicode-enabled file space named `\\your-node\h$`. The new Unicode-enabled file space now contains only the `\logs` directory and the `*.log` files specified in the operation. Tivoli Storage Manager stores all subsequent full and partial incremental, selective backup, and archive data in the new Unicode-enabled file spaces.

Renamed file spaces remain on the server as stabilized file spaces. *These file spaces contain all the original data, which you can restore as long as they remain on the server.*

**Note:** When an existing file space is renamed during Unicode conversion, any access rules defined for the file space remain applicable to the original file space. New access rules must be defined to apply to the new Unicode file space.

After installation, perform a full incremental backup and rename all existing file spaces that are not Unicode-enabled and back up the files and directories within them under the new Unicode-enabled file spaces. This operation requires increased processing time and storage on the server.

File spaces that are not Unicode-enabled can be viewed in the character set of the locale from which Tivoli Storage Manager backed up the files. A workstation running in a different locale might be unable to view or restore from these file spaces. Unicode-enabled file spaces that are backed up in one locale are visible in all other locales, provided that the workstation has the proper fonts installed. .

To restore or retrieve from a file space that is not Unicode-enabled, specify the source on the server and the destination on the client. See “Restoring from file spaces that are not Unicode-enabled” on page 524 for information on how to restore from file spaces that are not Unicode-enabled. See “Retrieving from file spaces that are not Unicode-enabled” on page 554 for information on how to retrieve from file spaces that are not Unicode-enabled.



## Supported Clients

This option is valid for all Windows clients. The server can define the *autofsrename* option and override the *autofsrename* setting on the client. The Tivoli Storage Manager client API does not support this option.

## Options File

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

## Syntax



## Parameters

*Yes* Specifies that Tivoli Storage Manager automatically renames all file spaces that are not Unicode-enabled in the current backup or archive operation.

*No* Specifies that Tivoli Storage Manager does not rename file spaces that are not Unicode-enabled in the current backup or archive operation.

### *Prompt*

Specifies that Tivoli Storage Manager prompts you whether to rename the file spaces that are not Unicode-enabled in the current operation. This is the default.

### Considerations:

- This option applies *only* when the server sets the *autofsrename* option to client.
- When the client scheduler is running, the default behavior is to not prompt you. The next interactive session prompts you to rename the file space.
- The client prompts you *only* one time per file space. If you specify no at the prompt, the client cannot rename the file spaces later. However, the Tivoli Storage Manager administrator can rename the file spaces on the server.
- When backing up files to a file space that is not Unicode-enabled, the Unicode-enabled client skips the files and directories with names containing characters from a code page that is different from the current locale.
- If files and directories with names containing characters from a code page other than the current locale were previously backed up with a client that was not Unicode-enabled, they might be expired. The Unicode-enabled client expires these files if you do not migrate the file space to a Unicode-enabled file space. You can back up and archive these files to a Unicode-enabled file space.

## Examples

### Options file:

```
autofsrename yes
```

---

## Backupsetname

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 systemstate**
- **restore**
- **restore backupset**
- **restore image**
- **restore systemobject**
- **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 c:\dir1\file.1 -backupsetname=YEAR_END_ACCOUNTING.1234  
dsmc query image -backupsetname=WEEKLY_BSET.1234
```

---

## Casesensitiveaware

The *casesensitiveaware* option specifies whether the Windows Backup-Archive client should filter out file and directory objects which have names conflicting in case only. For example, a set of objects called 'MyWork.xls', 'MYWORK.xls', and 'mywork.xls'.

NTFS volumes are case sensitive and allow case sensitive file names to be stored. Although the Windows operating system is not case sensitive, applications such as Windows Services for UNIX® (SFU) exploit POSIX and allow case sensitive file names. SFU is typically included with Windows operating systems such as Windows Powered OS and Windows Storage Server 2003. These operating systems are typically deployed on hardware (for example, NAS hardware) which is acting as a dedicated file server in a heterogeneous environment. If there are UNIX clients storing files on NTFS volumes in these Windows file server environments, it is highly recommended to use the *casesensitiveaware* option. If this option is not used in these environments unpredictable results will occur during backup and archive operations if case sensitive file name conflicts are encountered. For homogeneous Windows file server environments the *casesensitiveaware* option is not necessary.

Because the Windows operating system is not case sensitive, applications cannot distinguish between two objects named 'mywork.xls' and 'MyWork.xls'. For this reason, the Tivoli Storage Manager Windows Backup-Archive client cannot guarantee the restore integrity of such objects. When a name casing conflict arises, the Tivoli Storage Manager Backup-Archive client can only guarantee the restore integrity of the first file in an alphabetical sort. On an ASCII-based operating system such as Windows, this means that capital letters come first alphabetically before their lower-case counterparts, so 'MySwork.xls' would alphabetically precede 'mywork.xls'. In this example, if the *casesensitiveaware* option is used, only 'MyWork.xls' would be processed. An error message will be issued for 'mywork.xls' and it will be skipped. If 'mywork.xls' is a directory, then the directory subtree 'mywork.xls' would be skipped. In all cases, messages will be written to both the local error log and to the Tivoli Storage Manager server console to indicate the exact file names of the objects being skipped.

## Supported Clients

This option is valid for all Windows clients.

## Options File

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

## Syntax



## Parameters

*yes*

Specifies that the Tivoli Storage Manager should attempt to identify object names which differ in casing only and filter out objects which have casing conflicts and cannot be guaranteed to be restored properly.

*no* Specifies that the Tivoli Storage Manager should not attempt to identify object names which differ in casing only. This is the default.

---

## Changingretries

The *changingretries* option specifies how many additional times you want the client to attempt to back up or archive a file that is in use. Use this option with the **archive**, **incremental**, and **selective** commands.

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. Open file support can be used to back up files that are locked or in use.

## Supported Clients

This option is valid for all Windows 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 options file (dsm.opt). You can set this option on the **Backup** tab, **Number of retries if file is in use** field of the Preferences editor.

## Syntax

▶▶—CHAngingretries— *numberretries* —▶▶

## Parameters

*numberretries*

Specifies the number of times a backup or archive operation is attempted if the file is in use. The range of values is zero through 4; the default is 4.

## Examples

**Options file:**

```
changingretries 3
```

**Command line:**

```
-cha=3
```

---

## Class

The *class* option specifies whether to display a list of NAS or client objects when using the 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 all Windows clients. 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
```

---

## Clientview

The *clientview* option is available to users who have upgraded from the Tivoli Storage Manager Express backup client to the Tivoli Storage Manager enterprise client. You must be connected to a Tivoli Storage Manager Version 5.4 or higher server to use this option. The *clientview* option allows you to choose either the express view or the standard view of the client graphical user interface (GUI).

### Supported Clients

This option is valid for all Windows clients.

### Options File

Place this option in the dsm.opt file. To switch to the Express view:

1. In the IBM Tivoli Storage Manager window, select **Edit** → **Preferences** from the menu bar.
2. From the **General** tab of the Preferences editor, in the **Client View** field, click **Express**.
3. Click **OK** to save your change.

To switch to the Standard view:

1. In the IBM Tivoli Storage Manager window, click **Modify Settings**.
2. From the **General** tab of the Preferences editor, in the **Client View** field, click **Standard**.
3. Click **OK** to save your change.

### Syntax

▶▶ CLIENTVIEW =  $\left. \begin{array}{l} \textit{standard} \\ \textit{express} \end{array} \right\}$  ▶▶▶

### Parameters

#### *standard*

Specifies that the standard, or enterprise, view of the Tivoli Storage Manager client GUI should be used. The standard view contains the advanced features of the Tivoli Storage Manager backup-archive client GUI. This is the default.

#### *express*

Specifies that the express view of the Tivoli Storage Manager client GUI should be used. The express view contains the same features as the Express backup client GUI.





**Options file:**

```
CLUSTERNODE NO (default)
CLUSTERDISKSONLY YES (default)
DOMAIN ALL-LOCAL (default)
EXCLUDE c:\...\file.txt
```

**Scenario 1b: Back up a node which manages the local (non-clustered) drives and the system state information and bypass enumeration of cluster resources**

This is a scenario similar to scenario 1 which can be deployed if the Tivoli Storage Manager Backup-Archive client takes an inappropriate amount of time during startup processing. During initialization of the Backup-Archive client, all of the cluster resources are enumerated to determine which resources represent cluster disk devices. This processing can be skipped by coding the *clusterdisksonly* parameter to *no*.

**Options file:**

```
CLUSTERNODE NO (default)
CLUSTERDISKSONLY NO
DOMAIN C: D: (local drives must be explicitly enumerated)
EXCLUDE c:\...\file.txt
```

**Scenario 2: Back up a node which manages the local (non-clustered) drives and the system state information in a clustered environment using volume mount points as cluster resources**

This is the node that is basically dedicated to the restoration of the physical system should there be a hardware failure. There are clustered drives which appear as volume mount points in this environment (for example, IBM Tivoli SANergy, Windows 2003 server). Ensure that you remove any volume mount points from the incremental processing domain. For example, you have a volume mount point on a clustered drive (x:\mnt) and a non-clustered drive (c:\mnt):

**Options file:**

```
CLUSTERNODE NO (default)
CLUSTERDISKSONLY YES (default)
DOMAIN ALL-LOCAL -c:\mnt -x:\mnt
EXCLUDE c:\...\file.txt
```

**Scenario 3: Back up a node which manages the clustered drives within a cluster resource group**

In this scenario, it is assumed that the node that is responsible for backing up a cluster resource group has two drives, f: and g:. There are no clustered drives which appear as volume mount points. Ensure that you define the incremental processing domain as only the volumes within a cluster resource group. If you have multiple cluster resource groups, assign a unique Tivoli Storage Manager node to manage each cluster resource group.

**Options file:**

```
CLUSTERNODE YES
CLUSTERDISKSONLY YES (default)
DOMAIN f: g:
EXCLUDE f:\...\file.txt
```

**Scenario 3b: Back up a node which manages the clustered drives within a cluster resource group and bypass enumeration of cluster resources**

This is a scenario similar to scenario 3 which can be deployed if the Tivoli Storage Manager Backup-Archive client takes an inappropriate amount of time during startup processing. During initialization of the Backup-Archive client, all of the cluster resources are enumerated to determine which

resources represent cluster disk devices. This processing can be skipped by coding the *clusterdisksonly* parameter to *no*.

**Options file:**

```
CLUSTERNODE YES
CLUSTERDISKSONLY NO
DOMAIN f: g:
EXCLUDE f:\...\file.txt
```

**Scenario 4: Back up a node which manages the clustered drives within a cluster resource group, using volume mount points as cluster resources**

In this scenario, it is assumed that the node is responsible for backing up a cluster resource group that has two drives, f: and f:\mnt. There are clustered drives which appear as volume mount points (IBM Tivoli SANergy, Windows 2003 server). Ensure that you define the incremental processing domain as only the volumes within a cluster resource group. If you have multiple cluster resource groups, assign a unique Tivoli Storage Manager node to manage each cluster resource group.

**Options file:**

```
CLUSTERNODE YES
CLUSTERDISKSONLY NO
DOMAIN f: f:\mnt
EXCLUDE f:\mnt\...\file.txt
```

Table 49 shows the list of *clusternode* and *clusterdisksonly* combinations.

Table 49. *Clusternode* and *clusterdisksonly* combinations

Clusternode	Clusterdisksonly	When to use
no	yes	This is the default behavior if nothing is specified; since <i>clusterdisksonly</i> is set to <i>yes</i> , the cluster disk map will be built.
yes	yes	This is the basic way to run in a cluster node to back up cluster disks; the cluster disk map will be built.
yes	no	This is the way to run in a cluster environment if you have SANergy disks or cluster disks mounted on volume mount points. This is also a recommended way to run if you want to do both of the following: <ol style="list-style-type: none"> <li>1. Back up cluster resources</li> <li>2. Skip the processing at startup time</li> </ol> <p>This means that ALL-LOCAL cannot be used and you must explicitly code the DOMAIN statement.</p>

Table 49. **Clusternode** and **clusterdiskonly** combinations (continued)

Clusternode	Clusterdiskonly	When to use
no	no	<p>This is the recommended way to run if you want to do both of the following:</p> <ol style="list-style-type: none"> <li>1. Back up local resources in a cluster environment</li> <li>2. Skip the processing at startup time</li> </ol> <p>This means that ALL-LOCAL cannot be used and you must explicitly code the DOMAIN statement.</p>

---

## Clusternode

The *clusternode* option specifies whether Tivoli Storage Manager manages cluster drives in a Microsoft Cluster Server (MSCS) or Veritas Cluster Server (VCS) environment. For information on how to configure a Tivoli Storage Manager server to manage a cluster configured client, see Appendix D, “Configuring the backup-archive client in a cluster server environment,” on page 609.

When the *clusternode* option is set to *yes*, only shared cluster drives are available for backup and archive processing. When you set the *clusternode* option to *yes*, the node name will default to the cluster name.

To back up local drives or Windows Server 2003 systemstate, you must set the *clusternode* option to *no*.

**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 318 for more information.

## Supported Clients

This option is valid for Windows Server 2003.

## Options File

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

## Syntax



## Parameters

*Yes* Specifies that you want Tivoli Storage Manager to manage cluster drives in a MSCS or VCS environment..

*No* Specifies that you want to back up local disks. 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 Windows clients. The server can also define this option.

## Options File

Place this option in the client 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

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


### Supported Clients

This option is valid for all Windows clients.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Communication** tab, **Communication Method** list of the Preferences editor.

### Syntax



### Parameters

#### *TCPip*

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

#### *V6Tcpip*

Indicates that either TCP/IP Version 4 or Version 6 should be used, depending on the system configuration and the results of a domain name service lookup. A valid DNS environment must be available.

#### *NAMedpipes*

The interprocess communication method that permits message data streams to pass between a client and a server. Use this communication method with a Tivoli Storage Manager server running on the same workstation as the client.

#### *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.

**Note:** Use of this communication method requires that both client and server run under the same Windows account.

### Examples

#### Options file:

```
commMethod tcpip
```

Use only TCP/IP Version 4

```
commMethod V6Tcpip
```

Use both TCP/IP Version 4 and Version 6, depending on how the system is configured, and the results of a domain name service lookup.

| **Note:** The dsmsc schedule command cannot be used when both  
| *SCHEDMODE PRompt* and *COMMMethod V6Tcip* are specified.

| **Command line:**

| -commm=tcpi

| -commm=V6Tcip

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



---

## Commrestartduration

The *commrestartduration* option specifies the maximum number of minutes you want the client to try to reconnect to a Tivoli Storage Manager server after a communication error occurs.

**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 Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Communication** tab, **Common Options** section of the Preferences editor.

## Syntax

►►—COMMRESTARTDuration— *minutes* —————►►

## Parameters

*minutes*

The maximum number of minutes you want the client to attempt to reconnect with a server after a communication failure occurs. The range of values is zero through 9999; the default is 60.

## Examples

**Options file:**

```
commrestartduration 90
```

**Command line:**

Does not apply.

---

## Commrestartinterval

The *commrestartinterval* option specifies the number of seconds you want the client to wait between attempts to reconnect to a Tivoli Storage Manager server after a communication error occurs.

**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 Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Communication** tab, **Common Options** section of the Preferences editor.

## Syntax

►—COMMRESTARTInterval— *seconds* —————►

## Parameters

*seconds*

The number of seconds you want the client to wait between attempts to reconnect with a server after a communication failure occurs. The range of values is zero through 65535; the default is 15.

## Examples

**Options file:**

```
commrestartinterval 30
```

**Command line:**

Does not apply.

---

## Compressalways

The *compressalways* option specifies whether to continue compressing an object if it grows during compression. Use this option with the *compression* option.

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


## Supported Clients

This option is valid for all Windows clients. The server can also define this option.

## Options File

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

## Syntax



## Parameters

- Yes* File compression continues even if the file grows as a result of compression. This is the default.
- No* Backup-archive client objects are resent uncompressed if they grow during compression. API behavior depends on the application. Application backups might fail.

## Examples

**Options file:**  
compressalways yes

**Command line:**  
-compressa=no

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

---

## Compression

The *compression* option compresses files *before* you send them to the server. Compressing your files reduces data storage for backup versions and archive copies of your files. It can, however, affect Tivoli Storage Manager throughput. A fast processor on a slow network connection benefits from compression, but a slow processor on a fast network connection does not.

Use the *compression* option with the **archive**, **incremental**, and **selective** commands. The **backup image** command uses the *compression* option value specified in the dsm.opt. 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 options file (dsm.opt) to exclude specific files or groups of files from compression processing. See “Exclude options” on page 257 for more information.
- Use the *include.compression* option in your client options file (dsm.opt) to include files within a broad group of excluded files for compression processing. See “Include options” on page 280 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 Windows clients. The server can also define this option.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Backup** tab, **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.

---

## Computername

Use the *computername* option to specify the name of the computer for which you are performing system state recovery in Automated System Recovery (ASR) mode. See “Performing a Windows XP or Windows Server 2003 system recovery” on page 114 for complete instructions to perform this task.

Use the *computername* option with restore commands generated in the asr.sif file by the **backup asr** command only. Do not use this option outside the context of ASR recovery mode.

## Supported Clients

This option is valid for Windows XP and Windows Server 2003 clients only.

## Syntax

►►—COMPUTERNAME =—*computername*—————◄◄

## Parameters

*computername*

Specifies the name of the computer for which you are restoring the system state.

## Examples

### Command line:

```
restore systemstate -asrmode=yes -computername=originaljoe
```

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.
- ENV - Environment variables.
- ERRORLOG - The Tivoli Storage Manager error log file.
- FILE - Attributes for the file name that you specify.
- FILESNOTTOBACKUP - Enumeration of Windows Registry key:

```
HKEY_LOCAL_MACHINE\  
  SYSTEM\  
    CurrentControlSet\  
      BackupRestore\  
        FilesNotToBackup
```

This key specifies those files that backup products should not back up. The **query inclexcl** command will indicate that these files are excluded per the operating system.

- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- KEYSNOTTORESTORE - Enumeration of Windows Registry key:

```
HKEY_LOCAL_MACHINE\  
  SYSTEM\  
    ControlSet001\  
      BackupRestore\  
        KeysNotToRestore
```

This key specifies those Windows Registry keys that backup products should not restore.

- MSINFO - Windows system information (output from MSINFO32.EXE).
- OPTIONS - Compiled options.
- OSINFO - Name and version of the client operating system
- POLICY - Policy set dump.
- REGISTRY - Windows Tivoli Storage Manager-related Windows Registry entries.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually dsm Sched.log).
- SFP - The list of files protected by Windows System File Protection, and for each file, indicates whether that file exists. These files are backed up as part of the SYSFILES system object.
- SFP=<filename> - Indicates whether the specified file (*filename*) is protected by Windows System File Protection. For example:  
SFP=C:\WINNT\SYSTEM32\MSVCRT.DLL
- SYSTEMOBJECT - Windows system object information.
- CLUSTER - Windows 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 264 for more information.

## Supported Clients

This option is valid for all Windows 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.

Use this option if you want to change the default date format for the language of the message repository you are using.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time 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**
- **query systemobject**
- **query systemstate**
- **restore**
- **restore image**
- **restore nas**
- **retrieve**
- **restore registry**
- **restore systemobject**
- **restore systemstate**
- **set event**

## Supported Clients

This option is valid for all Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Regional Settings** tab, **Date Format** drop-down list of the Preferences editor.

## Syntax

►►—DATEformat— *format\_number* —————►►

## Parameters

*format\_number*

Displays the date using one of the following formats. Select the number that corresponds to the date format you want to use:

1 MM/DD/YYYY

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

## 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.

---

## 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 Windows 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 c:\foo\*.c -deletefiles
dsmc rest image c: -incre -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="2002 Budget for Proj 1" c:\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 Windows 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" c:\allproj\*.x
```

You can then use the description to retrieve all of the files.

## Examples

**Command line:**

```
dsmc archive -des="2003 Budget for Proj 1" c:\foo\ *.prj
```

---

## Detail

Use the *detail* option to display management class, file space, backup, and archive information depending on the command with which it is used.

Use the *detail* option with the **query mgmtclass** command to display detailed information about each management class in your active policy set. If you do not use the *detail* option, only the management class name and a brief description are displayed on the screen. If you specify the *detail* option, information about attributes in each copy group contained in each management class is displayed on the screen. A management class can contain a backup copy group, an archive copy group, both, or neither.

A Unicode-enabled file space might not display correctly if the server cannot display the Unicode name. In this case, use the file space identifier (fsID) of the file space to identify these file spaces on the server. Use the *detail* option with the **delete filesystem** and **query filesystem** 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 creation date of the file you specify.

- **delete filesystem**
- **query archive**
- **query backup**
- **query filesystem**
- **query mgmtclass**

## Supported Clients

This option is valid for all Windows 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 filesystem -detail
```

---

## Dfsbackupmntpnt

The *dfsbackupmntpnt* option specifies whether Tivoli Storage Manager views a Microsoft Dfs junction residing on an NTFS or FAT drive as a junction or a directory. If Tivoli Storage Manager views Microsoft Dfs junction as a junction, only the metadata of the junction is backed up or archived. The subtree under the junction point is not backed up or archived.

This option is effective only when you back up or archive a Microsoft Dfs root and is ignored when you back up or archive a Microsoft Dfs junction. To restore a Dfs tree, the root of the tree must already exist.

For more information on backing up a Dfs root, see “Backing up Microsoft Dfs files” on page 101.

### Supported Clients

This option is valid for Windows Server 2003 *only*.

### Options File

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

### Syntax



### Parameters

- Yes* Specifies that Tivoli Storage Manager views all Microsoft Dfs junctions as *junctions* and backs up only the name of any mounted junction it encounters during a backup operation. This is the default.
- No* Specifies that Tivoli Storage Manager views all Microsoft Dfs junctions as *directories* and backs up the contents of files and subdirectories of any junction it encounters during a backup operation.

### Examples

**Options file:**  
dfsbackupmntpnt no

**Command line:**  
Does not apply.

---

## Dirmc

The *dirmc* option specifies the management class you want to use for directories. If you do not specify this option to associate a management class with directories, the client program uses the management class in the active policy set of your policy domain with the longest retention period. 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 159 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 Windows clients. The server can also define this option.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Backup** tab, **Directory Management Class** section in the Preferences editor.

## Syntax

►—DIRMc— *mgmtclassname* —◄

## Parameters

*mgmtclassname*

Specifies the name of the management class 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 Windows 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 c:*
```

## Disablenqr

The *disablenqr* option specifies whether the Tivoli Storage Manager Backup-Archive Client can use the "no query restore" method for restoring files and directories from the server.

If you set the *disablenqr* option to *no* (the default), the client can use the "no query restore" process. Refer to "Standard query restore, no query restore, and restartable restore" on page 109 for more information about the "no query restore" process.

If you set the *disablenqr* option to *yes*, the client can use only the standard restore process (also known as "classic restore"). Refer to "Standard restore process" on page 109 for more information about the standard restore process.

**Note:** There is no option or value to specify that the client can use only "no query restore" method.

## Supported Clients

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

## Options File

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

## Syntax



## Parameters

*No* Specifies that Tivoli Storage Manager can use the "no query restore" method. This is the default.

*Yes* Specifies that the client uses only the "standard restore" method. The "no query restore" method is not allowed.

## Examples

**Options file:**  
`disablenqr yes`

**Command line**  
`-disablenqr=yes`

---

## Diskbuffsize

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

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

Optimal backup, archive migration client performance can usually be achieved if the value for this option is equal to or smaller than the amount of file read ahead provided by the client file system. A larger buffer will require more memory and might not improve performance.

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

## Supported Clients

This option is valid for all Windows clients.

## Options File

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

## Syntax

►►—DISKBuffsize— *size* —————►►

## Parameters

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

## Examples

**Options file:**  
diskbuffsize 64

---

## Diskcachelocation

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.opt* file or with the *include.fs* option. See “Include options” on page 280 for more information. If the *diskcachelocation* option appears in the option file, its value will be used for all file systems not represented by an *include.fs* option containing the *diskcachelocation* option.

The disk cache is a temporary file which is deleted after the **incremental** command is run. Use this option to select one of the following:

1. A location that has more free disk space if, when you are using *memoryefficientbackup=diskcachemethod*, you get the message that the disk cache file cannot be created because you do not have enough disk space.
2. A location on a different physical volume to reduce contention for the disk access mechanism, and therefore improve performance.

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

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

## Supported Clients

This option is valid for all clients. The server can also define this option.

## Options File

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

## 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 c:\temp  
diskcachelocation c:\tivoli\data
```

### Command line:

Does not apply.

---

## Domain

The *domain* option specifies the drives that you want to include for incremental backup in your client domain.

Use the *domain* option in your client 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 local drives to process during an incremental backup:

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

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

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

```
domain c: d: e:
```

and the following on the command line:

```
dsmc incremental -domain="g: h:"
```

Tivoli Storage Manager performs an incremental backup of the c: d: e: g: and h: local drives.

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 drives that you specify in the file specification. For example, if you enter:

```
dsmc incremental e: f: -domain="g: h:"
```

Tivoli Storage Manager performs an incremental backup for the e: and f: drives only.

You can also exclude drives, the systemobject domain, the system state (Windows Server 2003 and Windows Vista) domain by specifying the dash (-) operator before the drive or the systemobject domain. For example, in the following options Tivoli Storage Manager will process all local drives except for the c: drive, systemobject, and systemstate domain:

```
domain ALL-LOCAL -c: -systemobject  
domain ALL-LOCAL -c: -systemstate
```

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


## Supported Clients

This option is valid for all Windows 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 options file (dsm.opt). You can set this option on the Backup tab, **Domain for Backup** section of the Preferences editor.

## Syntax



## Parameters

### **all-local**

Backs up all local hard drives, the systemobject domain (Windows XP), and the systemstate domain (Windows Server 2003 and Windows Vista). This is the default.

### *domain*

Defines the drives to include in your default client domain.

When you use *domain* with the **incremental** command, it processes these drives in addition to those you specify in your default client domain.

### *-domain*

Defines the drives, the systemobject domain (Windows XP), and the systemstate domain (Windows Server 2003 and Windows Vista) to exclude in your default client domain.

### **systemobject**

Backs up all relevant system objects. Systemobject is included if **all-local** is specified. This is valid for Windows XP *only*.

### **systemstate**

Backs up Windows Server 2003 and Windows Vista system state. Systemstate is included if **all-local** is specified. This is valid for Windows Server 2003 and Windows Vista *only*.

## Examples

### Options file:

```
domain c: d: e:
domain c: systemobject
domain ALL-LOCAL -systemobject
domain ALL-LOCAL -c:
domain ALL-LOCAL -\\florence\e$
```

### Command line:

```
-domain="c: d:"
-domain="ALL-LOCAL -c: -systemobject"
```

---

## Domain.image

The *domain.image* option specifies the 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 Windows 32-bit platforms *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 options file (dsm.opt). You can set this option on the **Backup** tab → **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 d: e: f: domain.image f:\mnt\raw\rawmnt1  
f:\mnt\fs\fsmnt1
```

**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 options file (dsm.opt), 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.opt file. For example, if you enter the following in your dsm.opt 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.opt file, *all-nas* takes precedence.

## Supported Clients

This option is valid for all Windows clients. The server can also define this option.

## Options File

Place this option in the client options file (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.

This option is ignored for the Windows client. This option is always off, even if you explicitly specify *yes*. This is because the client uses the command-line history capabilities of the Windows command-line console.

## Supported Clients

This option is valid for all Windows clients.

## Options File

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

## Syntax



## Parameters

*Yes* Turns on the CLI editor and command retrieve capability. This is the default. However, for all Windows clients, the value for this option is always *editor=No*.

*No* Turns off the CLI editor and command retrieve capability.

## Examples

**Options file:**  
editor yes

**Command line:**  
-editor=yes

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

---

## Enable8dot3namesupport

The *enable8dot3namesupport* option specifies whether the client backs up and restores short 8.3 names for files that have long names on NTFS file systems.

### Supported Clients

This option is valid for all Windows clients.

A file with a long file name might not have a short 8.3 name if short name generation is disabled on the Windows system. This option is effective only for NTFS file systems.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the General tab of the Preferences editor.

### Syntax



### Parameters

*No* Short 8.3 names for files with long file names will not be backed up or restored. This is the default.

*Yes* Short 8.3 names for files with long file names will be backed up and restored. Each short name will use up to 14 additional bytes in the server database. Although this is a small number, if there are a large number of files with short 8.3 names on a large number of Windows systems, this can impact the size of the Tivoli Storage Manager server database.

**Recommendation:** Consult with your Tivoli Storage Manager server administrator before using this option.

The first backup that runs with this option will cause all files that have short 8.3 names to be updated on the Tivoli Storage Manager server, even if the files have not otherwise changed. This is because the client is adding the short 8.3 names to the active backup versions.

If this option is enabled for restore, the client will attempt to set the short 8.3 name for restored files, even if short name generation is disabled on the Windows system. The client must run under a Windows account that possesses the SE\_RESTORE\_NAME privilege in order for this option to be effective. See your system administrator if you have questions about account privileges.

During restore, a file's short 8.3 name will not be restored if another object in the same directory already has the same short 8.3 name. In this case, the file will be restored and an informational message will be logged indicating that the short name could not be set. If the file must be restored with its original short name, you need to resolve the conflict with the existing file, and then try the restore again.

**Attention:** This parameter can cause unexpected results in some cases. For example, if a file's short name happens to change between the last time the file was backed up and the time it is restored, and there is a link or registry entry that refers to the newer short name, then restoring the file with the older short name will invalidate the references to the newer short name.

## Examples

**Options file:**

```
enable8dot3namesupport yes
```

**Command line:**

```
-enable8dot3namesupport=yes
```

---

## Enablearchiveretentionprotection

The *enablearchiveretentionprotection* option allows the client to connect to a Tivoli Storage Manager data retention server to ensure that archive objects will not be deleted from the server until policy-based retention requirements for that object have been satisfied. This option is ignored if the client connects to a server that is not retention protection enabled. If the option is 'no' (the default) and an attempt is made to connect to a data retention server, the connection will be refused.

The data retention server is specially configured for that task, so normal backup or restore processing is rejected by the server. When the client is connected to a data retention server, the following commands will not be available. If you attempt to use these commands, a message will be displayed indicating that they are not valid with this server.

- incremental
- backup (all subcommands)
- selective
- restore (all subcommands except **restore backupset -location=file** or **-location=tape**)

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

- restart restore
- delete backup
- delete group
- expire
- All queries *except*:
  - query access
  - query archive
  - query filespace
  - query inclexcl
  - query managementclass
  - query node
  - query options
  - query schedule
  - query session
  - query systeminfo
  - query tracestatus
  - query TSA


### Supported Clients

This option is valid for all Windows clients.

### Options File

This option is valid only in client options file (dsm.opt) 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

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 *IBM Tivoli Storage Manager for Windows Storage Agent User's Guide*, SC32-0133.

### 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 529 for more information.

To specify a communication protocol between the Tivoli Storage Manager client and Storage Agent, see "Lanfreecommmethod" on page 291 for more information.

## Supported Clients

This option is valid for all Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **General** tab → **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, providing a stronger form of data encryption than DES 56-bit data encryption. The encryption type only affects backup and archive operations. The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received. During restore and retrieve operations, the encrypted data will be decrypted with the proper encryption algorithm, regardless of the setting for this option.

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

## Supported Clients

This option is valid for all Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can also set this option on the **Authorization** tab of the Preferences editor. The server can override this.

## Syntax



## Parameters

*AES128*

AES 128-bit data encryption. AES 128-bit data encryption provides a stronger form of data encryption than DES 56-bit data encryption. This is the default.

*DES56*

DES 56-bit data encryption.

## Examples

**Options file:**

```
encryptiontype aes128
```

**Command line:**

Does not apply.



## Encryptkey

The Tivoli Storage Manager client supports the option to encrypt files being backed up or archived to the Tivoli Storage Manager server. This option is enabled with the *include.encrypt* option. All files matching the pattern on the *include.encrypt* specification will be encrypted before the data is sent to the server. There are three options for managing the key used to encrypt the files (prompt, save, and generate). All three options can be used with either the Backup-Archive client or the Tivoli Storage Manager API.

### Notes:

1. The Tivoli Storage Manager API has an alternate way of specifying *encryptkey=generate*; the previous *enableclientencryptkey=yes* option can also be specified to request generate encryption processing.
2. The *enableclientencryptkey=yes* API option is still supported, so it is possible when using the API to specify two conflicting options. For example, *enableclientencryptkey=yes* and *encryptkey=prompt* or *encryptkey=save*.
3. When conflicting values are specified, the Tivoli Storage Manager API will return an error message.

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

## Supported Clients

This option is valid for all clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Authorization** tab, **Encryption Key Password** section of the Preferences editor.

## Syntax



## Parameters

### *save*

The encryption key password is saved in the Tivoli Storage Manager client's password file. A prompt is issued for an initial encryption key password, and after the initial prompt, the saved encryption key password in the password file is used for the backups and archives of files matching the *include.encrypt* specification. This key can be up to 63 bytes in length. The key is retrieved from the password file on restore and retrieve operations.

When the *save* option is specified for an API application, the initial key password must be provided by the application using the API in the *dsmInitEx* function call. The API itself does not issue a prompt to the user but relies on the application to prompt the user as necessary.

This is the default.

|

| **Restrictions:**

- This option can only be used when *passwordaccess generate* is also specified.
- The root user or a Tivoli Storage Manager authorized user must specify the initial encryption key password.
- The default for *encryptkey* is *save*, and the default for *passwordaccess* is *prompt*. However, if you specify *encryptkey=save*, you must also specify *passwordaccess generate*, when you specify include encrypt.

| *prompt*

| The management of the encryption key password is provided by the user. The user is prompted for the encryption key password when the Tivoli Storage Manager client begins a backup or archive. A prompt for the same key is issued when restoring or retrieving the encrypted file. This key can be up to 63 bytes in length.

| When the *prompt* option is specified for an API application, the key password must be provided by the application using the API in the *dsmlnitEx* function call. The API itself does not issue a prompt to the user but relies on the application to prompt the user as necessary.

| **Restriction:** This option can only be used by the root user or a Tivoli Storage Manager authorized user.

| *generate*

| An encryption key password is dynamically generated when the Tivoli Storage Manager client begins a backup or archive. This generated key password is used for the backups of files matching the *include.encrypt* specification. The generated key password, in an encrypted form, is kept on the Tivoli Storage Manager server. The key password is returned to the Tivoli Storage Manager client to enable the file to be decrypted on restore and retrieve operations.

|

## Examples

| **Options file:**

| encryptkey prompt

| **Command line:**

| Does not apply.

---

## Errorlogmax

The *errorlogmax* option specifies the maximum size of the error log, in megabytes.

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 Windows clients.

## Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **General** tab, **Select Error Log** button of the Preferences editor.

## Syntax

►—ERRORLOGMAX— *size* —————◄◄

## Parameters

*size*

Specifies the maximum size, in megabytes, for the log file. The range of values is 0 to 2047; the default is 0, which disables log file wrapping and allows the log file to grow indefinitely.

## Examples

**Options file:**

```
errorlogmax 2000
```

**Command line:**

```
-errorlogmax=2000
```

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

---

## Errorlogname

This option specifies the fully qualified path and file name of the file 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 Windows clients.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the **General** tab, **Select Error Log** button of the Preferences editor.

### Syntax

►—ERRORLOGName— *filespec* —————►

### Parameters

*filespec*

The fully qualified path and file name in which to store error log information. If any part of the path you specify does not exist, Tivoli Storage Manager attempts to create it.

### Examples

**Options file:**

```
errorlogname c:\temp\dsmerror.log
```

**Command line:**

```
-errorlogname=c:\temp\dsmerror.log
```

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

The location of the log file specified using the Client Service Configuration Utility or the client configuration wizard overrides the location specified in the client options file (dsm.opt).

---

## Errorlogretention

The *errorlogretention* option specifies how many days to maintain error log entries before pruning, and whether to save the pruned entries. The error log is pruned when the first error is written to the log after a Tivoli Storage Manager session is started. If the only session you run is the client scheduler, and you run it twenty-four hours a day, the error log might not be pruned according to your expectations. Stop the session and start it again to allow the scheduler to prune the error log.

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 Windows clients.

## Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **General** tab, **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

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.

You can exclude remotely accessed files by specifying Universal Naming Convention (**UNC**) names in your exclude statement. See “Excluding files with UNC names” on page 37 for examples of statements using UNC file names.

### Notes:

1. When you exclude a file that was previously included, existing backup versions become inactive during the next incremental backup.
2. The *exclude* statements are not case sensitive.
3. The server can define exclude options with the *inclexcl* option.
4. As with other include-exclude statements, you can use the *inclexcl* option to specify a file that can be in Unicode format, containing exclude statements with file names in Unicode. See “Inclexcl” on page 278 for more information.

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.

**Attention:** See “Excluding system files” on page 37 for a list of files that you should always exclude.

Use wildcard characters to exclude a broad range of files. See “Including and excluding groups of files” on page 37 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 c:\any\test
```

To exclude subdirectories that begin with test under the any directory, enter the following:

```
exclude.dir c:\any\test*
```

**Note:** If you define an exclude statement without using a drive letter, such as `exclude.dir code`, this will exclude the code directory on any drive from processing.

## Supported Clients

This option is valid for all Windows clients.

## Options File

Place these options in the client options file (`dsm.opt`). You can set these options on the **Include-Exclude** tab, **Define Include-Exclude Options** section of the Preferences editor.

## Syntax

►—*options pattern*—►

### **exclude, exclude.backup, exclude.file, exclude.file.backup**

Use these options to exclude a file or group of files from backup services.

### **exclude.archive**

Excludes a file or a group of files that match the pattern from archive services *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 c:\test\dan\data1` excludes the `c:\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.

**Note:** If you define an exclude statement without using a drive letter, such as `exclude.dir code`, this will exclude the code directory on any drive from processing.

### **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.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.dir* statements. This option is valid for all Windows clients.

### **exclude.image**

Excludes mounted file systems and raw logical volumes that match the specified pattern from full image backup operations. Incremental image backup operations are unaffected by *exclude.image*. This option is valid for Windows XP clients *only*.

### **exclude.subfile**

Excludes files from adaptive subfile backup processing. This option does not apply to archive processing.



### **exclude.systemobject**

Excludes individual system objects from backup services. Excluded system object types that you backed up previously are not expired during subsequent backups. This option only excludes the system object types that you specify from subsequent backups. This option is valid for Windows XP clients only.

### **exclude.systemservice**

Excludes individual system services components from system state backup. Input can be the keyword or component name to be excluded from processing. If you specify the component name and there are spaces, enclose the name in quotation marks. Table 50 shows valid system services components and their corresponding keywords:

Table 50. System services components and corresponding keywords

Component	Keyword
Background Intelligent Transfer Service	BITS
Event log	EVENTLOG
Removable Storage Management	RSM
Cluster Database	CLUSTERDB
Remote Storage Service	RSS
Terminal Server Licensing	TLS
Windows Management Instrumentation	WMI
Internet Information Services (IIS) metabase	IIS
DHCP database	DHCP
Wins database	WINSDB

If you are specifying a component that is not listed in Table 50, you can specify the component name. This option is valid for Windows Server 2003 and Windows Vista *only*.

## **Parameters**

### *pattern*

Specifies the file or group of files that you want to exclude.

**Note:** For NAS file systems: You must prefix the NAS node name to the file specification to specify the file server to which the exclude statement applies. If you do not specify a NAS node name, the file system identified refers to the NAS node name specified in the client options file (dsm.opt) 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 file system or raw logical volume.
- The valid values for the *exclude.systemobject* option are ACTIVEDIRECTORY, CERTSERVERDB, CLUSTERDB, COMPLUSDB, EVENTLOG, FRS, REGISTRY, RSM, SYSFILES, SYSVOL, and WMI.

## Examples

### Options file:

```
exclude ?:\...\swapper.dat
exclude "*/ea data. sf"
exclude ?:\io.sys
exclude ?:\...\spart.par
exclude c:\*\budget.fin
exclude c:\devel\*
exclude.dir c:\home\jodda
exclude.archive c:\home\*.obj
exclude.encrypt c:\system32\mydocs\*
exclude.compression c:\test\file.txt
exclude.subfile c:\test\file.txt
exclude.fs.nas netappsj/vol/vol0
exclude.image c:
exclude.systemobject sysfiles
exclude.systemservice eventlog
```

### Command line:

Does not apply.

## 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 222 for more information.

If you set the *compression* option to *yes* and no *exclude.compression* statements exist, Tivoli Storage Manager considers all files for compression processing.

- Tivoli Storage Manager processes *exclude.dir* and other include-exclude statements first. Tivoli Storage Manager then considers any *exclude.compression* statements. For example, consider the following include-exclude list:

```
exclude c:\test\*.*
exclude.compression c:\test\file.txt
include c:\test\file.txt
```

Tivoli Storage Manager examines the statements (reading from bottom to top) and determines that the c:\test\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.

## Adaptive subfile backup processing

If you want to exclude files from adaptive subfile backup processing using the *exclude.subfile* option, consider the following:

- You must set the *subfilebackup* option to *yes* to enable include-exclude adaptive subfile backup processing. If you do not specify the *subfilebackup* option or you set the *subfilebackup* option to *no*, Tivoli Storage Manager does not perform adaptive subfile backup processing. See “Subfilebackup” on page 383 for more information.

If you set the *subfilebackup* option to *yes* and no *exclude.subfile* statements exist, Tivoli Storage Manager considers all files for adaptive subfile backup processing.

- Tivoli Storage Manager processes *exclude.dir* and other include-exclude statements first. Tivoli Storage Manager then considers any *exclude.subfile* statements. For example, consider the following include-exclude list:

```
exclude c:\test\*.*
exclude.subfile c:\test\file.txt
include c:\test\file.txt
```

Tivoli Storage Manager examines the statements (reading from bottom to top) and determines that the `c:\test\file.txt` file is a candidate for backup, but is not a candidate for adaptive subfile backup.

- Include-exclude adaptive subfile backup processing is valid for backup and restore processing *only*. The *exclude.subfile* option does not affect whether files are excluded from backup or restore processing, only whether they are excluded from adaptive subfile backup processing.
- As with other include-exclude statements, you can specify *exclude.subfile* statements in Unicode using the *inclxcl* option. See “Inclxcl” on page 278 for more information.

| The *subfilebackup* option does not work correctly for migrated files. If you use a  
| combination of subfilebackup and non-subfilebackup for migrated files, your data  
| might be corrupted on the server.

## 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.
- The filelist can be an MBCS file or a Unicode file with all Unicode entries.
- Tivoli Storage Manager ignores any entry that is not valid.

The following is an example of a list of files within a filelist:

```
c:\myfiles\directory\file1
c:\tivoli\mydir\yourfile.doc
..\notes\avi\dir1
..\fs1\dir2\file3
"d:\fs2\Ha Ha Ha\file.txt"
"d:\fs3\file.txt"
```

You can use the *filelist* option during an open file support operation. In this case, Tivoli Storage Manager processes the entries in the filelist from the virtual volume instead of the real volume.

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=c:\filelist.txt d:\dir\`

the file specification `d:\dir\` represents the restore destination for all entries in the filelist. However, in the **selective** command:

```
selective -filelist=c:\filelist.txt d:\dir\
```

the file specification `d:\dir\` 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 `c:\dir0\myfile` and `c:\dir0`. If the `\dir0` directory does not exist on the server, the `c:\dir0` directory is sent to the server a second time.

## Supported Clients

This option is valid for all Windows 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=c:\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.
- ENV - Environment variables.
- ERRORLOG - The Tivoli Storage Manager error log file.
- FILE - Attributes for the file name that you specify.
- FILESNOTTOBACKUP - Enumeration of Windows Registry key:

```
HKEY_LOCAL_MACHINE\  
  SYSTEM\  
    CurrentControlSet\  
      BackupRestore\  
        FilesNotToBackup
```

This key specifies those files that backup products should not back up. The **query inclexcl** command will indicate that these files are excluded per the operating system.

- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- KEYSNOTTORESTORE - Enumeration of Windows Registry key:

```
HKEY_LOCAL_MACHINE\  
  SYSTEM\  
    ControlSet001\  
      BackupRestore\  
        KeysNotToRestore
```

This key specifies those Windows Registry keys that backup products should not restore.

- MSINFO - Windows system information (output from MSINFO32.EXE).
- OPTIONS - Compiled options.
- OSINFO - Name and version of the client operating system
- POLICY - Policy set dump.
- REGISTRY - Windows Tivoli Storage Manager-related Windows Registry entries.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually dmsched.log).
- SFP - The list of files protected by Windows System File Protection, and for each file, indicates whether that file exists. These files are backed up as part of the SYFILES system object.
- SFP=<filename> - Indicates whether the specified file (*filename*) is protected by Windows System File Protection. For example:  

```
SFP=C:\WINNT\SYSTEM32\MSVCRT.DLL
```
- SYSTEMOBJECT - Windows system object information.
- CLUSTER - Windows 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 225 for more information.

## Supported Clients

This option is valid for all Windows 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 dsmpoptfile 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 Windows 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
```



---

## 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 Windows 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 c:\Windows\Program  
Files\*.exe
```

---

## 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 Windows 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 d:\
```

---

## 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 Windows 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 c:\*
```

---

## Frsprimaryrestore

The *frsprimaryrestore* option specifies whether Tivoli Storage Manager allows the primary restoration of the SYSVOL system object if all replication partners have been lost. For example, in a disaster recovery situation when all systems containing the SYSVOL system object are lost, it is necessary to restore the first SYSVOL in this manner. In this case, set the *frsprimaryrestore* option to *yes* in your client options file (dsm.opt). If you are restoring a single system into an existing active directory tree, this option is not necessary.

### Supported Clients

This option is valid for Windows Server 2003 clients only.

### Options File

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

### Syntax



### Parameters

- Yes* Specifies that Tivoli Storage Manager allows the primary restoration of the SYSVOL system object if all replication partners have been lost.
- No* Specifies that Tivoli Storage Manager *does not* allow the primary restoration of the SYSVOL system object if all replication partners have been lost. This is the default.

### Examples

**Options file:**  
frsprimaryrestore yes

**Command line:**  
Does not apply.

---

## 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 Windows 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=c:\dir1\filelist1 -groupname=group1  
-virtualfsname=\virtfs -mode=full
```

---

## 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 Windows clients.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the **General** tab, **Return to tree window after function completed** check box of the Preferences editor.

### Syntax



### Parameters

- No* Returns you to the Tivoli Storage Manager main window after a successful operation completes. This is the default.
- Yes* Returns you to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.

### Examples

**Options file:**

```
guitreeviewafterbackup yes
```

**Command line:**

Does not apply.

---

## Httpport

The *httpport* option specifies a TCP/IP port address for the Web client.

### Supported Clients

This option is valid for all Windows clients. 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 **Web Client** tab → **HTTP Port** field of the Preferences editor.

### Syntax

▶▶—HTTPport— *port\_address* —————▶▶

### Parameters

*port\_address*

Specifies the TCP/IP port address that is used to communicate with the Web client. The range of values is 1000 through 32767; the default is 1581.

### Examples

**Options file:**

```
httpport 1502
```

**Command line:**

```
-httpport=1502
```

---

## 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 Windows 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 -ifnewer d:\logs\*.log
```



---

## Imagegapsize

Use the *imagegapsize* option with the **backup image** command, in the `dsm.opt` file, or with the *include.image* option to specify the minimum size of empty regions on a volume that you want to skip during backup. For example, if you specify a gap size of 10, this means that an empty region on the disk that is larger than 10k in size will not be backed up. Gaps that are exactly 10k will be backed up. A smaller image gap size means less data will need to be transferred, but with potentially decreased throughput. A larger image gap size results in more data being transferred, but with potentially better throughput. Use this option for LAN-based and LAN-free image backup.

Place the *include.image* statement containing the *imagegapsize* value in your `dsm.opt` file.

## Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option.

## Options File

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

## Syntax

►—IMAGEGapsize— *size* —————►

## Parameters

*size*

Specifies the minimum size of empty regions in a formatted logical volume that should be skipped during an image backup. You can specify k (kilobytes) m (megabytes) or g (gigabytes) qualifiers with the value. Without a qualifier, the value is interpreted in kilobytes. Valid values are 0 through 4294967295. If you specify a value of 0, all blocks, including unused blocks at the end of the volume, will be backed up. If you specify any value other than 0, unused blocks at the end of the volume will *not* be backed up. For LAN-based and LAN-free image backup the default value is 32k.

**Note:** Because of operating system limitations, use this option for NTFS file systems only. If you specify an `imagegapsize` that is greater than 0 for a file system other than NTFS, the error log receives a warning message.

## Examples

**Options file:**

```
imagegapsize 1m
```

**Include-exclude list example:** `include.image h: MYMC imagegapsize=1m`

**Command line:**

```
-imagegapsize=64k
```

---

## Imagetofile

Use the *imagetofile* option with the **restore image** command to specify that you want to restore the source image to a file. You might need to restore the image to a file if bad sectors are present on the target volume, or if you want to manipulate the image data. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume.

### Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option.

### Syntax

▶▶—IMAGETOfile—————▶▶

### Parameters

There are no parameters for this option.

### Examples

**Command line:**

```
dsmc restore image d: e:\diskD.img -imagetofile
```

---

## 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 systemobject**
- **query systemstate**
- **query was**
- **restore**
- **restore group**
- **restore image**
- **restore nas**
- **restore systemstate**
- **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 Windows 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 -inactive c:\id\projecta\ -pick
```

---

## Incl excl

The *incl excl* option specifies the path and file name of an include-exclude options file.

Multiple *incl excl* statements are permitted. However, you must specify this option for each include-exclude file.

Ensure that you store your include-exclude options file in a directory to which all users have read access.

When processing occurs, the include-exclude statements within the include-exclude file are placed in the list position occupied by the *incl excl* option, in the same order, and processed accordingly.

### Considerations for Unicode-enabled clients

The include-exclude file can be in Unicode or non-Unicode format. If you specify a non-Unicode include-exclude file, that file must be in the same code page that the client is running. For example, a non-Unicode include-exclude file on an English Windows system cannot contain Japanese characters.

A Unicode include-exclude file provides the following benefits:

- Names with characters from another code page no longer have to be wildcarded.
- File names and directories from any code page can be *fully specified* for the Unicode-enabled client to process.

To create an include-exclude file in Unicode format, perform the following steps:

1. Open Notepad.
2. Enter your *include* and *exclude* statements using the appropriate include-exclude options in “Using include-exclude options” on page 32. You might need to copy file names with characters from other code pages using Microsoft Windows Explorer.
3. Click **File** and then click **Save As**. The **Save As** window displays.
4. Select the **Save as Unicode** check box, specify the file and target directory, and then save the file.
5. Place an *incl excl* option specifying the include-exclude file you just created in your client options file (dsm.opt).
6. Restart the Tivoli Storage Manager client.

For more information about creating an include-exclude options file, see “Creating an include-exclude list (optional)” on page 31.

## Supported Clients

This option is valid for all Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Include-Exclude** tab of the Preferences editor.

## Syntax

▶▶—INCLExcl— *filespec* —▶▶

## Parameters

*filespec*

Specifies the path and file name of *one* include-exclude options file.

## Examples

### Options file:

```
incl excl c:\dsm\backup.excl
```

### Command line:

Does not apply.

---

## Include options

The include options specify any of the following:

- Objects within a broad group of excluded objects that you want to include for backup, archive, and image 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.
- Files that are included for backup processing that you also want to include for adaptive subfile backup 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.

You can include remotely accessed files by specifying Universal Naming Convention (UNC) names in your include statement. See “Excluding files with UNC names” on page 37 for example statements using UNC file names.

Remember that Tivoli Storage Manager compares the files it processes against the patterns specified in the include-exclude statements, reading from the bottom to the top of the options file.

### Notes:

1. The *exclude.dir* statement overrides all include statements that match the pattern.
2. The *include* statements are not case sensitive.
3. The server can also define these options with the *inlexcl* option.

## Supported Clients

This option is valid for all Windows clients.

## Options File

Place these options in the client options file (dsm.opt). You can set these options on the **Include-Exclude** tab, **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 c:\test\*\ archmc
include.backup c:\test\*
```

#### **include.archive**

Includes files or assigns management classes for archive processing.

#### **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**

If open file support has been configured (see “Configuring Open File Support (OFS)” on page 29), Tivoli Storage Manager performs a snapshot backup or archive of files that are locked (or “in use”) by other applications. The snapshot allows the backup to be taken from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to the file system are not included in the backup. You can set the *snapshotproviderfs* parameter of the *include.fs* option to *none* to specify which drives do not use open file support.

To control an open file support operation with LVSA, you can specify these additional options in your *dsm.opt* file or as values of the *include.fs* option: *snapshotproviderfs*, *presnapshotcmd*, *postsnapshotcmd*. Also, if you selected the LVSA as your snapshot provider, the following additional options can be set: *snapshotcachelocation*, *snapshotcachesize*, *snapshotfsidleretries*, *snapshotfsidlewait*. See Chapter 9, “Using processing options,” on page 165 for more information about these options.

To control how Tivoli Storage Manager processes your file space for incremental backup, you can specify these additional options in 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:**

```
include.fs d: memoryefficientbackup=diskcachem
diskcachelocation=e:\temp
include.fs e: memoryefficientbackup=diskcachem
diskcachelocation=c:\temp
```

#### **include.fs.nas**

Use the *include.fs.nas* option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system

image backup, using the *toc* option with the *include.fs.nas* option in your client options file (dsm.opt). See “Toc” on page 399 for more information.

#### **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.

By default, Tivoli Storage Manager performs an offline image backup. To enable and control an online image operation, you can specify these options in your dsm.opt file or as values of the *include.image* option:

*snapshotproviderimage*, *presnapshotcmd*, *postsnapshotcmd*. Also, if you selected the LVSA as your snapshot provider, the following additional options can be set: *snapshotcachelocation*, *snapshotcachesize*, *snapshotfsidleretries*, *snapshotfsidlewait*

#### **include.subfile**

Includes files for adaptive subfile backup processing. This option does not apply to archive processing. This option is valid for all Windows clients.

#### **include.systemobject**

Assigns management classes for backup of Windows XP system objects. By default, Tivoli Storage Manager binds all system objects to the default management class. You cannot use this option to bind individual systemobject components to a different management class. You cannot use this option to include or exclude a system object from processing. This option is valid for Windows XP *only*.

**Note:** Other include-exclude statements do not affect system object processing. It is unnecessary to explicitly include the Windows Registry staging directory `include c:\adsm.sys\...\*` for back up, to ensure that the Windows Registry is backed up properly. If you accidentally exclude a directory that is critical to a system object backup, the system object backup is not affected.

#### **include.systemstate**

Assigns management classes for back up of the system state. The default is to bind the system state object to the default management class.

## **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 options file (dsm.opt) 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 file system or raw logical volume.



**Note:** When using *include.systemobject* and *include.systemstate* the only valid pattern is ALL (all types of system objects and system state objects). By default, Tivoli Storage Manager binds all objects to the default management class.

*mgmtclassname*

Specifies the name of the management class to assign to the objects. If a management class is not specified, the default management class is used.

Use the following command to bind a management class to a group:

```
include <virtualfname>\<groupname> <mgmt class name>
```

Here is a sample command:

```
include myvfs\mygroup mymgmtclass
```

## Examples

### Options file:

```
include c:\proj\text\devel.*
include c:\proj\text\* textfiles
include ?:* managall
include WAS_ND_NDNODE mgmtclass
include WAS_APPNODE mgmtclass
include.backup c:\win98\system\* mybackupclass
include.archive c:\win98\system\* myarchiveclass
include.encrypt c:\win98\proj\gordon\*
include.compress c:\test\file.txt
include.subfile c:\test\file.txt
include.image h: MGMTCLASSNAME
    snapshotproviderimage=lvsa
    snapshotcachesize=40
include.image x:
    snapshotproviderimage=none
include.image y:
    snapshotproviderimage=vss
include.image z: MGMTCLASSNAME
    snapshotproviderimage=none
include.fs.nas netappsj1/vol/vol0 homemgmtclass
include.fs c:
    snapshotproviderfs=lvsa snapshotcachesize=40
include.systemobject ALL mgmtc2
include.systemstate ALL mgmtc3
```

### Command line:

Does not apply.

## Compression, encryption, and adaptive subfile backup processing

If you want to include specific files or groups of files for compression, encryption, and adaptive subfile backup 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 222 for more information.
- You must set the *subfilebackup* option to *yes* to enable include-exclude adaptive subfile backup processing. If you do not specify the *subfilebackup* option or you set the *subfilebackup* option to *no*, Tivoli Storage Manager does not perform adaptive subfile backup processing. See “Subfilebackup” on page 383 for more information.

- Tivoli Storage Manager processes *exclude.dir* and other include-exclude statements first. Tivoli Storage Manager then considers any *include.compression*, *include.encrypt* and *include.subfile* statements. For example, consider the following include-exclude list:

```
exclude c:\test\file.txt
include.compression c:\test\file.txt
include.encrypt c:\test\file.txt
include.subfile c:\test\file.txt
```

Tivoli Storage Manager examines the `exclude c:\test\file.txt` statement first and determines that `c:\test\file.txt` is excluded from backup processing and is, therefore, not a candidate for compression, encryption, or adaptive subfile backup processing.

- Include-exclude compression and encryption processing is valid for backup and archive processing *only*.
- Include-exclude adaptive subfile backup processing is valid for backup and restore processing *only*.
- As with other include-exclude statements, you can use the *incl excl* option to specify a file that is in Unicode format, which contains *include.compression*, *include.encrypt* and *include.subfile* statements specifying Unicode files. See “Incl excl” on page 278 for more information.

|  
|  
|  
|

---

## 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.
- NAS file system designations on the command line require brace delimiters {} around the file system names, such as: {/vol/vol0}. Do not use brace delimiters in the option file.

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 399 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 file systems

You can use the *include* option in your client options file (dsm.opt) to assign a management class to a WebSphere Application Server 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)” on page 31 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 NTFS security information, 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 Windows 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.

## Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option.

## Syntax

►►—INCREmental—————►►

## Examples

**Command line:**

```
res i d: -incremental
```

---

## Incrthreshold

The *incrthreshold* option specifies the threshold value for the number of directories in any journaled file space that might have active objects on the server, but no equivalent object on the workstation.

When a Windows client deletes a file or directory with a long name, it sometimes reports this using a compressed name. After the object is deleted, the compressed name might be reused and the deletion notice can no longer identify a unique object. During a journaled incremental backup of a file space, this can result in the *no active version* response from the server resulting in an unsuccessful expire for an object.

The *incrthreshold* option allows you to specify what to do when this condition arises:

- If you set the *incrthreshold* option to 0 (the default), Tivoli Storage Manager takes no action. The primary consequence is that, during a restore of such a directory, these objects might be inadvertently restored. When the next non-journaled incremental is run on this directory, Tivoli Storage Manager expires all objects in the directory that exist on the server but not on the workstation.
- If you specify a value greater than zero, Tivoli Storage Manager saves an object's directory name in the journal during journaled backups. During a full file space journaled incremental backup, if the number of directories in the file space is greater than or equal to this value, a full incremental backup of each directory occurs. This takes place automatically after completion of the journaled backup and does not require entry of another command.
- If you set the *incrthreshold* option to 1, Tivoli Storage Manager performs a full incremental backup of these directories whenever a *no active version* response is received during a full file space journaled incremental backup.

See "Incremental" on page 475 for more information about journaled backups.

## Supported Clients

This option is for all Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Backup** tab → **Threshold for non-journal incremental backup** field of the Preferences editor.

## Syntax

►►—INCRThreshoLd—*numberdirectories*—◄◄

## Parameters

*numberdirectories*

Specifies the threshold value for the number of directories in any journaled file space that might contain active files that should be expired. When this threshold is reached during a full file space journaled incremental, the client

initiates an incremental backup on each such directory at the completion of the journaled backup. The range of values is 0 through 2,000,000,000; the default is 0.

## Examples

**Options file:**

```
incrthreshold 1
```

**Command line:**

```
-increthreshold=1
```

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

---

## Journalpipe

The *journalpipe* option specifies the pipe name of a journal daemon session manager to which the backup clients will attach.

See Appendix C, “Journal service configuration,” on page 601 for more information about using the *journalpipe* option with the **dsmcutil** command.

### Supported Clients

This option is for all Windows clients.

### Options File

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

```
JournalPipe \\.\pipe\jnlService1
```

You can also specify this option in the journal configuration file (tsmjbbd.ini).

```
[JournalSettings]  
JournalPipe=\\.\pipe\jnlService1
```



---

## Lanfrecommmethod

The *lanfrecommmethod* 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 *lanfrecommmethod TCPip* in the client options file (dsm.opt).

### Supported Clients

This option is valid for all Windows clients.

### Options File

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

### Syntax

▶▶—LANFREECommmethod— *commmethod* —————▶▶▶

### Parameters

*commmethod*

Specifies the supported protocol for your Tivoli Storage Manager client:

*TCPip*

The Transmission Control Protocol/Internet Protocol (TCP/IP) communication method.

Use the *lanfreetcpport* option to specify the TCP/IP port number where the Storage Agent is listening. See “Lanfreetcpport” on page 294 for more information.

*V6Tcpip*

Indicates that either TCP/IP Version 4 or Version 6 should be used, depending on the system configuration and results of a domain name service lookup. The only time this is not true is when ‘dsmc schedule’ is used *and schedmode* is prompt. A valid DNS environment must be available.

*NAMedpipes*

The interprocess communication method that permits message data streams to pass between a client and a server. This is the default. Do not specify the *lanfreetcpport* option if you want to use the NAMedpipes communication method for LAN-free communication.

*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.

### Examples

**Options file:**

```
lanfrecommmethod tcp
```

```
Use only TCP/IP Version 4
```

```
lanfrecommmethod V6Tcpip
```

| Use both TCP/IP Version 4 or Version 6, depending on how the system is  
| configured and the results of a domain name service lookup.

| **Command line:**  
| -lanfreec=tcp  
| -lanfreec=V6Tcpip

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

---

## Lanfreeshmport

Use the *lanfreeshmport* option when *lanfreecommmethod*=*SHAREdmem* is specified for communication between the Tivoli Storage Manager client and Storage Agent. This enables processing between the client and the SAN-attached storage device. See “Lanfreecommmethod” on page 291 for more information about the *lanfreecommmethod* option.

### Supported Clients

This option is valid for all Windows clients *only*.

### Options File

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

### 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.

---

## Lanfreetcpport

The *lanfreetcpport* 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 *lanfreetcpport* option if you want to use the NAMedpipes communication method for LAN-free communication. See “Lanfreecommmethod” on page 291 for more information about the *lanfreecommmethod* option.

### Supported Clients

This option is valid for all Windows clients.

### Options File

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

### Syntax

▶▶—LANFREETCPport— *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 *lanfreetcpport* value must match Storage Agent *tcpport* value for communications with the Storage Agent (virtual server). The client *tcpport* value must match the server *tcpport* value for communications with the actual server.

### Examples

**Options file:**

```
lanfreetcp 1520
```

**Command line:**

```
-lanfreetcp=1520
```

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

---

## Lanfreetcpserveraddress

The *lanfreetcpserveraddress* option specifies the TCP/IP address for a Tivoli Storage Manager Storage Agent. Use this option when you specify *lanfreecommmethod=TCPIP* or *V6Tcip* 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 all Windows clients, except the 64-bit Windows Server 2003.

### Options File

Place this option in the client system options file.

### Syntax

▶▶—LANFREETCPServeraddress— *stagent\_address*—————▶▶

### Parameters

| *stagent\_address*  
| Specifies a 1 to 64 character TCP/IP address for a server. Specify a TCP/IP  
| domain name or a numeric IP address. The numeric IP address can be either a  
| TCP/IP v4 or TCP/IP v6 address. The default value is 127.0.0.1 (localhost).

### Examples

#### Options file:

```
LANFREETCPServeraddress stagent.bayamon.ibm.com
```

```
LANFREETCPServeraddress 192.168.1.50
```

#### Command line:

Does not apply.

---

## Language

The *language* option specifies the national language in which to present client messages.

You can use US English (ENU) with all clients.

The Tivoli Storage Manager client automatically detects the language of the system locale and displays Tivoli Storage Manager for that language. For example, a supported operating system on a Windows system running a French locale will display Tivoli Storage Manager in French by default, without specifying the *language* option; provided that the French language pack is installed. If Tivoli Storage Manager cannot load the French message catalog, it will default to the US English language pack. If the client is running on an unsupported language/locale combination, such as French/Canada or Spanish/Mexico, Tivoli Storage Manager defaults to US English. You can override the default language by specifying the *language* option.

**Note:** The *language* option does not affect the Web client. The Web client displays in the language associated with the locale of the browser. If the browser is running in a locale that Tivoli Storage Manager does not support, the Web client displays in US English.

## Supported Clients

This option is valid for all Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Regional Settings** tab, **Language** drop-down list of the Preferences editor.

## Syntax

►►—LANGUage— *language* —————►►

## Parameters

*language*

Specifies the language you want to use. The available languages include:

- ENU (English, United States)
- PTB (Brazilian Portuguese)
- CHS (Chinese, Simplified)
- CHT (Chinese, Traditional)
- FRA (Standard French)
- DEU (Standard German)
- ITA (Standard Italian)
- JPN (Japanese)
- KOR (Korean)
- ESP (Standard Spanish)
- CSY (Czech)
- HUN (Hungarian)
- PLK (Polish)
- RUS (Russian)

## Examples

**Options file:**

language enu

**Command line:**

Does not apply.

---

## Latest

Use the *latest* option with the following commands to restore the most recent backup version of a file, even if the backup is inactive:

- **restore**
- **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 Windows 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 c:\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` 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 529 for more information.

### Supported Clients

This option is valid for all Windows clients.

### Options File

Place this option in the `dsm.opt` 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 Windows 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 clients.

**Note:** 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/>

## Examples

### Command line:

```
restore backupset
  backupsetname=monthly_financial_data.1234
  -location=server

restore backupset backupsetname=c:
  \BudgetData\weekly_budget_data.ost
  -location=file

restore backupset
  backupsetname=\\.TAPE1 -location=tape

restore f:\dir1\file.1 -backupsetname=
  c:\BudgetData\monthly_budget_data.ost
  -location=file
```

```
query backupset backupsetname=  
    weekly_financial_data.1234 -location=server  
  
query backup f:\* -subdir=yes  
    -backupsetname=\\.\TAPE1 -location=tape
```

---

## Managedservices

The *managedservices* option specifies whether the Tivoli Storage Manager client acceptor service manages the scheduler, the Web client, or both.

See “Configuring the client scheduler” on page 21 for instructions to set up the client acceptor daemon to manage the scheduler.

**Attention:** You cannot use the *dsmcad* for scheduling when you set the *sessioninitiation* option to *serveronly*. Refer to “Sessioninitiation” on page 361 for more information.

The client acceptor daemon serves as an external timer for the scheduler. When the scheduler is started, it queries the server for the next scheduled event. The event is either executed immediately or the scheduler exits. The client acceptor daemon restarts the scheduler when it is time to execute the scheduled event.

### Notes:

1. If you set the *schedmode* option to *prompt*, the server prompts the client acceptor daemon when it is time to run the schedule. The scheduler will connect and disconnect to the server when the client acceptor daemon is first started.

The *dsmc schedule* command cannot be used when both *SCHEDMODE* *PRompt* and *COMMMethod* *V6Tcpip* are specified.

2. Set the *passwordaccess* option to *generate* in your client options file (*dsm.opt*) and generate a password, so Tivoli Storage Manager can manage your password automatically. See “Passwordaccess” on page 320 for more information.

Using the client acceptor daemon to manage the scheduler service can provide the following benefits:

- Memory retention problems that can occur when using traditional methods of running the scheduler are resolved. Using the client acceptor daemon to manage the scheduler requires very little memory between scheduled operations.
- The client acceptor daemon can manage both the scheduler program and the Web client, reducing the number of background processes on your workstation.
- By default, if you do not specify the *managedservices* option, the client acceptor daemon manages the Web client only.

## Supported Clients

This option is valid for all Windows clients. 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 **Web Client** tab of the Preferences editor.

## Syntax

►►—MANAGEDServices—mode—webclient  
schedule—►►

## Parameters

*mode*

Specifies whether the client acceptor daemon manages the scheduler, the Web client, or both.

*webclient*

Specifies that the client acceptor daemon manages the Web client. This is the default.

*schedule*

Specifies that the client acceptor daemon manages the scheduler.

## Examples

### Options file:

The following are examples of how you might specify the *managedservices* option in your client options file (dsm.opt).

**Task** Specify that the client acceptor daemon manages the Web client *only*.

```
managedservices webclient
```

**Task** Specify that the client acceptor daemon manages the scheduler *only*.

```
managedservices schedule
```

**Task** Specify that the client acceptor daemon manages both the Web client and the scheduler.

```
managedservices schedule webclient
```

**Note:** The order in which these values are specified is not important.

### Command line:

Does not apply.

---

## Maxcmdretries

The *maxcmdretries* option specifies the maximum number of times the client scheduler (on your workstation) attempts to process a scheduled command that fails. The command retry starts *only if* the client scheduler has not yet backed up a file, never connected to the server, or failed before backing up a file. This option is only used when the scheduler is running.

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 Windows clients. 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 **Scheduler** tab, **Maximum command retries** field of the Preferences editor.

### Syntax

▶▶—MAXCMDRetries— *maxcmdretries* —————▶▶

### Parameters

*maxcmdretries*

Specifies the number of times the client scheduler can attempt to process a scheduled command that fails. The range of values is zero through 9999; the default is 2.

### Examples

**Options file:**

```
maxcmdr 4
```

**Command line:**

```
-maxcmdretries=4
```

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

---

## Memoryefficientbackup

The *memoryefficientbackup* option specifies the memory-conserving algorithm to use for processing full file space backups. One method backs up one directory at a time, using less memory. The other method uses much less memory, but requires more disk space. Use the *memoryefficientbackup* option with the **incremental** command when your workstation is memory constrained. You can also use this option as a parameter to the *include.fs* option in order to select the algorithm Tivoli Storage Manager uses on a per-file space basis. See "Include options" on page 280 for more details. If the *memoryefficientbackup* option appears in the option file, its value will be used for all file systems not represented by an *include.fs* option containing the *memoryefficientbackup* parameter. If a directory-level incremental is being performed, and *memoryefficientbackup=diskcachemethod* is present, *memoryefficientbackup=yes* will be used.

Use *memoryefficientbackup=diskcachemethod* for any file space that has too many files for Tivoli Storage Manager to complete the incremental backup with either the default setting, *memoryefficientbackup=no*, or with *memoryefficientbackup=yes*. The disk cache file created by the initial disk cache incremental backup can require up to 5 gigabytes of disk space for each million files or directories being backed up.

The disk cache file created by subsequent disk cache incremental backups should require considerably less disk space. The actual amount of disk space required for the disk cache file created by subsequent disk cache incremental backups depends on the number of files and directories included in the backup and on the length of the longest file or directory name which is backed up. Since the disk cache file can become very large, ensure that large file support is enabled on the file system that is being used for the disk cache file.


## 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*), or on the initial command line. You can also set this option on the **Performance Tuning** tab, 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 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 all Windows clients *only*. The Tivoli Storage Manager client API does not support this option.

## Syntax


*For image backup of client file systems ( Windows XP)*



*For image backup of NAS file systems (all Windows clients)*



*For group and Websphere Application Server backups (all Windows clients)*



## 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, Websphere Application Server, or group objects. This is the default for Websphere Application Server and group backups.

### *differential*

This is the default for NAS objects. Specifies that you want to perform a NAS, Websphere Application Server, or group backup of files that changed since the last full backup. If there is no copy of a full image stored on the Tivoli Storage Manager server, a full backup occurs. If a full image exists, whether it is restorable, or expired and being maintained because of dependent differential images, specifying **MODE =differential** sends a differential image backup. If a full image is sent during a *differential* backup, it is reflected as a full image using the **query nasbackup** server command.

A full image can be eligible for expiration based on versioning or retention (verexists retextra), but still be maintained on the Tivoli Storage Manager server to allow for restoring dependent differential images. A full image that is eligible for expiration *cannot* be selected for restore, so it is not displayed using the **query nasbackup** server command. The differential image backups that depend on an "expired" full image can be restored.

## Examples

**Task** Perform the NAS image backup of the entire file system.

**Command:** `dsmc backup nas -mode=full -nasnodename=nas1 {/vol/vol0} {/vol/vol1}`

**Task** Back up the c: drive using an image incremental backup that backs up only new and changed files after the last full image backup.

**Command:** `dsmc backup image c: -mode=incremental`

**Task** Perform a full backup of all the files in filelist c:\dir1\filelist1 to the virtual file space name \virtfs containing the group leader c:\group1 file.

**Command:**

`backup group -filelist=c:\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 all Windows clients.

## 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}
```

---

## Namedpipename

The *namedpipename* option specifies the name of a named pipe to use for communications between a client and a server on the same Windows server domain.

### Supported Clients

This option is valid for all Windows clients.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Communication** tab of the Preferences editor.

### Syntax

▶▶—`Namedpipename`— *name* —————▶▶

### Parameters

*name*

The name of a named pipe. The default is `\\.\pipe\Server1`.

### Examples

**Options file:**

```
namedpipename \\.\pipe\dsmser1
```

**Command line:**

```
-namedpipename=\\.\pipe\dsmser1
```

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

---

## 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 options file (dsm.opt).

You can override the default value in the dsm.opt file by entering a different value on the command line. If you do not specify the *nasnodename* option in the dsm.opt file, you *must* specify this option on the command line when processing NAS file systems.

You can use the *nasnodename* option with the following commands:

- **backup nas**
- **delete filespace**
- **query backup**
- **query filespace**
- **restore nas**

## Supported Clients

This option is valid for the all Windows clients. 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 **General** tab of the Preferences editor.

## Syntax

▶▶—NASNodename— *nodename* —————▶▶

## Parameters

*nodename*

Specifies the node name for the NAS file server.

## Examples

**Options file:**

```
nasnodename nas2
```

**Command line:**

```
-nasnodename=nas2
```

---

## Nodename

Use the *nodename* option in your client options file (dsm.opt) 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 410 for more information.

If you are working from a different workstation, you can use the *nodename* option even if the *passwordaccess* option is set to *generate*, however, the password will be stored in the Windows Registry. To prevent this, use the *virtualnodename* option instead of *nodename*.

The node name is not necessarily the TCP/IP host name.

- In the absence of a *nodename* entry in the dsm.opt file, or a *virtualnodename* entry in the client 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.opt file, the *nodename* entry overrides the name that the **hostname** command returns.
- If a *virtualnodename* entry exists in the client 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 Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **General** tab, **Node Name** field of the Preferences editor.

## Syntax

►►—NODename— *nodename* —————►►

## Parameters

*nodename*

Specifies a 1 to 64 character node name for which you want to request Tivoli Storage Manager services. The default is the value returned with the **hostname** command. If you set the *clusternode* option to *yes*, the default is the cluster name.

Not specifying a node name will permit the node name to default to the host name of the workstation

## Examples

**Options file:**

```
nodename cougar
```

**Command line:**

```
-nodename=cougar
```

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

---

## 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).
- Attribute changes to an object require a back up of the entire object.

For these reasons, you might want to use the *nojournal* option periodically to perform a traditional full incremental backup.

See “Journal-based backup” on page 476 for more information about journal-based incremental backup.

### Supported Clients

This option is valid for all Windows clients, except for clients running on Windows Server 2003 for Itanium-based Systems. The Tivoli Storage Manager client API does not support this option.

### Syntax

►►—NOJournal—◄◄

### Parameters

There are no parameters for this option.

### Examples

**Command line:**

```
dsmc incr c: -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 Windows 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 c:\home\project\*
```

---

## Numberformat

The *numberformat* option specifies the format you want to use to display numbers.

Use this option if you want to change the default number format for the language of the message repository you are using.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time the client is called. Consult the documentation on your local system for details about setting up your locale definition.

**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**
- **query systemobject**
- **query systemstate**
- **restore**
- **restore image**
- **restore nas**
- **retrieve**
- **restore registry**
- **restore systemobject**
- **restore systemstate**
- **set event**

## Supported Clients

This option is valid for all Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Regional Settings** tab, **Number Format** field of the Preferences editor.

## Syntax

▶—NUMBERformat— *number* —▶

## Parameters

*number*

Displays numbers using any one of the following formats. Specify the number (0–6) that corresponds to the number format you want to use.

0 Use the locale-specified date format. This is the default.

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

## 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 options file you want to use when you start a Tivoli Storage Manager session.

### Supported Clients

This option is valid for all Windows 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.

### Examples

**Command line:**

```
dsmc query session -optfile=myopts.opt
```

**client acceptor daemon:**

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 Windows clients.

## Options File

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

## Syntax

►►—PASsword— *password* —————►►

## Parameters

### *password*

Specifies a 1 to 63 character password. A password is not case-sensitive. Valid characters include:

#### Characters

	Description
A–Z	Any letter, A through Z, uppercase or lowercase
0–9	Any number, 0 through 9
+	Plus
.	Period
_	Underscore
-	Hyphen
&	Ampersand

## Examples

### Options file:

```
password secretword
```

### Command line:

```
-password=secretword
```

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

---

## Passwordaccess

The *passwordaccess* option specifies whether you want to generate your password automatically or set as a user prompt. Your administrator can require a password for your client node by enabling the authentication feature. Ask your administrator if a password is required for your client node.

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.

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 Web client.
- When performing NAS operations.
- When using Tivoli Continuous Data Protection for Files.

## Supported Clients

This option is valid for all Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Authorization** tab, **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.

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. The new password is randomly generated by the Tivoli Storage Manager client

|  
|  
|

A password prompt is displayed when registering a workstation with a server using open registration or if your administrator changes your password manually.

## Examples

**Options file:**

passwordaccess generate

**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 asr**
- **restore group**
- **restore image**
- **restore nas**
- **restore was**
- **retrieve**

## Supported Clients

This option is valid for all Windows 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 c:\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 systemstate**
- **query was**
- **restore**
- **restore group**
- **restore was**
- **restore image**
- **restore nas**
- **restore systemstate**
- All query and restore system object commands

When *pitdate* is used, the *inactive* and *latest* options are implicit.

## Supported Clients

This option is valid for all Windows 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 -pitdate=08/01/2003 c:\myfiles\
```

---

## 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**
- **query systemstate**
- **restore**
- **restore image**
- **restore nas**
- **restore systemstate**
- **restore was**
- All query and restore system object commands

## Supported Clients

This option is valid for all Windows 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 query backup -pitt=06:00:00 -pitd=08/01/2003 c:\myfiles\
```

---

## Postschedulecmd/Postnschedulecmd

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 Windows clients. 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 **Scheduler** tab, **Schedule Command** button of the Preferences editor.

## Syntax



## 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.

Specify the command string just as you would enter it from the operating system command prompt. If the command string contains any blank spaces, enclose the command string in single quotes. For example:

```
'net stop someservice'
```

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:

```
posts startdb.cmd
posts 'rename c:\myapp\logfile.log logfile.new'
posts 'net start "simple service"'
posts 'rename "c:\myapp\log file.log" "log file.new"'
posts '"C:\Program Files\MyTools\runreport.bat"
log1.in log2.in'
```

### 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 a command or script after a snapshot is started. This option can be used in conjunction with the *presnapshotcmd* option to allow you to quiesce an application while a snapshot is created, and then to restart that application after the snapshot is started. This option is only valid if OFS or online image backup has been configured.

For an online image backup, use this option with the **backup image** command, the *include.image* option, or in the *dsm.opt* file.

For open file support operations, use the *postsnapshotcmd* option in an *include.fs* statement or in your client options file (*dsm.opt*). See “Include options” on page 280 for more information.

If the *postsnapshotcmd* fails the operation continues, but appropriate warnings are logged.

## Supported Clients

This option is valid for all Windows clients. 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 also set this option on the **Image-Snapshot** tab of the Preferences editor.

## Syntax

►—POSTSNAPSHOTCMD— "*cmdstring*" —◄

## Parameters

*"cmdstring"*

Specifies the quiesce command to process.

Use a blank, or null, string for *"cmdstring"* if you want to prevent any commands from running that the administrator uses for *postsnapshotcmd*. If you specify a blank or null string, it prevents the administrator from using a command on this option. If your administrator uses a blank or null string on the *postsnapshotcmd* option, you cannot run a post-snapshot command. Use the *srprepresentsnapdisabled* option to prevent the Tivoli Storage Manager server administrator from executing operating system commands on the client system. See “Srvprepresentscheddisabled” on page 378 for details on this option.

If the command string contains blanks, enclose the command string in quotation marks:

```
"resume database myDb"
```

If you placed quotation marks within the command string, then enclose the entire command string in single quotes:

```
'resume database "myDb"'
```

## Examples

**Options file:**

```
postsnapshotcmd "restart application"
```

The command string is a valid command for restarting your application.

**Command line:**

```
backup image -postsnapshotcmd="restart application"
```

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

---

## Preschedulecmd/Prenschedulecmd

The *preschedulecmd* option specifies a command that the client program processes before it runs a schedule. The client program waits for the command to complete before it starts the schedule.

If you do not want it to wait, specify *prenschedulecmd*.

### 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 Windows clients. 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 **Scheduler** tab, **Schedule Command** button of the Preferences editor.

## Syntax

►► — PRESchedulecmd — "cmdstring" ————— ◀◀  
    └─ PRENSchedulecmd ─┘

## Parameters

"cmdstring"

Specifies the command to process. Use only one *preschedulecmd* option. You can enter a command to be executed before a schedule using this option.

Specify the command string just as you would enter it from the operating system command prompt. If the command string contains any blank spaces, enclose the command string in single quotes. For example:

```
'net stop someservice'
```

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:

```
pres stopdb.cmd
pres 'rename c:\myapp\logfile.log logfile.old'
pres 'net stop "simple service"'
pres 'rename "c:\myapp\log file.log" "log file.old"'
pres '"C:\Program Files\MyTools\runreport.bat"
log1.in log2.in'
```

### 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), 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. This option requires extra processing time during backup and archive for each file that is sent to the Tivoli Storage Manager server, so it should only be enabled when necessary. If you have enabled open file support, the last access date for files will always be preserved regardless of the setting for *preservelastaccessdate*. When using open file support, leave this option disabled.

Use this option with the **incremental**, **selective**, or **archive** commands.

### Notes:

1. This option only applies to files; it does not apply to directories.
2. Resetting the last access date incurs additional overhead that impact backup and archive performance. The last access date should be reset only if you are using other another application, such as a Storage Resource Management (SRM) that relies on accurate last access dates.
3. The last access date cannot be preserved on files which are write protected either by the read-only attribute or by a restrictive NTFS security permission.
4. 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 Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Backup** tab of the Preferences editor.

## Syntax

►► PRESERVELastaccessdate  No  Yes ◀◀

## Parameters

- No* Specifies that the Tivoli Storage Manager client *will not* reset the last access date of any backed up or archived files to their original value following the backup or archive operation. This is the default.

*Yes* Specifies that the Tivoli Storage Manager *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:**

```
dsmc incr c: e: f: -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 Windows 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 the server file space contains the following backup copies:

```
c:\h1\m1\file.a
c:\h1\m1\file.b
c:\h1\m1\l1\file.x
c:\h1\m1\l1\file.y
```

**This command:**

```
dsmc res backupset my.backupset.file /fs/h1/m1/ /u/ann/ -su=yes
creates a local backupset file named "my.backupset.file".
```

**Restores these directories and files:**

```
c:\ann\h1\m1\file.a
c:\ann\h1\m1\file.b
```

**This command:**

```
dsmc res c:\h1\m1\ c:\ann\ -preser=nobase.
```

**Restores these directories and files:**

```
c:\ann\file.a
c:\ann\file.b
```

**This command:**

```
dsmc res c:\h1\m1\ c:\ann\ -preser=subtree.
```

**Restores these directories and files:**

```
c:\ann\m1\file.a
c:\ann\m1\file.b
```

**This command:**

```
dsmc res c:\h1\m1\ c:\ann\ -preser=none.
```

**Restores these directories and files:**

```
c:\ann\file.a
c:\ann\file.b
```

**This command:**

```
dsmc res c:\h1\m1\ c:\ann\ -su=yes -preser=
complete
```

**Restores these directories and files:**

```
c:\ann\h1\m1\file.a
c:\ann\h1\m1\file.b
c:\ann\h1\m1\l1\file.x
c:\ann\h1\m1\l1\file.y
```

**This command:**

```
dsmc res c:\h1\m1\ c:\ann\ -su=yes -preser=nobase.
```

**Restores these directories and files:**

```
c:\ann\file.a
c:\ann\file.b
c:\ann\l1\file.x
c:\ann\l1\file.y
```

**This command:**

```
dsmc res c:\h1\m1\ c:\ann\ -su=yes -preser=subtree.
```

**Restores these directories and files:**

```
c:\ann\m1\file.a
c:\ann\m1\file.b
c:\ann\m1\l1\file.x
c:\ann\m1\l1\file.y
```

**This command:**

```
dsmc res c:\h1\m1\ c:\ann\ -su=yes -preser=none.
```

**Restores these directories and files:**

c:\ann\file.a  
c:\ann\file.b  
c:\ann\file.x  
c:\ann\file.y

**This command:**

```
dsmc res backupset c:\h1\m1\ c:\ann\ -su=yes  
-preser=nobase -loc=file
```

**Restores these directories and files:**

c:\ann\file.a  
c:\ann\file.b  
c:\ann\file.x  
c:\ann\file.y

---

## Presnapshotcmd

The *presnapshotcmd* option allows you to run a command or script before a snapshot is started. This option can be used in conjunction with the *postsnapshotcmd* option to allow you to quiesce an application while a snapshot is created, and then to restart that application after the snapshot is started. This option is only valid if OFS or online image backup has been configured.

For an online image backup, use this option with the **backup image** command, the *include.image* option, or in the dsm.opt file.

For open file support operations, use the *presnapshotcmd* option in an *include.fs* statement or in your client options file (dsm.opt). See "Include options" on page 280 for more information.

If the *presnapshotcmd* fails it is assumed that the application is not in a consistent state and the client will stop the operation and display the appropriate error message.

### Supported Clients

This option is valid for all Windows clients. 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 also this option on the **Image-Snapshot** tab of the Preferences editor.

### Syntax

►—PRESnapshotcmd— "*cmdstring*" —————►

### Parameters

*"cmdstring"*

Specifies the quiesce command to process.

Use a blank, or null, string for *"cmdstring"* if you want to prevent any commands from running that the administrator uses for *presnapshotcmd*. If you specify a blank or null string, it prevents the administrator from using a command on this option. If your administrator uses a blank or null string on the *presnapshotcmd* option, you cannot run a pre-snapshot command.

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 378 for details on this option.

If the command string contains blanks, enclose the command string in quotation marks:

```
"quiesce database myDb"
```

If you placed quotation marks within the command string, then enclose the entire command string in single quotes:

```
'resume database "myDb"'
```

## Examples

### Options file:

```
presnapshotcmd "<insert your application quiesce command here>  
application"
```

The command string is a valid command for quiescing your application.

### Command line:

```
backup image -presnapshotcmd="<insert your application quiesce command  
here> application"
```

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

---

## Queryschedperiod

The *queryschedperiod* option specifies the number of hours you want the client scheduler to wait between attempts to contact the server for scheduled work. This option applies only when you set the *schedmode* option to *polling*. This option is used only when the **scheduler** is running.

Your administrator can also set this option. If your administrator specifies a value for this option, that value overrides the value set in your client options file after your client node successfully contacts the server.

### Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Scheduler** tab, **Query Schedule Interval** field of the Preferences editor.

### Syntax

►►—QUERYSChedperiod— *hours* —————►►

### Parameters

*hours*

Specifies the number of hours the client scheduler waits between attempts to contact the server for scheduled work. The range of values is 1 through 9999; the default is 12.

### Examples

**Options file:**

```
querysch 6
```

**Command line:**

```
-queryschedperiod=8
```

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



---

## 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 Windows 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 options file (dsm.opt). You can set this option on the **Command Line** tab, **Do not display process information on screen** checkbox of the Preferences editor.

### Syntax

▶▶—QUIET—▶▶

### Parameters

There are no parameters for this option.

### Examples

**Options file:**  
quiet

**Command line:**  
-quiet

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

---

## 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 Windows clients. 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 **Restore** tab, **Action for files that already exist** section of the Preferences editor.

## Syntax



## Parameters

- | *Prompt*  
| For nonscheduled operations, you specify whether to overwrite existing files.  
| For scheduled operations, existing files are not overwritten and no prompts are  
| displayed. This is the default.
- | *All*  
| All existing files are overwritten, including read-only files. All locked files are  
| replaced when the system is rebooted. If access to a file is denied, you are  
| prompted to skip or overwrite the file. No action is taken on the file until there  
| is a response to the prompt.
- | *Yes*  
| Existing files are overwritten, *except* read-only files. For nonscheduled  
| operations, you specify whether to overwrite existing read-only files. For  
| scheduled operations, existing read-only files are not overwritten and no  
| prompts are displayed. If access to a file is denied, the file is skipped.
- | *No* Existing files are not overwritten. No prompts will be displayed.

**Note:** You can choose to replace locked files when the system is rebooted. Tivoli Storage Manager cannot perform an in-place restore of active files. However, it will stage restored versions of active files for replacement during the next reboot, except files containing named streams, sparse files, and directories. You can only restore these files if they are unlocked.

## 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.

---

## Resetarchiveattribute

Use the *resetarchiveattribute* option to specify whether Tivoli Storage Manager resets the Windows archive attribute on files that are successfully backed up to a Tivoli Storage Manager server. Tivoli Storage Manager will also reset the archive attribute during incremental backups if it is determined that there is already an active object on the Tivoli Storage Manager server. The *resetarchiveattribute* option is useful in conjunction with applications, such as the IBM Tivoli Storage Resource Manager, as a simple way to report on the backup status of files.

The Windows archive attribute is used to indicate that a file has changed since the last backup. After Tivoli Storage Manager resets the archive attribute, the Windows operating system will turn the attribute back to *ON* after the file has been modified. Tivoli Storage Manager does not use the Windows archive attribute to determine if a file is a candidate for incremental backup, but only manipulates this attribute for reporting purposes. Tivoli Storage Manager uses a much more sophisticated method to determine candidacy for incremental backup. See “Full and partial incremental backup” on page 65 for more information.

There are several applications which manipulate or examine the Windows archive attribute. Be aware of the ramifications of using the *resetarchiveattribute* option in conjunction with these products.

If you set the *resetarchiveattribute* option to *yes*, after a file has been successfully backed up to the Tivoli Storage Manager server, the Tivoli Storage Manager client resets the Windows archive attribute on the local file system:

- The Windows archive attribute will be reset during incremental and selective backups after the file has been successfully committed to the Tivoli Storage Manager database. This attribute is not reset for archive, or image operations.
- The Windows archive attribute is not reset when processing system objects or system state objects.
- The Windows archive attribute is not reset for directory entries.

In addition, in order for the local file system to reflect the current active object inventory on the Tivoli Storage Manager server, the *resetarchiveattribute* option will instruct the Tivoli Storage Manager client to reset the Windows archive attribute on the local file system if it is determined during incremental backup that a valid, active backup copy of the file already exists on the Tivoli Storage Manager server. This behavior will not be displayed in the following cases:

- Incremental backup operations which do not examine the stored client attributes on the Tivoli Storage Manager server, such as journal-based backup or incremental-by-date processing.
- Files that are not examined during an incremental backup operation because they are excluded from backup processing.

The Tivoli Storage Manager client does not guarantee the accuracy of the current setting of the Windows archive attribute. For example, if the *resetarchiveattribute* option is set to *yes* and a file examined by a reporting product indicates that the Windows archive attribute is *OFF* for a particular file, this does not necessarily mean that a valid, active backup copy of the file exists on the Tivoli Storage Manager server. Factors that could contribute to this type of situation include:

- A third-party product is manipulating the Windows archive attribute
- A file space was deleted from the Tivoli Storage Manager server
- A backup tape was lost or destroyed

There should be no significant performance degradation when using the *resetarchiveattribute* option. The *resetarchiveattribute* option does not affect restore processing.

## Supported Clients

This option is valid for all Windows clients. The server can also define this option.

## Options File

This option is valid in the client options file (*dsm.opt*) or server client options set. You can set this option on the **Backup** tab of the Preferences editor.

## Syntax



## Parameters

- Yes* Specifies that you want to reset the Windows archive attribute for files during a backup operation.
- No* Specifies that you do not want to reset the Windows archive attribute for files during a backup operation. This is the default.

## Examples

**Options file:**  
resetarchiveattribute yes

---

## Resourceutilization

Use the *resourceutilization* option in your client options file `dsm.opt` 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.

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 215 for more information.

### Regulating restore sessions

When you request a restore, the default is to use a maximum of one session. Additional restore sessions are based on the following factors:

- Resourceutilization value
- How many tapes on which the requested data is stored
- How many tape drives are available
- The maximum number of mount points allowed for the node

**Notes:**

1. If all of the files are on disk, only one session is used. There is no multi-session for a pure disk storage pool restore. However, if you are performing a restore in which the files reside on 4 tapes and some on disk, you could use up to 5 sessions during the restore.
2. The Tivoli Storage Manager server can set the maximum number of mount points a node can use on the server using the `MAXNUMMP` parameter. If the *resourceutilization* option value exceeds the value of the `MAXNUMMP` on the server for a node, the backup can fail with an *Unknown System Error* message.
3. You can get a multi-session restore from your single **restore** command, and from a single volume on the server, if that volume is device class `FILE`.

For example, if the data you want to restore is on 5 different tape volumes, the maximum number of mount points is 5 for your node, and *resourceutilization* is set to 3, then 3 sessions will be used for the restore. If you increase the *resourceutilization* setting to 5, then 5 sessions will be used for the restore. There is a 1 to 1 relationship to the number of restore sessions allowed for the *resourceutilization* setting. Multiple restore sessions are only allowed for no query restore operations.

## Considerations

The following factors can affect the throughput of multiple sessions:

- The server's ability to handle multiple client sessions. Is there sufficient memory, multiple storage volumes, and CPU cycles to increase backup throughput?
- The client's ability to drive multiple sessions (sufficient CPU, memory, etc.).
- The configuration of the client storage subsystem. File systems that are striped across multiple disks, using either software striping or RAID-5 can better handle an increase in random read requests than a single drive file system. Additionally, a single drive file system might not see performance improvement if it attempts to handle many random concurrent read requests.
- Sufficient bandwidth in the network to support the increased traffic.

Potentially undesirable aspects of running multiple sessions include:

- The client could produce multiple accounting records.
- The server might not start enough concurrent sessions. To avoid this, the server *maxsessions* parameter must be reviewed and possibly changed.
- A query node command might not summarize client activity.

## Supported Clients

This option is valid for all Windows 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 options file (*dsm.opt*). You can set this option on the **General** tab, **Resource Utilization** field of the Preferences editor.

## Syntax

►►—RESOURceutilization— *number* —————►►

## Parameters

*number*

Specifies the level of resources the Tivoli Storage Manager server and client can use during processing. The range of values that you can specify is 1 through 10.

## Examples

**Options file:**

```
resourceutilization 7
```

**Command line:**

```
-resourceutilization=7
```

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



---

## Retryperiod

The *retryperiod* option specifies the number of minutes the client scheduler waits between attempts to process a scheduled command that fails, or between unsuccessful attempts to report results to the server. Use this option only when the scheduler is running.

Your administrator can also set this option. If your administrator specifies a value for this option, that value overrides the value in your client options file after your client node successfully contacts the server.

### Supported Clients

This option is valid for all Windows clients. 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 **Scheduler** tab, **Retry period** field of the Preferences editor.

### Syntax

►►—RETRYPeriod— *minutes* —————►►

### Parameters

*minutes*

Specifies the number of minutes the client scheduler waits between attempts to contact the server, or to process a scheduled command that fails. The range of values is 1 through 9999; the default is 20.

### Examples

**Options file:**

```
retryp 10
```

**Command line:**

```
-retryperiod=10
```

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

---

## Revokeremoteaccess

The *revokeremoteaccess* option restricts an administrator with client access privilege from accessing a client workstation that is running the Web client. This option does not restrict administrators with client owner, system, or policy privilege from accessing your workstation through the Web client.


### Supported Clients

This option is valid for all Windows clients. 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 **Web Client** tab 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

---

## Runasservice


The *runasservice* option forces the client command process to continue running, even if the account that started the client logs off. Use this option with the **AT** command and the **dsmc sched** command when you schedule client command batch jobs. The *runasservice* option is *not* valid in any options file (dsm.opt or tsmasr.opt).

**Recommendation:** Use the scheduler service when running Tivoli Storage Manager services unattended. Set *runasservice=yes* only to schedule Tivoli Storage Manager commands using the Windows **AT** command. Setting *runasservice=yes* might interfere with other interactive uses of the Tivoli Storage Manager client.

## Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option.

## Syntax



## Parameters

*No* Does not force the client command process to continue running, even if the account that started the client logs off. This is the default.

*Yes*

Forces the client command process to continue running, even if the account that started the client logs off.

### Restrictions:

1. When *runasservice=yes*, the setting for the REPLACE is always overridden to the behavior of *replace=no*.
2. The option *runasservice=yes* cannot be used with *passwordaccess=prompt*.
3. Backup, archive, restore and retrieve options performed with *runasservice=yes* that encounter prompts will always fail. To avoid this problem, either save the encryption key password with *encryptkey=save*, or turn off the *runasservice* option.

## Examples

### Command line:

```
-runasservice=yes
```

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

---

## Schedcmddisabled

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 511 for more information.

### Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option.

### Options File

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

### 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

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 Windows clients.

## Options File

Place this option in the client options file (*dsm.opt*). You can set this option on the **Scheduler** tab, **Schedule Log** button of the Preferences editor.

## Syntax

▶—SCHEDLOGMAX— *size* —▶

## Parameters

*size*

Specifies the maximum size, in megabytes, for the log file. The range of values is 0 to 2047; the default is 0, which disables log file wrapping and allows the log file to grow indefinitely.

## Examples

**Options file:**

```
schedlogmax 100
```

**Command line:**

```
-schedlogmax=100
```

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

---

## Schedlogname

The *schedlogname* option specifies the path and file name where you want to store schedule log information. Use this option only when you want to store schedule log information. This option applies only when the scheduler is running. If this option is not used, the `dsmsched.log` file is created in the same directory as the `dsmerror.log` file. See “Errorlogname” on page 254 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 Windows clients. 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 **Scheduler** tab, **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.

## 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`.

## Examples

### Options file:

```
schedlogname c:\mydir\schedlog.jan
```

### Command line:

```
-schedlogn=c:\mydir\schedlog.jan
```

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

---

## Schedlogretention

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 Windows clients.

### Options File

Place this option in the client options file (*dsm.opt*). You can set this option on the Scheduler tab, **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

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.opt* 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.opt* file or on the **schedule** command; the client can then be contacted at either an address or a port of your choice (useful for client systems with multiple network interface cards).

### Notes:

1. When changing the setting of this option in the client options file (*dsm.opt*) 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 Windows clients.

## Options File

Place this option in the client options file (*dsm.opt*). You can set this option on the **Scheduler** tab, **Schedule Mode** section of the Preferences editor.

## Syntax



## Parameters

### *POLLing*

The client scheduler queries the server for scheduled work at prescribed time intervals. This is the default. You can set the time intervals using the *querschedperiod* option.

### *PRompted*

The client scheduler waits for the server to contact your client node when scheduled work needs to be done.

**Note:** If you use the **dsmc schedule** command and both *SCHEDMODE PRompt* and *COMMMethod V6Tcpip* are specified, the client and Tivoli Storage Manager server must be configured for IPv6. Additionally, the client host name must be set up for the IPv6 address.

## Examples

**Options file:**

    schedmode prompted

**Command line:**

    -schedmod=po

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

---

## Schedrestretrdisabled

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 client options file (dsm.opt) for the scheduler. You can set this option on the **Scheduler** tab, **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 Windows 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 options file (dsm.opt). You can set this option on the **Command Line** tab, **Number of lines to display on screen** field of the Preferences editor.

## Syntax

►►—SCROLLLines— *number* —————▶▶▶

## Parameters

*number*

Specifies the number of lines of information that are displayed on your screen at one time. The range of values is 1 through 80; the default is 20.

## Examples

**Options file:**

```
scrolllines 25
```

**Command line:**

```
-scroll=25
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This 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 Windows 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 options file (dsm.opt). You can set this option on the **Command Line** tab, **Pause after displaying the following number of lines** field of the Preferences editor.

## Syntax



## 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.

---

## Sessioninitiation

Use the *sessioninitiation* option to control whether the server or client initiates sessions through a firewall. The default is that the client initiates sessions. You can use this option with the **schedule** command.

For the client scheduler, it is unnecessary to open *any* ports on the firewall. If you set the *sessioninitiation* option to *serveronly*, the client will not attempt to contact the server. *All sessions must be initiated by server prompted scheduling* on the port defined on the client with the *tcpclientport* option. The *sessioninitiation* option only affects the behavior of the client scheduler running in the prompted mode. If you set the *sessioninitiation* option to *serveronly*, with the exception of client acceptor daemon-managed schedulers, the command-line client, backup-archive client 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 scheduler from the command line (`dsmc schedule`) and enter the password for your node when prompted or use the following **dsmcutil** command to write the password into the Windows Registry:

```
dsmcutil updatepw /node:nnn /password:ppp /validate:no
```

To avoid this problem, when configuring the client scheduler using the setup wizard, ensure that the **Contact the TSM Server to validate password** checkbox on the TSM Authentication page is unchecked. See “Configuring the client scheduler” on page 21 for more information.

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 Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the Scheduler tab, **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 392. Except for client acceptor daemon-managed schedulers, the command-line client, backup-archive client 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

The *shmport* option specifies the TCP/IP port address of a server when using shared memory. All shared memory communications start with a TCP/IP connection.

**Note:** The value specified for the *shmport* option in the client options file (dsm.opt) must match the value specified for *shmport* in the server options file.

### Supported Clients

This option is valid for all Windows clients *only*.

### Options File

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

### Syntax

▶▶—SHMPort *port\_number*—————▶▶

### Parameters

*port\_number*

Specifies the port number. You can specify a value from 1000 to 32767. The default value is 1510.

### Examples

**Options file:**

shmport 1580

**Command line:**

Does not apply.

---

## Showmembers

Use the *showmembers* option with the following commands to display all members of a group:

- **query group**
- **query was**
- **query systemstate**
- **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 Windows clients.

## Syntax

▶▶—SHOWMembers—————▶▶

## Parameters

There are no parameters for this option.

## Examples

**Command line:**

```
restore group {virtfs}\* -pick -showmembers
```

---

## Skipntpermissions

The *skipntpermissions* option bypasses processing of NTFS security information. Select this option for incremental backups, selective backups, archives, retrieves and restores. Use this option with the following commands:

- **archive**
- **incremental**
- **restore**
- **retrieve**
- **selective**

## Supported Clients

This option is valid for all Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **General** tab of the Preferences editor.

## Syntax



## Parameters

*No* If you specify *No*, the NTFS security information is backed up, restored, archived or retrieved. This is the default.

*Yes* If you specify *Yes*, the NTFS security information is not backed up, restored, archived or retrieved.

## Examples

**Options file:**  
skipntp yes

**Command line:**  
-skipntp=yes

---

## Skipntsecuritycrc

The *skipntsecuritycrc* option controls the computation of the security cyclic redundancy check (CRC) for a comparison of NTFS security information during an incremental or selective backup, archive, restore, or retrieve operation. If you set the *skipntsecuritycrc* option to *no* (the default), performance might be slower because the program must retrieve all the security descriptors.

Use this option with the following commands:

- **archive**
- **incremental**
- **restore**
- **retrieve**
- **selective**

## Supported Clients

This option is valid for all Windows clients.

## Options File

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

## Syntax



## Parameters

- No* If you specify *No*, the security CRC is generated during a backup. This is the default.
- Yes* If you specify *Yes*, the security CRC is not generated during a backup. All the permissions are backed up, but the program will not be able to determine if the permissions are changed during the next incremental backup. When the *skipntpermissions* option is set to *yes*, the *skipntsecuritycrc* option does not apply.

## Examples

- Options file:**  
skipnts no
- Command line:**  
-skipnts=no

---

## Snapshotcachelocation

Use the *snapshotcachelocation* option during an online image backup or open file support operation to specify the location where the LVSA places the cache files. The cache files contain changes made to a volume by other system applications during an online image backup or open file support operation. For an online image backup, the cache files are backed up to the server and the changes are applied to the volume image when you restore the volume. This option is only valid if OFS or online image backup has been configured and the LVSA has been selected as the snapshot provider.

**Recommendation:** Do not specify the *snapshotcachelocation* option on a different volume than the target. For example, do not use non-resident cache.

For an online image backup, use the *snapshotcachelocation* option with the **backup image** command, in the *dsm.opt* file, or with the *include.image* option.

For open file support operations, you can specify the *snapshotcachelocation* option in your *dsm.opt* file or the with the *include.fs* option. See “Include options” on page 280 for more information.

## Supported Clients

This option is valid for the all supported Windows clients *only*.

## 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

▶▶—SNAPSHOTCACHELocation— *path* —▶▶

## Parameters

*path*

Specifies the location where the LVSA will place the cache files. The path must specify a formatted volume. With the single drive enhancements of LVSA in Tivoli Storage Manager Version 5.3, cache files can be placed on the same volume that is being backed up. The default location is to create the OBF files on the volume on which the snapshot is taken.

## Examples

**Options file:**

```
snapshotcachelocation c:\temp\
```

**Command line:**

```
-snapshotcachelocation=c:\temp\
```

---

## Snapshotcachesize

Use the *snapshotcachesize* option to determine the size of the cache file containing the original data for blocks that change during an online image backup or open file support operation. The value is the percent of the used space on the volume on which the snapshot is taken. On a volume with high file system 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. This option is only valid if OFS or online image backup has been configured and the LVSA has been selected as the snapshot provider.

For online image backups, use the *snapshotcachesize* option with the **backup image** command, the *include.image* option, or in your *dsm.opt* file.

For open file support operations, use the *snapshotcachesize* option in your *dsm.opt* file or with the *include.fs* option. See “Include options” on page 280 for more information.

## Supported Clients

This option is valid for all Windows clients. 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

▶—SNAPSHOTCACHESize— *size* —▶

## Parameters

*size*

Specifies the maximum size of the OBF. The value is a percent of the total size of the volume where the OBF resides, as determined by the *snapshotcachelocation* option. The range of values is 1 through 100 percent; the default is 1 percent, but not less than 500 MB.

If the calculated cache file size exceeds the free space on the drive where the cache file is written, Tivoli Storage Manager truncates the maximum cache file size to 90 percent of the available free space, and the following error message is logged in the *dsmerror.log*:

Snapshot cache size adjusted to accommodate insufficient space

## Examples

**Options file:**

```
snapshotcachesize 40
```

**Command line:**

```
-snapshotcachesize=40
```

---

## Snapshotfsidleretries

Use the *snapshotfsidleretries* option to specify the number of additional times the LVSA should try to achieve the snapshot file system idle wait time before the online image backup or open file support operation fails. This option is only valid if OFS or online image backup has been configured, and the LVSA has been selected as the snapshot provider.

While the snapshot provider gives a point-in-time image of the volume, it does not guarantee that there is no disk related activity while the snapshot is taken. This might lead to an image backup that will require you to run `chkdsk` before the image backup is accessed. A message will indicate that a `chkdsk` is necessary when you access the volume after the restore. In the case of OFS, a file might not be able to be read. If the file cannot be read, a warning of the file read problem will appear in the `dsmerror.log` as a failed backup file, and the file will be backed up during the next incremental backup.

For an online image backup, use the *snapshotfsidleretries* option with the **backup image** command, the *include.image* option, or in your `dsm.opt` file. Place the *include.image* statement containing the *snapshotfsidleretries* value in your `dsm.opt` file.

For open file support operations, use the *snapshotfsidleretries* option in your `dsm.opt` file or the with the *include.fs* option. See “Include options” on page 280 for more information.

If the value you specified with the *snapshotfsidlewait* option cannot be satisfied in the number of retries you specify with the *snapshotfsidleretries* option, the snapshot and the image backup will fail. If this occurs, you must adjust the *snapshotfsidlewait* and *snapshotfsidleretries* values until the snapshot is successful. You can also use the *presnapshotcmd* and *postsnapshotcmd* options to temporarily quiesce applications until the snapshot starts. See “Snapshotfsidlewait” on page 371, “Presnapshotcmd” on page 336, and “Postsnapshotcmd” on page 327 for more information. You can also use the Preferences editor to specify these options.

## Supported Clients

This option is valid for the Windows 32-bit 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

→—SNAPSHOTFSIDLERetries— *retrynumber* —→

## Parameters

	<i>retrynumber</i>
	Specifies the number of additional times the LVSA should try to achieve the
	snapshot file system idle wait time before the online image backup operation
	fails. The range of values is zero through 99; the default is 99.

## Examples

**Options file:**

```
include.image h: MYMC snapshotfsidledretries=2
```

**Command line:**

```
-snapshotfsidleretries=2
```



---

## Snapshotfsidlewait

Use the *snapshotfsidlewait* option during an online image backup or open file support operation to specify the amount of time that must pass in which there is no write activity (read activity is ignored) on a volume before a snapshot can be taken. This allows the disk to reach an idle state. This option is only valid if OFS or online image backup has been configured, and the LVSA has been selected as the snapshot provider.

While the snapshot provider gives a point-in-time image of the volume, it does not guarantee that there is no disk related activity while the snapshot is taken. This might lead to an image backup that will require you to run `chkdsk` before the image backup is accessed. A message will indicate that a `chkdsk` is necessary when you access the volume after the restore. In the case of OFS, a file might not be able to be read. If the file cannot be read, a warning of the file read problem will appear in the `dsmerror.log` as a failed backup file, and the file will be backed up during the next incremental backup. The snapshot file system idle wait period can decrease the likelihood of creating a snapshot while disk activity is occurring on the volume by delaying the creation of the snapshot until the disk is idle (with no write activity) for a specified amount of time. For example, if you specify a value of 2 seconds, a period of 2 seconds must pass without write activity on a volume before a snapshot can occur.

This option is only valid if the LVSA is installed and configured on your system allowing you to perform an online image backup or open file support operation.

For an online image backup, you can use the *snapshotfsidlewait* with the **backup image** command, the *include.image* option, or in your `dsm.opt` file when performing an online image backup. Place the *include.image* statement containing the *snapshotfsidlewait* value in your `dsm.opt` file.

For open file support operations, use the *snapshotfsidlewait* option in your `dsm.opt` file or the with the *include.fs* option. See “Include options” on page 280 for more information.

## Supported Clients

This option is valid for the Windows 32-bit 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

▶—SNAPSHOTFSIDLEwait— *maximum wait time* —————▶  
  └ *minimum wait time* ─┘

## Parameters

| *maximum wait time*  
|       Specifies the maximum amount of time that must pass during which there is  
|       no write activity on a volume before a snapshot can be taken. You can specify  
|       **ms** (milliseconds) or **s** (seconds) qualifiers with the value. If you do not specify

a qualifier, the value you specify is interpreted as seconds. The range of values is zero through 999; the default value is 2 s. A value of zero (0) indicates that the snapshot will occur immediately without waiting for a period of inactivity. You can specify a *maximum wait time* value *without* specifying a *minimum wait time* value.

#### *minimum wait time*

Specifies the minimum amount of time that must pass during which there is no write activity on a volume before a snapshot can be taken. You can specify **ms** (milliseconds) or **s** (seconds) qualifiers with the value. If you do not specify a qualifier, the value you specify is interpreted as seconds. The range of values is zero through 999; the default value is 50 **ms**.

If you set *snapshotfsidlewait* to 6,1 and *snapshotfsidleretries* to 5, Tivoli Storage Manager waits for i/o inactivity for 5 seconds, then retries 5, 4, 3, 2, and finally 1. The maximum value, minimum value, and the value of *snapshotfsidleretries* are used to determine the value to decrease by on each retry. This allows the longer wait times to be attempted first but still allows snapshots to occur if the activity is too high. If *snapshotfsidleretries* is set to zero (0), no retries are attempted, even if a minimum value is specified.

If you specify both maximum and minimum wait times, the following rules apply:

#### **Notes:**

1. You must specify the *maximum wait time*.
2. If you do not specify a *minimum wait time*, the *maximum wait time* is used until the value of *snapshotfsidleretries* is met.
3. The *maximum wait time* must be greater than or equal to the *minimum wait time* or you will receive an invalid option error.
4. The *maximum wait time* qualifier can be different than the *minimum wait time* qualifier. For example, SNAPSHOTFSIDLEWait 4s,500ms is valid.
5. You can use either a comma or a space to separate the *maximum wait time* and *minimum wait time*.

## Examples

#### **Options file:**

```
SNAPSHOTFSIDLEWait 5
SNAPSHOTFSIDLEWait 5,1
SNAPSHOTFSIDLEWait 4S,500MS
SNAPSHOTFSIDLEWait 7S,1
SNAPSHOTFSIDLEWait 7,1S
SNAPSHOTFSIDLEWait 7,500MS
```

```
Include.Fs
include.fs f: snapshotproviderfs=lvsa
    snapshotfsidlewait=4s,500ms
```

```
Include.Image
include.image f: snapshotproviderimage=lvsa
    snapshotfsidlewait=7s,1s
```

#### **Command line:**

```
-snapshotfsidlewait=2s
-snapshotfsidlewait=4s,500ms
```

---

## Snapshotproviderfs

Use the *snapshotproviderfs* option to enable snapshot-based file backup and archive operations, and to specify a snapshot provider.

### Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager API does not support this option.

### Options File

Specify this option in the client options file, *dsm.opt*, to enable snapshots. You can override the client-wide option for a specific operation by specifying this option on the command line for the backup and archive commands. You can also override the client-wide option for a specific file system by using the *include.fs* statement in the *dsm.opt* file. You can also set this option using the Preferences editor.

### Parameters

*value*

Specifies one of the following values:

#### LVSA

Specifies that the LVSA should be used to provide OFS support.

**Note:** LVSA must be installed before it can be specified. Install LVSA during the install process or use the setup wizard. Refer to “Configuring Open File Support (OFS)” on page 29 for more information.

#### VSS

Specifies that VSS should be used to provide OFS support (VSS is not supported on Windows XP).

#### NONE

Specifies that no snapshot provider should be used; This turns off OFS support. This is the default.

### Examples

#### Options file:

```
snapshotproviderfs VSS
include.fs d: snapshotproviderfs=vss
```

#### Command line:

```
-snapshotproviderfs=VSS
```

---

## Snapshotproviderimage

Use the *snapshotproviderimage* option to enable snapshot-based online image backup, and to specify a snapshot provider. You must be a root user to perform a snapshot-based image backup operation.

### Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager API does not support this option.

### Options File

Specify this option in the client options file, *dsm.opt*, to enable snapshots for all the file systems on the client. You can override the client-wide option for a specific operation by specifying this option on the command line for the **backup image** command. You can also override the client-wide option for a specific file system using the *include.image* statement in the *dsm.opt* file. You can also set this option using the Preferences editor.

### Parameters

*value*

Specifies one of the following values:

#### LVSA

Specifies that the LVSA should be used to provide online image support.

**Note:** LVSA must be installed before it can be specified. Install LVSA during the install process or use the setup wizard. Refer to “Configuring Open File Support (OFS)” on page 29 for more information.

#### VSS

Specifies that the VSS should be used to provide online image support (VSS is not supported on Windows XP).

#### NONE

Specifies that no snapshot provider should be used. This turns off online image support. This is the default

### Examples

#### Options file:

```
snapshotproviderfs VSS
include.fs d: snapshotproviderfs=vss
```

#### Command line:

```
-snapshotproviderfs=NONE
```

---

## 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 network share mounted file systems. Both the back up specification (source) and the *snapshotroot* value can be a network share mounted file specification. For example, the *snapshotroot* option can be used to back up a network share file system hosted on a network-attached storage (NAS) that supports snapshot.

In the following example, `c:\snapshots\snapshot.0` is network share mounted from a NAS filer and `\\florance\c$` represents the snapshot created at the NAS filer.

```
dsmc incr \\florance\c$ -snapshotroot=c:\shapshots
\snapshot.0
```

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 c: drive and mounts it as the NTFS junction point `\\floreance\c$\snapshots\snapshot.0`. If you back up this data using the following command:

```
dsmc incremental \\floreance\c$\snapshots\snapshot.0
```

a unique file space called `\\floreance\c$\snapshots\snapshot.0` is created on the server. However, you might want to associate the snapshot data with the data already processed for the c: drive (`\\floreance\c$`). Using the *snapshotroot* option, you can associate the data with the file space corresponding to the c: drive (`\\floreance\c$`) on the Tivoli Storage Manager server:

```
dsmc incr c: -snapshotroot=\\floreance\c$\snapshots\snapshot.0
-or-
dsmc incr \\floreance\c$ -snapshotroot=\\floreance\c$\snapshots\
snapshot.0
```

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 incr c: -snapshotroot=\\floreance\c$\snapshots\snapshot.1
```

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 incr c:\dir1\* -subdir=yes -snapshotroot=\\floreance\c$\
snapshots\snapshot.1
dsmc sel c:\dir1\sub1\file.txt -snapshotroot=\\floreance\c$\
snapshots\snapshot.1
dsmc archive c:\mydocs\*.doc -snapshotroot=\\floreance\c$\
snapshots\snapshot.1
```

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 c: drive), and *not* the name of the target of the snapshot (\\florence\c\$\snapshots\snapshot.1). 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 c:\dir1\...\*.txt 1yrmgmtclass
exclude \\florence\c$\mydocs\*.doc
```

The following include-exclude statements are not valid because they contain the name of the snapshot:

```
include \\florence\c$\snapshots\snapshot.1\dir1\...\
*.txt 1yrmgmtclass
exclude \\florence\c$\mydocs\*.doc
```

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 incr c: -snapshotroot=\\florence\c$\snapshots\snapshot.0
dsmc incr c:\dir1\* -snapshotroot=\\florence\c$\snapshots\
snapshot.0
```

The following command is invalid because it contains two file specifications:

```
dsmc incr c:\dir1\* e:\dir1\* -snapshotroot=\\florence\c$\
snapshots\snapshot.0
```

The following command is invalid because it contains no file specification:

```
dsmc incr -snapshotroot=\\florence\c$\snapshots\snapshot.0
```

#### 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 cannot use the *snapshotroot* option with any backup command, such as **backup image**, **backup systemobjects**, etc.
4. If you specify the *snapshotroot* option in conjunction with the Tivoli Storage Manager Open File Support feature provided by the Logical Volume Snapshot Agent (LVSA), the Tivoli Storage Manager Open File Support feature will not be used for the operation. It is strongly recommended that you use the Tivoli Storage Manager LVSA to provide open file support instead of using the *snapshotroot* option in conjunction with another snapshot provider.

**Attention:** Use the *snapshotroot* option with caution in conjunction with the Tivoli Storage Manager journal-based backup feature. Since there is no coordination between the Tivoli Storage Manager journal and the third party snapshot provider, unwanted behavior can occur with journal notifications received after the snapshot occurs. For example, files might not be backed up, or they might be backed up redundantly to the Tivoli Storage Manager server. If you want to use the Tivoli Storage Manager journal-based backup feature in conjunction with a snapshot provider, it is strongly recommend that you use it in conjunction with the Tivoli Storage Manager LVSA which provides full integration with the journal-based backup feature.

5. 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 Windows 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 incr c: -snapshotroot=\\florence\c$\snapshots\snapshot.0
```

---

## Srvprepostscheddisabled

The *srvprepostscheddisabled* option specifies whether to prevent the pre-schedule and post-schedule commands specified by the Tivoli Storage Manager administrator from executing on the client system, when performing scheduled operations. The *srvprepostscheddisabled* option can be used in conjunction with the *schedcmddisabled* and *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 client options file (dsm.opt) for the scheduler. You can set this option on the **Schedule Command** button of the Preferences editor.

### Syntax



### Parameters

- No* Specifies that Tivoli Storage Manager allows pre-schedule and post-schedule commands defined by the Tivoli Storage Manager administrator to execute on the client system, when performing scheduled operations. If a pre-schedule or a post-schedule command is defined by both the client and the Tivoli Storage Manager administrator, the command defined by the administrator will override the corresponding command defined in the client option file. This is the default.
- Yes* Specifies that Tivoli Storage Manager prevents pre-schedule and post-schedule commands defined by the Tivoli Storage Manager administrator to execute on the client system, when performing scheduled operations. If a pre-schedule or a post-schedule command is defined by both the client and the Tivoli Storage Manager administrator, the command defined by the administrator will *not* override the corresponding command defined in the client option file. This option can be used in conjunction with the *schedcmddisabled* 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 system, 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 Windows 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 client options file (dsm.opt) for the scheduler. You can set this option on the **Image-Snapshot** tab 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 system, when performing scheduled image snapshot backup operations. If a pre-snapshot or a post-snapshot command is defined by both the client and the Tivoli Storage 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 system, when performing scheduled image snapshot backup operations. If a pre-snapshot or a post-snapshot command is defined by both the client and the Tivoli Storage 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.

---

## Ssl

### Authorized User

The *ssl* option is used to enable Secure Socket Layer (SSL), to provide secure client and server communications.

### Supported Clients

This option is valid for Windows clients.

### Options File

Place this option in the client options file (*dsm.opt*). You can set this option on the **Communications** tab of the Preferences editor.

### Syntax



### Parameters

*No* Specifies that the Tivoli Storage Manager client disables SSL. This is the default.

*Yes* Specifies that the Tivoli Storage Manager client enables SSL.

In addition to specifying *ssl=yes* to enable SSL, you should also change the value of the *tcppport* option, because the Tivoli Storage Manager server is usually set up to listen for SSL connections on a separate port.

### Examples

#### Options file:

```
ssl 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 Windows 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 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 333.

## 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:

```
rest \path\dir1\* \path2\ -su=yes
rest \path\dir1\file* \path2\ -su=yes
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.

---

## Subfilebackup

The *subfilebackup* option specifies whether to enable adaptive subfile backup.

### Notes:

1. To use the *subfilebackup* option, the server must have "subfile YES" or subfile client. Ask your Tivoli Storage Manager administrator to set the proper values.
2. You can use open file support in conjunction with subfile backup operations.
3. The *subfilebackup* option does not work correctly for migrated files. If you use a combination of subfilebackup and non-subfilebackup for migrated files, your data might be corrupted on the server.

## Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Backup** tab of the Preferences editor.

## Syntax

►►—SUBFILEBackup no  
yes—►►

## Parameters

*No* Specifies that Tivoli Storage Manager does not use adaptive subfile backup. This is the default.

*Yes* Specifies that Tivoli Storage Manager uses adaptive subfile backup.

## Examples

### Options file:

```
subfileb yes
```

### Command line:

```
-subfileb=yes
```

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

---

## Subfilecachepath

The *subfilecachepath* option specifies the path where the client cache resides for adaptive subfile backup processing. If you do not specify a path, Tivoli Storage Manager creates a path called \cache under the directory where the Tivoli Storage Manager executables reside.

All directories and subdirectories in the path name you specify with the *subfilecachepath* option must exist. For example, if you specify c:\temp\cache, the c:\temp directory must already exist. You should exclude the subfilecache directory from backup.

### Supported Clients

This option is valid for all Windows 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 options file (dsm.opt). You can set this option on the **Backup** tab of the Preferences editor.

### Syntax

▶▶—SUBFILECACHEPath— *path\_name* —————▶▶

### Parameters

*path\_name*

Specifies the path name where the client cache resides for adaptive subfile backup processing.

### Examples

#### Options file:

```
subfilecachep c:\temp\cc_tsm
```

#### Command line:

```
-subfilecachep=c:\temp\cc_tsm
```

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

---

## Subfilecachesize

The *subfilecachesize* option specifies the client cache size for adaptive subfile backup. If the cache size is too small, base files for some files will not be cached and subfile processing will not apply for them. However, setting the value too large can take up more disk space than can be spared. The files maintained in the cache should closely reflect the files used on a regular basis.

### Supported Clients

This option is valid for all Windows clients. The server can also define this option.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Backup** tab of the Preferences editor.

### Syntax

►►—SUBFILECACHESize— *size* —————►►

### Parameters

*size*

Specifies the size, in megabytes, of the client cache for adaptive subfile backup processing. The range is 1 through 1024 (1 GB); the default is 10.

### Examples

**Options file:**

```
subfilecaches 10
```

**Command line:**

```
-subfilecaches=10
```

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

---

## Tapeprompt

The *tapeprompt* option specifies whether you want Tivoli Storage Manager to wait for a tape mount if it is required for a backup, archive, restore, or retrieve process, or to be prompted for a choice.

In the Tivoli Storage Manager GUI, the Media Mount dialog can display the **Information Not Available** value in the Device and Volume Label fields if you perform a standard (also known as classic) restore or retrieve operation. This value means that this information is only available for no query restore or retrieve operations; not a standard restore or retrieve operation. The Device field displays the name of the device on which to mount the media needed to process an object. The Volume Label field displays the name of the volume needed to process an object. See “Standard query restore, no query restore, and restartable restore” on page 109 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 Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **General** tab, **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

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 Windows 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 options file (dsm.opt). You can set this option on the **Communication** tab, **Admin Port** field of the Preferences editor. See the *IBM Tivoli Storage Manager Administrator's Reference* for your operating system, for more information.

## Syntax

▶▶—TCPADMINPort [admin\_port\_address] ▶▶

## Parameters

*admin\_port\_address*  
Specifies the port number of the server. The default value is the value of the *tcpport* option.

## Examples

**Options file:**  
tcpadminport 1502

---

## Tcpbuffsize

The *tcpbuffsize* option specifies the size of the internal TCP/IP communication buffer used to transfer data between the client node and server. Although it uses more memory, a larger buffer can improve communication performance.

### Supported Clients

This option is valid for all Windows clients.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Communication** tab, **Buffer Size** field of the Preferences editor.

### Syntax

▶▶—TCPBuffsize— *size* —————▶▶

### Parameters

*size*

Specifies the size, in kilobytes, that you want to use for the internal TCP/IP communication buffer. The range of values is 1 through 512; the default is 32.

Depending on the operating system communication settings, your system might not accept all values in the range of 1 through 512.

### Examples

**Options file:**

```
tcpb 32
```

**Command line:**

```
-tcpbuffsize=32
```

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

---

## Tpcadaddress

The *tpccadaddress* option specifies a TCP/IP address for dsmcad. Normally, this option is not needed. Use this option only if your client node has more than one TCP/IP address, or if TCP/IP is not the default communication method (see “Commmethod” on page 217).

### Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

### Options File

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

### Syntax

►►—TCPCADAddress— *cad\_address*—————►►

### Parameters

*cad\_address*

| Specifies a TCP/IP Internet domain name or a numeric IP address. You can  
| specify TCPip or V6Tcpip. This option uses TCP/IP Version 4 if TCPip is  
| specified and uses Version 6 if V6Tcpip is specified. You can only use IPv6  
| addresses if you specified the *commmethod V6Tcpip* option.

### Examples

#### Options file:

```
tcpcada dsmclnt.sanjose.ibm.com
```

#### Command line:

```
-tcpcadaddress=128.33.10.249  
-tcpcadaddress=khoyt.mycompany.mydomain.com  
-tcpcadaddress=2002:92b:111:221:128:33:10:249
```

This option is valid only on the initial command line of the dsmcad program. It is not valid with other dsm modules.

---

## Tcpclientaddress

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 Windows clients. 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 **Scheduler** tab, **Your TCP/IP address** field of the Preferences editor.

### Syntax

▶—TCPCLIENTAddress— *client\_address* —▶

### Parameters

*client\_address*

Specifies the TCP/IP address you want the server to use to contact your client node. Specify a TCP/IP Internet domain name or a numeric IP address. The numeric IP address can be either a TCP/IPv4 or TCP/IPv6 address. You can only use IPv6 addresses if you specified the *commmethod V6Tcpip* option.

### Examples

**Options file:**

```
tcpclienta dsmclnt.sanjose.ibm.com
```

**Command line:**

```
-tcpclientaddress=128.33.10.249  
-tcpclientaddress=khoyt.mycompany.mydomain.com  
-tcpclientaddress=2002:92b:111:221:128:33:10:249
```

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

---

## Tcpclientport

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 Windows clients. 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 **Scheduler** tab, **Your TCP/IP port** field of the Preferences editor.

### Syntax

▶▶—TCPCLIENTPort— *client\_port\_address* —————▶▶

### Parameters

*client\_port\_address*

Specifies the TCP/IP port address you want the server to use to contact your client node. The range of values is 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

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 Windows clients.

### Options File

Place this option in the client options file (*dsm.opt*). You can set this option on the **Communication** tab 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

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 Windows clients.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Communication** tab, **Server Port** field of the Preferences editor.

### Syntax

▶▶—TCPPort— *port\_address* —————▶▶

### Parameters

*port\_address*

Specifies the TCP/IP port address that is used to communicate with a server. The range of values is 1000 through 32767; the default is 1500.

### Examples

**Options file:**

tcpp 1501

**Command line:**

-tcpport=1501

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



---

## Tcpserveraddress

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 Windows clients.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Communication** tab, **Server Address** field of the Preferences editor.

### Syntax

▶▶—TCPServeraddress— *server\_address*————▶▶

### Parameters

|                   *server\_address*  
|                   Specifies a 1 to 64 character TCP/IP address for a server. Specify a TCP/IP  
|                   domain name or a numeric IP address. The numeric IP address can be either a  
|                   TCP/IP v4 or TCP/IP v6 address. You can only use IPv6 addresses if you  
|                   specified the *commmethod V6Tcip* option.

### Examples

**Options file:**

tcps dsmchost.endicott.ibm.com

**Command line:**

-tcpserveraddress=129.33.24.99

-tcpserveraddress=2002:92b:111:221:128:33:10:249

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

---

## Tcpwindowsize

Use the *tcpwindowsize* option to specify, in kilobytes, the size you want to use for the TCP/IP sliding window for your client node. The sending host cannot send more data until it receives an acknowledgment and a TCP receive window update. Each TCP packet contains the advertised TCP receive window on the connection. A larger window allows the sender to continue sending data and can improve communication performance.

### Supported Clients

This option is valid for all Windows clients.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Communication** tab, **Window Size** field of the Preferences editor.

### Syntax

►—TCPWindowsize— *window\_size* —————►

### Parameters

*window\_size*

Specifies the size, in kilobytes, to use for your client node TCP/IP sliding window. The range of values is 0 through 2048. A value of 0 allows Tivoli Storage Manager to use the operating system default TCP window size. Values from 1 to 2048 indicate that the window size is in the range of 1KB to 2MB. 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.
4. Windows XP provides a larger TCP receive window size when communicating with hosts that also provide this support, known as RFC1323. In these environments, a value greater than 63 can be useful.

### 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.

Use this option if you want to change the default time format for the language of the message repository you are using.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time the client is called. Consult the documentation on your local system for details about setting up your locale definition.

**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**
- **query systemobject**
- **query systemstate**
- **restore**
- **restore image**
- **restore nas**
- **retrieve**
- **restore registry**
- **restore systemobject**
- **restore systemstate**
- **set event**

## Supported Clients

This option is valid for all Windows clients.

## Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Regional Settings** tab, **Time Format** field of the Preferences editor.

## Syntax

▶—TIMEformat— *format\_number* —▶

## Parameters

*format\_number*

Displays time in one of the formats listed below. Select the format number that

corresponds to the format you want to use. When you include the *timeformat* option in a command, it must precede the *fromtime*, *pittime*, and *totime* options.

- 1 23:00:00
- 2 23,00,00
- 3 23.00.00
- 4 12:00:00 A/P

## Examples

**Options file:**

```
timeformat 4
```

**Command line:**

```
-time=3
```

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 all Windows clients. 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 client options file (dsm.opt).

## 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 Windows 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 -todate=12/11/2003 c:\myfiles\
```

---

## 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 Windows 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 query backup -totime=23:59:00 -todate=06/30/2003 c:\mybackups\
```



---

## Txnbytelimit

The *txnbytelimit* option specifies the number of kilobytes the client program buffers before it sends a transaction to the server. A *transaction* is the unit of work exchanged between the client and server. 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 Windows clients.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the **General** tab **Transaction Buffer Size** field of the Preferences editor.

### Syntax

▶▶—TXNBytelimit— *number* —————▶▶

### Parameters

*number*

Specifies the number of kilobytes the client program 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 all Windows clients. 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

---

## Usedirectory

The *usedirectory* option queries the Active Directory for the communication method and server with which to connect. This option overrides the *commmethod* parameters specified in the client options file (*dsm.opt*). Optimally, the administrator enables only one server and one specific communication protocol for a given client node. The specification of this information in Active Directory is done using the Tivoli Storage Manager server on Windows, which has a wizard to assist with this configuration. If a node is registered to more than one server published in Active Directory, the first server returned in the Active Directory query will be used. If the client cannot contact the server, the client session will fail.

### Supported Clients

This option is valid for all Windows clients. 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 **Communication** tab of the Preferences editor.

### Syntax



### Parameters

- Yes* Specifies that the client will ignore *commmethod* parameters set in the client options file and query the Active Directory for the communication method and server with which to connect.
- No* Specifies that the client will use the communication method specified in the option file. If there is no communication method specified in the option file the default communication method and server are used.

### Examples

**Options file:**  
`usedirectory no`

**Command line:**  
`-usedir=yes`

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

---

## 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 Windows 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 c:\relx\dir1\ -v2archive -su=y
```

**Archives these files:**

```
c:\relx\dir1\file1  
c:\relx\dir1\file2  
c:\relx\dir1\file3  
c:\relx\dir1\dir2\file4  
c:\relx\dir1\dir2\file5
```

Note: Tivoli Storage Manager does not archive `c:\relx\dir1` and `c:\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 Windows 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 options file (*dsm.opt*). You can set this option on the **Command Line** tab, **Do not display process information on screen** checkbox of the Preferences editor.

## Syntax

▶▶—VErbose—▶▶

## Parameters

There are no parameters for this option.

## Examples

**Options file:**  
verbose

**Command line:**  
-verbose

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

---

## Verifyimage

Use the *verifyimage* option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log.

### Supported Clients

This option is valid for Windows 32-bit clients *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 d: -verifyimage
```

---

## Virtualfsname

Use the *virtualfsname* option with the **backup group** command to specify the name of the virtual file space for the group on which you want to perform the operation. The *virtualfsname* cannot be the same as an existing file space name.

## Supported Clients

This option is valid for all Windows 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=c:\dir1\filelist1 -groupname=group1  
-virtualfsname=\virtfs -mode=full
```

---

## Virtualnodename

The *virtualnodename* option specifies the node name of your workstation when you want to restore or retrieve files to a different workstation.

When you use the *virtualnodename* option in your client options file, or with a command:

- You must specify the name you specified with the *nodename* option in your client options file (dsm.opt). 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*. The password is not stored in the Windows Registry even when the *passwordaccess* option is set to *generate*. Once a connection to the server is established, then access is permitted to any file backed up using this login ID.

For the Java GUI, the *virtualnodename* option will only work if you are root user. If you need to use the *virtualnodename* option with a non-root user account, use the command-line client (or native GUI, for platforms that still have one).

## Supported Clients

This option is valid for all Windows clients.

## Options File

Place this option in the client 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.

---

## Vmchost

Use the *VMCHost* option with the **backup VM** or **query VM** commands to specify the host name of the VMware VirtualCenter or ESX server where the VMware Consolidated Backup commands are directed.

**Recommendation:** Use the VirtualCenter if it is available. If you cannot use a VirtualCenter server and you need to perform backups of multiple systems on multiple ESX servers, do not specify this option, but instead specify the option with the command so that it can be varied for each ESX server.

### Supported Clients

This option is valid for Windows clients that are configured as a VMware Consolidated Backup proxy.

### Options File

Place this option in the client options file (*dsm.opt*), or on the command line.

### Syntax

▶—*VMCHost*— *hostname* —▶

### Parameters

*hostname*

Specifies the host name of the VMware VirtualCenter or ESX server where the VMware Consolidated Backup commands will be directed.

### Examples

**Options file:**

```
vmchost vcenter.storage.usca.ibm.com
```

**Command line:**

```
-vmchost=esx1.storage.usca.ibm.com  
-vmchost=esx2.storage.usca.ibm.com
```

---

## Vmcpw

Use the *VMCPW* option with the **backup VM** or **query VM** commands to specify the password of the VMware VirtualCenter or the ESX user ID that is specified with the *VMCUSER* option.

**Recommendation:** Use the VirtualCenter if it is available. If you cannot use a VirtualCenter server and you need to perform backups of multiple systems on multiple ESX servers, do not specify this option, but instead specify the option with the command so that it can be varied for each ESX server.

### Supported Clients

This option is valid for Windows clients that are configured as a VMware Consolidated Backup proxy.

### Options File

Place this option in the client options file (*dsm.opt*), or on the command line.

### Syntax

▶▶—VMCPW— *pwname* —————▶▶

### Parameters

*pwname*

Specifies the password for the VMware VirtualCenter or ESX server where the VMware Consolidated Backup commands will be directed.

### Examples

**Options file:**

vmcpw SECRET

**Command line:**

-vmcpw=SECRET

---

## Vmcuser

Use the **VMCUSER** option with the **backup VM** or **query VM** commands to specify the user name of the VMware VirtualCenter or ESX server where the VMware Consolidated Backup commands are directed.

**Recommendation:** Use the VirtualCenter if it is available. If you cannot use a VirtualCenter server and you need to perform backups of multiple systems on multiple ESX servers, do not specify this option, but instead specify the option with the command so that it can be varied for each ESX server.

### Supported Clients

This option is valid for Windows clients that are configured as a VMware Consolidated Backup proxy.

### Options File

Place this option in the client options file (dsm.opt), or on the command line.

### Syntax

▶▶—VMCUSER— *username* —▶▶

### Parameters

*username*

Specifies the user name of the VMware VirtualCenter or ESX server where the VMware Consolidated Backup commands will be directed.

When working directly with an ESX server, a user ID capable of logging onto the ESX service console, such as "root", is required. When working with a Virtual Center server, a Windows user ID that is capable of logging onto the Virtual Center server is required.

### Examples

**Options file:**

vmcuser administrator

**Command line:**

-vmcuser=domainname\administrator

---

## Vmlist

Use the *VMLIST* option with the **backup VM** or **query VM** commands to specify the host name or list of host names of the virtual machine to back up.

Multiple values can be specified (see the following examples).

### Supported Clients

This option is valid for Windows clients that are configured as a VMware Consolidated Backup proxy.

### Options File

Place this option in the client options file (dsm.opt), or on the command line.

### Syntax

▶▶—VMLIST— *vmname* —————▶▶  
                  └─ [ *tsmnode* ] ─┘

### Parameters

*vmname*

Specifies the host name of the virtual machine to back up on the VMware backup proxy, or a list of host names.

*tsmnode*

This optional parameter can be used in situations where the virtual machine's name and Tivoli Storage Manager node name do not match, or where DNS name resolution is not available. The option defaults to the hostname portion of the value specified with *vmname*.

### Examples

#### Options file:

```
vmlist vm1,vm2,vm3  
vmlist vm3.domain.com,vm8.domain.com  
vmlist 9.22.100.85[vm4],9.22.100.90[vm5]  
vmlist 9.22.100.99,9.22.100.101  
vmlist vm3[tsmnodeX]
```

#### Command line:

```
-vmlist=vm1,vm2,9.88.200.4
```

---

## 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 Windows Server 2003 clients.

### Options File

Place this option in the client 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 quotation marks.

### Examples

**Options file:**

```
wasExpHome "C:\Program Files\IBM\WebSphere\Express\AppServer"
```

---

## Washome

Use the *washome* option in your client 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 Websphere Application Server on a new system. See “Websphere Application Server instance restore procedures” on page 549 for more information.

If you do not specify a value for this option, Tivoli Storage Manager uses the current value in the Windows Registry. If there are multiple installations of the Application Server on the same system, use a different options file for each installation with the proper path to the installation directory.

### Supported Clients

This option is valid for Windows Server 2003.

### Options File

Place this option in the client 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 quotation marks.

### Examples

Options file:

```
washome "c:\program files\mydir\appserver"
```

---

## Wasndhome

Use the *wasndhome* option in your client 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 Websphere Application Server on a new system. See “Websphere Application Server instance restore procedures” on page 549 for more information.

If you do not specify a value for this option, Tivoli Storage Manager uses the current value in the Windows Registry. If there are multiple installations of the Network Deployment Manager on the same system, use a different options file for each installation with the proper path to the installation directory.

### Supported Clients

This option is valid for Windows Server 2003 clients.

### Options File

Place this option in the client 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 quotation marks.

### Examples

Options file:

```
wasndhome "c:\program files\mydir\DeploymentManager"
```



---

## Wasnode

Use the *wasnode* option with the **set waspassword** command to specify the Websphere Application Server node name when performing the operation on the Websphere Application Server Network Deployment Manager or Application Server.

## Supported Clients

This option is valid for Windows Server 2003 clients.

## Syntax

►►—WASNode =— *nodename*—————►►

## Parameters

*nodename*

Specifies the Websphere Application Server node name when performing the operation on the Websphere Application Server 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 Network Deployment Manager (contains setup, application files, and configuration information), the Application Server, or both.

## Supported Clients

This option is valid for Windows Server 2003 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 system. This parameter displays the instance name, hostname, soap port information, installed path, the type of the Websphere Application Server (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 Websphere Application Server security is enabled, use the *wasuser* option with the **set waspassword** command to set the Websphere Application Server user name for each installation of Websphere Application Server on your system.

## Supported Clients

This option is valid for Windows Server 2003 clients.

## Syntax

►►—WASuser =— *username*—————►►

## Parameters

*username*

Specifies the Websphere Application Server user name when performing the operation on the Websphere Application Server 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 service and Web Client Agent service for communications with the Web client.

Values for both the Client Acceptor service 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 service and the Web Client Agent service. The port value TCP/IP assigns is in the range of 1024 through 5000.

**Note:** The Tivoli Storage Manager client API does not support this option.

### Supported Clients

This option is valid for all Windows clients.

### Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Web Client** tab, **WEB Ports** fields of the Preferences editor.

### Syntax

►—WEBPorts— *cadport*— *agentport*—◀

### Parameters

*cadport*

Specifies the *required* Tivoli Storage Manager Client Acceptor 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.

*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:**

```
-webports=2123,2124
```

---

## 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 51 shows a list of tasks related to entering commands.

Table 51. Entering commands

Task	Page
Starting and ending a client command session	427
Entering client commands	428
Using wildcard characters	431

Table 52 provides an alphabetical list of the commands, a brief description, and where to locate more information.

Table 52. Commands

Command	Description	Page
<b>archive</b>	Archives files from a workstation to Tivoli Storage Manager storage.	433
<b>backup asr</b>	Generates Automated System Recovery (ASR) files in the adsm.sys\ASR staging directory and backs them up to the ASR file space on the server. This command is valid for the Windows XP and Windows Server 2003 clients <i>only</i> .	436
<b>backup certserverdb</b>	Backs up a Windows XP certificate server database.	437
<b>backup complusdb</b>	Backs up the Windows XP COM+ database.	438
<b>backup eventlog</b>	Backs up Windows XP event logs.	439
<b>backup group</b>	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.	440
<b>backup image</b>	Creates an image backup of one or more file systems or logical volumes that you specify.	442
<b>backup nas</b>	Creates an image backup of one or more file systems belonging to a Network Attached Storage (NAS) file server.	447
<b>backup registry</b>	Backs up the Windows XP registry.	450
<b>backup sysfiles</b>	Backs up Windows XP system and boot files.	451
<b>backup systemobject</b>	Backs up all valid Windows XP system objects.	452
<b>backup systemstate</b>	Backs up all bootable system state and system services components as one object to provide a consistent point-in-time snapshot of the system state. This command is valid for the Windows Server 2003 and Windows Vista clients <i>only</i> .	453

Table 52. Commands (continued)

Command	Description	Page
<b>backup was</b>	Backs up the WebSphere Application Server Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server to the Tivoli Storage Manager server.	455
<b>backup wmi</b>	Backs up a Windows XP Windows Management Instrumentation (WMI) repository.	459
<b>backup vm</b>	Backs up virtual machines specified in the <i>VMLIST</i> option.	457
<b>cancel process</b>	Displays a list of current NAS (if NDMP support is enabled) and server-free image backup and restore processes for which the administrative user has authority.	460
<b>cancel restore</b>	Displays a list of restartable restore sessions from which you can select one to cancel.	461
<b>delete access</b>	Deletes authorization rules for files or images that are stored on the server.	462
<b>delete archive</b>	Deletes archived files from Tivoli Storage Manager server storage.	463
<b>delete backup</b>	Deletes active and inactive backup files from Tivoli Storage Manager server storage.	465
<b>delete filespace</b>	Deletes file spaces in Tivoli Storage Manager server storage.	468
<b>delete group</b>	Deletes a group backup on the Tivoli Storage Manager server.	470
<b>expire</b>	Inactivates backup objects that you specify in the file specification or with the <i>filelist</i> option.	472
<b>help</b>	Displays a Table of Contents of help topics for the command-line client.	474
<b>incremental</b>	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.	475
<b>loop</b>	Starts an interactive command session.	482
<b>macro</b>	Executes commands within a macro file that you specify.	484
<b>monitor process</b>	Displays a list of current NAS image backup and restore processes from which you can select one to cancel	485
<b>preview</b>	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.	486
<b>query access</b>	Displays a list of current authorization rules.	487
<b>query archive</b>	Displays a list of archived files.	488
<b>query asr</b>	Displays information about the active backup of the Automated System Recovery files on the Tivoli Storage Manager server. This command is valid for the Windows XP and Windows 2003 clients only.	490
<b>query backup</b>	Displays a list of backup versions.	491

Table 52. Commands (continued)

Command	Description	Page
<b>query backupset</b>	Queries a backup set from a local file, tape device, or the Tivoli Storage Manager server.	495
<b>query complusdb</b>	Displays information about the active backup of the COM+ database system object on the Tivoli Storage Manager server. This command is valid for the Windows XP client only.	497
<b>query eventlog</b>	Displays information about the active backup of the Windows XP event log system object on the Tivoli Storage Manager server.	498
<b>query filespace</b>	Displays a list of file spaces in Tivoli Storage Manager storage. You can also specify a single file space name to query.	499
<b>query group</b>	Displays information about group backups and their members.	501
<b>query image</b>	Displays information about image backups.	503
<b>query inclexcl</b>	Displays a list of include-exclude statements in the order in which they are processed during backup and archive operations.	505
<b>query mgmtclass</b>	Displays information about available management classes.	506
<b>query node</b>	Displays all the nodes for which an administrative user ID has authority to perform operations.	507
<b>query options</b>	Displays all or part of your options and their current settings.	508
<b>query registry</b>	Displays information about the active backup of a Windows XP registry system object on the Tivoli Storage Manager server.	509
<b>query restore</b>	Displays a list of your restartable restore sessions in the server database.	510
<b>query schedule</b>	Displays information about scheduled events for your node.	511
<b>query session</b>	Displays information about your session, including the current node name, when the session was established, server information, and server connection information.	512
<b>query sysfiles</b>	Displays information about the active backup of the Windows XP system and boot files on the Tivoli Storage Manager server.	513
<b>query systeminfo</b>	Gathers Tivoli Storage Manager system information and outputs this information to a file or the console.	514
<b>query systemobject</b>	Displays information about the active backup of all valid Windows XP system objects on the Tivoli Storage Manager server.	516
<b>query systemstate</b>	Displays information about the backup of the system state on the Tivoli Storage Manager server. This command is valid for the Windows Server 2003 and Windows Vista clients only.	517

Table 52. Commands (continued)

Command	Description	Page
<b>query was</b>	Displays backups of the WebSphere Application Server Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server that match the node name and type of the Websphere Application Server group backup that you specify.	518
<b>query wmi</b>	Displays information about the active backup of the Windows Management Instrumentation (WMI) repository on the Tivoli Storage Manager server. This command is valid for Windows XP clients only.	520
<b>restart restore</b>	Displays a list of restartable restore sessions from which you can one to restart.	521
<b>restore</b>	Restores copies of backup versions of your files from a Tivoli Storage Manager server.	522
<b>restore asr</b>	Restores the active version of the Automated System Recovery (ASR) files from the Tivoli Storage Manager server. This command is valid for the Windows XP and Windows Server 2003 clients only.	528
<b>restore backupset</b>	Restores a backup set from the Tivoli Storage Manager server or a local file. You can also restore a backup from a tape device.	529
<b>restore complusdb</b>	Restores a Windows XP COM+ database from the Tivoli Storage Manager server.	532
<b>restore eventlog</b>	Restores the active Windows XP event logs from the Tivoli Storage Manager server.	533
<b>restore group</b>	Restores specific members or all members of a group backup.	534
<b>restore image</b>	Restores a file system or raw volume image backup.	536
<b>restore nas</b>	Restores the image of a file system belonging to a Network Attached Storage (NAS) file server.	540
<b>restore registry</b>	Restores the Windows XP registry.	542
<b>restore sysfiles</b>	Restores Windows XP system and boot files on the Tivoli Storage Manager server.	544
<b>restore systemobject</b>	Restores valid active or inactive Windows XP system objects stored on the Tivoli Storage Manager server.	545
<b>restore systemstate</b>	Restores a backup of the system state. This command is valid for the Windows Server 2003 and Windows Vista clients <i>only</i> .	547
<b>restore was</b>	Restores the WebSphere Application Server Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server from the Tivoli Storage Manager server.	549
<b>restore wmi</b>	Restores a Windows XP Windows Management Instrumentation (WMI) repository from the Tivoli Storage Manager server.	552



Table 52. Commands (continued)

Command	Description	Page
<b>retrieve</b>	Retrieves copies of archived files from the Tivoli Storage Manager server.	553
<b>schedule</b>	Starts the client scheduler on the workstation.	557
<b>selective</b>	Backs up selected files.	559
<b>set access</b>	Authorizes another user to access your backup versions, archived copies, or image backupsbackup or archive data.	563
<b>set event</b>	Allows you to specify the circumstances for when archived data is deleted.	566
<b>set password</b>	Changes the Tivoli Storage Manager password for your workstation.	568
<b>set waspassword</b>	When Websphere Application Server security is enabled, this command allows you to set your WebSphere node name, user name, and password for the Websphere Application Server Network Deployment Manager or the Application Server.	569

## Starting and ending a client command session

You can start or end a client command session in either batch mode or interactive mode. Use batch mode when you want to enter a *single* client command. Tivoli Storage Manager processes the command and returns to the command prompt.

Use interactive mode when you want to enter a *series* of commands. Since Tivoli Storage Manager establishes connection to the server only once for interactive mode, a series of commands can be processed more quickly. Tivoli Storage Manager processes the commands and returns to the **tsm>** prompt.

## Process commands in batch mode

Some options are valid *only* on the initial command line and not in interactive mode. These options generally affect the operation of the entire session. For example, the command **dsmc query session -errorlogname=myerror.log** is accepted and it does name the error log. However, it is accepted simply because it appears in the initial command, even though the option is not valid for the query command.

There are also some options that are always valid on the initial command line as well as on individual commands in interactive mode. Therefore, certain options are accepted on the initial command line even though they have no effect on the command being entered. For example, **dsmc query session -subdir=yes** is a valid command, but in this case the *-subdir* option has no effect on the command that was entered. Refer to Chapter 9, "Using processing options," on page 165 for the option that you are using to see if it is valid only on the initial command line.

When you enter a *single* command in batch mode, precede it with the executable program name, **dsmc**. For example, to process the **incremental** command in batch mode, you would enter:

```
dsmc incremental
```

Tivoli Storage Manager prompts you each time you enter a command if the *passwordaccess* option is set to *prompt* and authentication on the server is set to *On*. Type your password and press Enter.

You can also enter your password using the *password* option with a command, but your password appears on the screen. For example, if your password is *secret*, enter:

```
dsmc incremental -password=secret
```

If you set the *passwordaccess* option to *generate* in your *dsm.opt* file, you do not need to specify the password with the command. Tivoli Storage Manager *only* prompts you for your password if you are registering your workstation with a server or manually changing your password.

## Process commands in interactive mode

Use the *interactive* mode (or *loop* mode) to enter a series of commands. Enter **dsmc** on the command line and press Enter. When the **tsm>** command prompt appears, type the command name and press Enter. *Do not* precede each command with the executable program name, **dsmc**. Alternatively, you can enter **dsmc loop** on the command line to start a client command session in interactive mode. **Loop** is the default command for **dsmc**.

If a password is required, Tivoli Storage Manager prompts you before you enter the first command. Type your password and press Enter. You can also enter your password using the *password* option with the **loop** command, but your password appears on the screen. For example, if your password is *secret*, enter:

```
dsmc loop -password=secret
```

To end an interactive session, enter **quit** at the prompt.

---

## Entering client commands

A client command can include one or more of these components:

- Command name
- Options
- Parameters

The sections that follow describe each of these components.

### Command name

The first part of a command is the command name. The command name consists of a single word, such as **help** or **schedule**, or an action word and an object for that action, such as **query archive**. Enter the full command name, or its minimum abbreviation. For example, you can enter any of the following versions of the **query schedule** command:

```
query schedule
q sc
q sched
query sc
```

### Options

There are two groups of options that you can use with commands:

- **Client options:** The group of options that are set in your client options file (dsm.opt). To override an option in the client options file (dsm.opt), enter the option with a command. For detailed information about client options, see “Client options reference” on page 190.
- **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 190.

### 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 options file (dsm.opt). 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 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 c:\project directory, you would enter:

```
dsmc archive c:\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 c:\project\budget.fin file to c:\newproj\newbudg.fin, enter:

```
dsmc restore c:\project\budget.fin c:\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*:

- 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 c:\home\me, then the *destinationfilespec* would be c:\home\me\mydir in the following command:

```
dsmc restore c:\fs\dir1\ mydir\
```

- 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 c:\home\mydir\ c:\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 c:\home\mydir c:\away\yourdir
```

- When the file specification contains spaces, it must be enclosed in quotes. For example:

```
dsmsc sel "x:\dir one\file1"
```

However, when the filespec is a directory with a trailing backslash, an extra backslash (`\`) must be added. If an extra backslash is not added, the filespec will not be processed correctly, and the operation might cause unexpected results.

The following example is incorrect:

```
dsmsc sel "x:\dir one\"
```

The following example is correct:

```
dsmsc sel "x:\dir one\\"
```

Here is an example of restoring the contents of one directory to another, when both directories contain spaces:

```
dsmsc rest "x:\dir one\\" "x:\dir two\\"
```

- 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 `d:\test` directory, enter:

```
set access backup d:\test\* * *
set access backup d:\test\*\* * *
```

- The maximum number of bytes for a file name and file path combined is 8440. However, the file name itself cannot exceed 256 bytes and the path leading to the file cannot exceed 8184 bytes. Furthermore, directory names (including the directory delimiter) within a path are limited to 256 bytes. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.
- Microsoft Dfs volumes are accessed using the standard UNC names. The following are examples of valid syntax to access MS Dfs volumes:

```
\\Server_Name\Dfs_Root_Name\path
\\Fault_Tolerant_Name\Dfs_Root_Name\path
```

- 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*.
  - 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. Another way to prevent shell expansion from causing you to go over the 20-operand limit is by placing quotation marks around your source filespec expansion characters for restore commands.

**Note:** This has the side effect of causing a no-query restore. See "Filelist" on page 262 for more information.

---

## 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, drives, or directories. 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 53 shows examples of each wildcard.

Table 53. 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
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	abdefrs	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.

---

## 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.
- To back up quoted file specifications within a directory, the quoted file specifications must end with an extra backslash (\). For example, to back up the contents of directories *dir one* and *dir two*, use the following command:

```
dsmc sel "x:\dir one\\" "x:\dir two\\"
```

---

## 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 xii 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.

### Migrating to Unicode-enabled file spaces

See “Autofsrename” on page 202 for information about using the Unicode-enabled client.

### Open file support

If open file support has been configured (see “Configuring Open File Support (OFS)” on page 29), Tivoli Storage Manager performs a snapshot backup or archive of files that are locked (or “in use”) by other applications. The snapshot allows the archive to be taken from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to the file system are not included in the archive. You can set the *fileleveltype* parameter of the *include.fs* option to *dynamic* to specify which drives do not use open file support.

To control an open file support operation with LVSA, you can specify these additional options in your dsm.opt file or as values of the *include.fs* option: *snapshotcachelocation*, *snapshotcachesize*, *snapshotfsidleretries*, *snapshotfsidlewait*, *presnapshotcmd*, *postsnapshotcmd*. See “Include options” on page 280 for more information.

You can also use VSS for open file support. If you use VSS, you do not need to install LVSA.

#### Notes:

1. You can only set the *fileleveltype* option as a value for the *include.fs* option; not as a standalone option in your dsm.opt file.
2. You can use the *include.fs* option to set snapshot options on a per file system basis.
3. Use the *snapshotcachelocation* option to relocate the cache if necessary. You can specify a snapshotcachelocation for a specific drive using the *include.fs* option.
4. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with FAT, FAT32 or NTFS file systems. This support includes SAN-attached volumes that meet these requirements.
5. If the client is unable to create a snapshot, failover to non-OFS backup occurs; the same backup support that would be done if LVSA was not installed.
6. To enable open file support in a cluster environment, all systems in the cluster should have LVSA configured.

For information about Tivoli Storage Manager Open File Support restrictions and issues, search for the **TSM Client v5.2 Open File Support** document under the **Storage Management** product category at the following Web site:

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


## 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 375 for more information.

## Supported Clients

This command is valid for all Windows clients.

## Syntax



## Parameters

*options*

Table 54. Archive command: Related options

Option	Where to use	Page
<i>archmc</i>	Command line only.	193
<i>autofsrename</i>	Client options file (dsm.opt) only.	202
<i>changingretries</i>	Client options file (dsm.opt) or command line.	207
<i>compressalways</i>	Client options file (dsm.opt) or command line.	221
<i>compression</i>	Client options file (dsm.opt) or command line.	222
<i>deletefiles</i>	Command line only.	229
<i>description</i>	Command line only.	230
<i>dirsonly</i>	Command line only.	235
<i>encryptiontype</i>	Client options file (dsm.opt).	250
<i>encryptkey</i>	Client options file (dsm.opt).	251
<i>filelist</i>	Command line only.	262
<i>filesonly</i>	Command line only.	266
<i>postsnapshotcmd</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	327
<i>preservelastaccessdate</i>	Client options file (dsm.opt) or command line.	331
<i>presnapshotcmd</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	336
<i>skipntpermissions</i>	Client options file (dsm.opt) or command line.	365
<i>skipntsecuritycrc</i>	Client options file (dsm.opt) or command line.	366
<i>snapshotcachelocation</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	367
<i>snapshotcachesize</i>	Client options file (dsm.opt) or <i>include.fs</i> option.	368
<i>snapshotfsidleretries</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	369
<i>snapshotfsidlewait</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	371
<i>snapshotroot</i>	Command line only.	375



Table 54. Archive command: Related options (continued)

<i>subdir</i>	Client options file (dsm.opt) or command line.	381
<i>tapeprompt</i>	Client options file (dsm.opt) or command line.	386
<i>v2archive</i>	Command line only.	406

*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. You can specify as many file specifications as available resources or other operating system limits permit. Separate file specifications with a space. You can also use the *filelist* option to process a list of files. The Tivoli Storage Manager client opens the file you specify with this option and processes the list of files within according to the specific command. See “Filelist” on page 262 for more information. See “Maximum file size for operations” on page 63 for the maximum file size for archive processing.

## Examples

**Task** Archive a single file named budget.jan in the c:\plan\proj1 directory.

**Command:** archive c:\plan\proj1\budget.jan

**Task** Archive all files in the c:\plan\proj1 directory with a file extension of .txt.

**Command:** archive c:\plan\proj1\\*.txt

**Task** Archive all files in the c: drive.

**Command:** archive -subdir=yes c:\\*.\*

**Task** Archive all files in the Microsoft Dfs volume, MyDfsVolume. You must specify *subdir=yes* to archive *all* files in the volume.

**Command:** archive \\myserver\mydfsroot\mydfsvolume\\*.\* -subdir=yes

**Task** Assuming that you initiated a snapshot of the C: drive and mounted the snapshot as \\florence\c\$\snapshots\snapshot.0, archive the c:\dir1\sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name C:.

**Command:** dsmc archive c:\dir1\sub1\\* -subdir=yes  
-snapshotroot=\\florence\c\$\snapshots\snapshot.0

---

## Backup ASR

Use the **backup asr** command to generate Automated System Recovery (ASR) files in the `adsm.sys\ASR` staging directory and back them up to the ASR file space on the Tivoli Storage Manager server. This command generates the `asr.sif` file that contains instructions required for Tivoli Storage Manager to restore the files necessary during ASR recovery mode.

For Windows Server 2003, use the **backup systemstate** (in bold) (see `backup systemstate`). For Windows XP ...

**Note:** ASR recovery requires a valid system state backup, and this command does not perform that backup automatically. For Windows Server 2003, use the **backup systemstate** command (see “Backup Systemstate” on page 453). For Windows XP, use the **backup systemobject** command (see “Backup Systemobject” on page 452) to back up all valid system objects.

## Supported Clients

This command is valid for the Windows XP and Windows Server 2003 clients only.

## Syntax

▶▶—Backup ASR—◀◀

## Parameters

There are no parameters for this command.

## Examples

**Task** Back up the ASR files to the Tivoli Storage Manager server.

**Command:** `backup asr`

---

## Backup Certserverdb

The **backup certserverdb** command backs up a Windows XP certificate server database, specifically the server database of client x.509 certificates. The Certificate Server Database must be online to perform a backup. For information on system objects that you must back up together, see “Backing up system objects (Windows XP)” on page 79.

### Supported Clients

This command is valid for Windows XP clients only.

### Syntax

►►—Backup CERTServerdb—————►►

### Parameters

There are no parameters for this command.

### Examples

**Task** Back up the certificate server database.  
**Command:** back certs

---

## Backup Complusb

The **backup complusb** command backs up the Windows XP COM+ database. COM is a system component model that creates non-graphic entities without relying on language features of the entity's particular class. COM+ is the integration of Microsoft Transaction Server (MTS) and Microsoft Message Queuing (MSMQ) with the core operating system.

For information on system objects that you must back up together, see "Backing up system objects (Windows XP)" on page 79.

### Supported Clients

This command is valid for Windows XP operating systems only.

### Syntax

►►—Backup COMPLusb—◄◄

### Parameters

There are no parameters for this command.

### Examples

**Task** Back up the COM+ database.

**Command:** back comp

---

## Backup Eventlog

The **backup eventlog** command backs up the Windows XP event logs.

This command backs up all event logs available to the operating system. The back up of the Windows XP event log will determine which logs will process by enumerating the sub-keys of Windows Registry key:

- HKLM\SYSTEM\CurrentControlSet\Services\Eventlog

This ensures that Tivoli Storage Manager processes all Windows XP event logs.

Tivoli Storage Manager saves the log copies in the adsm.sys\eventlog directory under root directory of the system drive before sending them to the server for backup.

## Supported Clients

This command is valid for Windows XP clients only.

## Syntax

►►—Backup Eventlog—◄◄

## Parameters

There are no parameters for this command.

## Examples

**Task** Back up the entire event log.

**Command:** backup eventlog

**Task** Back up the security event log.

**Command:** backup eventlog

---

## Backup Group

Use the **backup group** command to create and back up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.

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 280 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 262 for more information.

**groupname**

Specifies the fully qualified name of the group containing a list of files. See “Groupname” on page 271 for more information.

**virtualfsname**

Specifies the name of the virtual file space for the group on which you want to perform the operation. The *virtualfsname* parameter cannot be the same as an existing file space name. See “Virtualfsname” on page 409 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 307 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 501 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 534 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 470 for more information.

5. 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 499 for more information.
6. A group backup can be added to a backup set. See “Restoring data from a backupset” on page 122 for more information about backup sets.

## Supported Clients

This command is valid for all Windows clients.

## Syntax

►—Backup GRoup— *options* —►

## Parameters

*options*

Table 55. Backup Group command: Related options

Option	Where to use	Page
<i>filelist</i>	Command line only.	262
<i>groupname</i>	Command line only.	271
<i>mode</i>	Command line only.	307
<i>snapshotproviderfs</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	373
<i>snapshotproviderimage</i>	Client options file (dsm.opt) or with <i>include.image</i> option.	374
<i>virtualfsname</i>	Command line only.	409

## Examples

**Task** Perform a full backup of all the files in the c:\dir1\filelist1 file to the virtual file space name \virtfs containing the group leader group1 file.

**Command:**

```
backup group -filelist=c:\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. These volumes can be formatted FAT, FAT32, NTFS, or unformatted RAW volumes. If a volume is NTFS-formatted, only those blocks used by the file system will be backed up. If you set the *imagegapsize* option to 0, all blocks, including unused blocks at the end of the volume, will be backed up. See “Imagegapsize” on page 275 for more information.

### Notes:

1. The account running the Tivoli Storage Manager client *must* have administrator authority to successfully perform any type of image backup.
2. The API must be installed to use 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.opt. You can also specify the *compression* option with the **backup image** command.

### Offline and online image backup

The traditional *offline* image backup prevents write access to the volume by other system applications during the operation.

If open file support has been configured (see “Configuring Open File Support (OFS)” on page 29, Tivoli Storage Manager performs a snapshot backup or archive of files that are locked (or “in use”) by other applications. See “Performing an image backup” on page 82, for information on how to install and configure the LVSA. To control an open file support operation with LVSA, you can specify these additional options in your dsm.opt file or as values of the *include.fs* option: *snapshotcachelocation*, *snapshotcachesize*, *snapshotfsidleretries*, *snapshotfsidlewait*, *snapshotproviderfs*, *presnapshotcmd*, *postsnapshotcmd*

You can also use VSS for open file support. If you use VSS, you do not need to install LVSA.

**Considerations:** The following considerations apply to offline and online image backups:

- If you create an image of the system drive, you cannot restore it to the original location. Restore of any image requires that the client have an exclusive lock of the volume you are restoring to, so the system drive cannot be restored since Tivoli Storage Manager is unable to lock the system drive. You can restore an image backup of the system drive to an alternate location.
- Because of different system component configurations, the system image not be consistent across components (such as Active Directory). Some of these components can be configured to use different volumes where parts are installed on the system drive and others to non-system volumes.
- We recommend that the Tivoli Storage Manager client program be installed on the system drive. The client cannot restore an image to the same volume where the client program is installed.
- Image backup is only supported on volumes that have a mount point or a drive letter assigned. Tivoli Storage Manager will not back up a volume without a mount point or drive letter.



- If bad disk sectors are detected on the source drive during a LAN-free or LAN-based image backup, data corruption occur. In this case, bad sectors are skipped when sending image data to the Tivoli Storage Manager server. If bad disk sectors are detected during the image backup, a warning message is issued after the image backup completes.

## 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 h:
```

2. Perform an image backup of the same file system, for example:

```
dsmc backup image h:
```

3. Periodically, perform incremental backups, for example:

```
dsmc incremental h:
```

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 h: -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 h:
```

2. Perform an incremental image backup of the file system, for example:

```
dsmc backup image h: -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 307.

3. Periodically, perform full image backups, for example:

```
dsmc backup image h:
```

4. Restore the image as follows:

```
dsmc restore image h: -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:

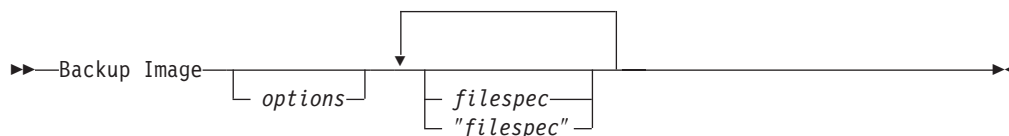
- The file system can have no previous full incremental backups produced by the **incremental** command.
- Incremental-by-date image backup does not inactivate files on the server; therefore, when files are restored, none can be deleted.
- If this is the first image backup for the file system, a full 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 85.

## Supported Clients

This command is valid for all Windows platforms.

## Syntax



## Parameters

*options*

Table 56. Backup Image command: Related options

Option	Where to use	Page
<i>compressalways</i>	Client options file (dsm.opt) or command line.	221
<i>compression</i>	Client options file (dsm.opt) or command line.	222
<i>compression</i>	Client options file (dsm.opt) or command line.	222
<i>compression</i>	Client options file (dsm.opt) or command line.	222
<i>imagegapsize</i>	Use with the <b>backup image</b> command, or the <i>include.image</i> option in your client options file (dsm.opt).	275
<i>mode</i>	Command line only.	307
<i>postsnapshotcmd</i>	Use with the <b>backup image</b> command, the <i>include.image</i> option, or in the dsm.opt file.	327
<i>presnapshotcmd</i>	Use with the <b>backup image</b> command, the <i>include.image</i> option, or in the dsm.opt file.	336
<i>snapshotcachelocation</i>	Use with the <b>backup image</b> command, the <i>include.image</i> option, or in the dsm.opt file.	367
<i>snapshotcachesize</i>	Use with the <b>backup image</b> command, the <i>include.image</i> option, or in the dsm.opt file.	368
<i>snapshotfsidleretries</i>	Use with the <b>backup image</b> command, the <i>include.image</i> option, or in the dsm.opt file.	369
<i>snapshotfsidlewait</i>	Use with the <b>backup image</b> command, the <i>include.image</i> option, or in the dsm.opt file.	371
<i>snapshotproviderimage</i>	Client options file (dsm.opt) or with <i>include.image</i> option.	374

*filespec*

Specifies the name of one or more logical volumes. *Image backup is only supported on a volume that has a mount point assigned or a drive letter assigned. A volume without a drive letter or mount point cannot be backed up.* 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.

## Examples

**Task** Back up a volume that has no drive letter but is mounted as a mount point.

**Command:** dsmc backup image m:\mnt\myntfs

**Task** Back up the h: drive using an image incremental backup that backs up only new and changed files after the last full image backup.

**Command:** dsmc backup image h: -mode=incremental

**Task** Perform an offline image backup of the f: drive.

**Command:** dsmc backup image f: -imagetype=static

**Task** Perform an online image backup of the f: drive.

**Command:**

```
dsmc backup image f: -imagetype=snapshot
```

**Task** Back up the f: drive which is mapped to a volume which has not been formatted with a file system.

**Command:** dsmc backup image f:

---

## 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 options file (dsm.opt). The value in the client options file is the default, but can be overridden on the command line. See “Nasnodename” on page 311 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 399 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 307 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 309.

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 460.

A NAS file system specification uses the following conventions:

- 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 all Windows clients.

## Syntax



## Parameters

*options*

Table 57. Backup NAS command: Related options

Option	Where to use	Page
<i>mode</i>	Command line only.	307
<i>monitor</i>	Command line only.	309
<i>nasnodename</i>	Client options file (dsm.opt) or command line.	311
<i>toc</i>	Command line or with the <i>include.fs.nas</i> option in your client options file (dsm.opt).	399

*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 242.

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 Registry

The **backup registry** command backs up the Windows XP registry. The registry is a centralized database that contains information about hardware, applications, and operating system settings for each workstation on the network. This command is processed automatically as part of the SYSTEMOBJECT domain.

Use this command if you must back up the Windows Registry immediately.

**Attention:** Your system might appear to stop when using the **backup registry** command because of the Microsoft RegSaveKey() API.

Registry backups are copied to the adm.sys directory under the root directory of the system drive. The registry is then copied to the server.

For information on system objects that must be backed up together, see “Backing up system objects (Windows XP)” on page 79.

### Supported Clients

This command is valid for Windows XP clients.

### Syntax

►►—Backup REgistry—◄◄

### Parameters

There are no parameters for this command. The **backup registry** command backs up the system registry hives listed in:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\  
Control\hivelist
```

### Examples

**Task** Backup the entire registry.

**Command:** backup registry



---

## Backup Sysfiles

The **backup sysfiles** command backs up Windows XP system and boot files.

System and boot files consist of the following:

- All files that are protected by system file protection
- All files that are in the system file protection service catalog directory
- Performance counter configuration files
- Internet Information Services (IIS) metabase

These files are backed up as a single system object; individual files within this object are backed up as normal files. On Windows XP this command is processed automatically as part of the SYSTEMOBJECT domain. The **backup sysfiles** command does not back up these boot files on 64 bit Windows XP: boot.ini, NTDETECT.COM, ntldr, BOOTSECT.DOS, Ntbootdd.sys.

For information on system objects that must be backed up together, see “Backing up system objects (Windows XP)” on page 79.

## Supported Clients

This command is valid for Windows XP clients.

## Syntax

►►—Backup SYSFiles—◄◄

## Parameters

There are no parameters for this command.

## Examples

**Task** Back up system and boot files.

**Command:** back sysf

---

## Backup Systemobject

The **backup systemobject** command backs up all valid system objects. The following system objects which are active on the system are processed:

- Certificate Server Database
- Cluster Database (cluster node only)
- COM+ database
- Windows Registry
- System and boot files
- System Volume
- Event logs (system, security and application)
- Removable Storage Management Database (RSM)
- Replicated file systems
- Windows Management Instrumentation (WMI)

This command is processed automatically as part of the SYSTEMOBJECT domain.

**Attention:** Your system might appear to stop when using the **backup systemobject** command because of the Microsoft RegSaveKey() API.

**Note:** Include-exclude statements do not affect system object processing. If you accidentally exclude a directory that is critical to a system object backup, the system object backup is not affected.

## Supported Clients

This command is valid for Windows XP clients.

## Syntax

►►—Backup SYSTEMObject—◄◄

## Parameters

There are no parameters for this command.

## Examples

**Task** Back up all valid system objects.

**Command:** backup systemobject

---

## Backup Systemstate

Use the **backup systemstate** command to back up all bootable system state and system services components as a single object, to provide a consistent point-in-time snapshot of the system state.

Bootable system state components can include the following:

- Active Directory (domain controller only)
- System volume (domain controller only)
- Certificate Server Database
- COM+ database
- Windows Registry
- System and boot files
- ASR (Windows 2003 only)

System services components can include the following:

- Background Intelligent Transfer Service (BITS)
- Event logs
- Removable Storage Management Database (RSM)
- Cluster Database (cluster node only)
- Remote Storage Service
- Terminal Server Licensing
- Windows Management Instrumentation (WMI)
- Internet Information Services (IIS) metabase
- DHCP database
- Wins database

The list of bootable system state and system services components is dynamic and can change depending on service pack and operating system features installed. Tivoli Storage Manager allows for the dynamic discovery and back up of these components.

### Notes:

1. To back up Windows system state, your client must be connected to a Tivoli Storage Manager Version 5.2.0 or higher server.
2. The system and boot files component of system state is backed up only if a member (file) of that component has changed since the last backup. If a member changes, the entire group of files that comprise that component are backed up.
3. The Tivoli Storage Manager Windows client does not allow the backup of any individual component.
4. Use the *include.systemstate* option in your client options file (dsm.opt) to assign management classes for backup of system state. The default is to bind the system state object to the default management class.
5. Use the **query systemstate** command to display information about a backup of the system state on the Tivoli Storage Manager server. See “Query Systemstate” on page 517 for more information.
6. Use the **restore systemstate** command to restore the backup of the system state from the Tivoli Storage Manager server. See “Restore Systemstate” on page 547 for more information.

## Supported Clients

This command is valid for the Windows Server 2003 and Windows Vista clients only.

## Syntax

▶▶ Backup SYSTEMState ◀◀

## Parameters

There are no parameters for this command.

## Examples

**Task** Back up the system state.

**Command:** backup systemstate

## Backup WAS

The **backup was** command specifies whether to back up the WebSphere Application Server 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 Websphere Application Server 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 Websphere Application Server on your system. You only need to perform this task once, unless you change your Websphere Application Server user name or password. See “Set Waspasword” on page 569 for more information.

To determine if Websphere Application Server security is enabled, enter the following command:

```
dsmc query was -wast=local
```

Tivoli Storage Manager displays the Websphere Application Server 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 307 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 Websphere Application Server that you want to back up. The default is ND. See “Wastype” on page 420 for more information.
5. Use the *include* option in your client options file (dsm.opt) to assign a management class to a Websphere Application Server group backup. For example:
  - For the Network Deployment Manager: `include /WAS_ND_NDNODE mgmtclass`
  - For the Application Server: `include WAS_APPNODE mgmtclass`
6. Websphere Application Server backups can also be added to a backup set. See “Restoring data from a backupset” on page 122 for more information about Websphere Application Server backups.

## Supported Clients

This command is valid for Windows Server 2003.

## Syntax

```
► Backup WAS [ options ] —nodename—►
```

## Parameters

*options*

Table 58. Backup WAS command: Related options

Option	Where to use	Page
--------	--------------	------

Table 58. Backup WAS command: Related options (continued)

<i>mode</i>	Command line only.	307
<i>wastype</i>	Command line only.	420

*nodename*

Specifies the node name of the instance of Websphere Application Server 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

## Backup VM

The **backup VM** command backs up virtual machines specified in the *VMLIST* option. The **backup VM** command is used to back up VMware virtual machines from the VMware Consolidated Backup proxy system.

One or more virtual machines are processed using VMware Consolidated Backup snapshot backups from the backup proxy, but the backup files are stored as if they were backed up from within the virtual machines. Before using this command, specific setup steps are required, which are documented in “Using VMware Consolidated Backup” on page 92.

## Supported Clients

This command is valid for Windows clients that are configured as a VMware Consolidated Backup proxy.

## Syntax

►► Backup VM [options]

## Parameters

*options*

Table 59. Backup VM command: Related options

Option	Where to use	Page
<i>VMCHost</i>	Command line or dsm.opt	412
<i>VMCPW</i>	Command line or dsm.opt	413
<i>VMCUSER</i>	Command line or dsm.opt	414
<i>VMLIST</i>	Command line or dsm.opt	415

## Examples

**Task** Back up virtual machine *vm1.ibm.com* using VMware VirtualCenter system *virctr.ibm.com* to the Tivoli Storage Manager Server, using system name *vm1*.

**Command:** `dsmc backup vm -vmlist=vm1 -vmchost=virctr -vmcuser=virctr\admin -vmcpw=xxxxx`

**Task** Back up virtual machine *vm1.ibm.com* and *vm2.ibm.com* using VMware ESX Server system *vmserve1.ibm.com* to the Tivoli Storage Manager Server, using system name *vm1* and *vm2*.

**Command:** `dsmc backup vm -vmlist=vm1,vm2 -vmchost=vmserve1 -vmcuser=root -vmcpw=xxxxx`

**Task** Back up virtual machines in an environment where DNS name resolution is not available.

**Command:** `dsmc backup vm -vmlist=9.100.45.101  
[vm1],9.100.45.102  
[vm1] -vmchost=9.100.45.100 -vmcuser=virctr\admin  
-vmcpw=xxxxx`

| **Task** Back up virtual machines from several ESX servers in an environment  
| where a virtual center server is not available.

| **Command:** dsmc backup vm -vmlist=vm1,vm2,vm3  
| -vmchost=esxserver1 -vmcuser=root

| **Command:** dsmc backup vm -vmlist=vm4,vm5,vm6  
| -vmchost=esxserver2 -vmcuser=root

| **Command:** dsmc backup vm -vmlist=vm7,vm8,vm9  
| -vmchost=esxserver3 -vmcuser=root



---

## Backup WMI

The **backup wmi** command backs up a Windows XP Windows Management Instrumentation (WMI) repository. You cannot back up a WMI repository to a file specification. Tivoli Storage Manager exports the WMI repository to `adsm.sys\wmi\wmidbfile` file and backs up the exported file.

### Supported Clients

This command is valid for Windows XP only.

### Syntax

▶▶—Backup WMI—▶▶

### Parameters

There are no parameters for this command.

### Examples

**Task** Back up the Windows Management Instrumentation repository.

**Command:** `back wmi`

---

## 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 all Windows clients.

## 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 select zero or more restartable restore sessions that you want to cancel. 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.
- You want to cancel restartable restore sessions.
- 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 Windows 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 Windows 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	daisy	c:\dev\proja\list.c
2	Archive	node3	marm	c:\fin\budg\depta.jan
3	Backup	node4	susie	c:\plan\exp\deptc.feb
4	Archive	node5	susies	c:\mfg\invn\parta.wip

Enter Index of rule(s) to delete, or quit to cancel:

To delete the authorization rules that allow *marm* and *susies* to access your files, type **2 4** or **(2,4)**, then 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 Windows clients.

## Syntax

```
▶▶ Delete ARchive [options] [filespec] ▶▶
```

└── options ─┘ └── {filespace}filespec ─┘

## Parameters

*options*

Table 60. Delete Archive command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client options file (dsm.opt) or command line.	227
<i>description</i>	Command line only.	230
<i>filelist</i>	Command line only.	262
<i>noprompt</i>	Command line only.	315
<i>numberformat</i>	Client options file (dsm.opt) or command line.	316
<i>pick</i>	Command line only.	322
<i>subdir</i>	Client options file (dsm.opt) or command line.	381
<i>tapeprompt</i>	Client options file (dsm.opt) or command line.	386
<i>timeformat</i>	Client options file (dsm.opt) or command line.	397

*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. You can specify as many file specifications as available resources or other operating system limits permit. Separate file specifications with a space. You can also use the *filelist* option to process a list of files. The Tivoli Storage Manager client opens the file you specify with this option and processes the list of files within according to the specific command. See “Filelist” on page 262 for more information.

**Note:** If you indicate *filespace*, do not include a drive letter in the file specification.

*{filespace}*

Specifies the file space (enclosed in braces) on the server that contains the file you want to delete. This is the name on the workstation drive from which the file was archived. You can specify UNC names; drive label names are only used for removable media.

Use the *filespace* if the name has changed, or if you are deleting files archived from another node with drive labels that are different from yours. You must specify a mixed or lowercase NTFS file space name enclosed in quotes and braces. For example, {"NTFSDrive"}.

Single quotes or quotation marks are valid in loop mode. For example, {"NTFSDrive"} and {'NTFSDrive'} are both valid. In batch mode, only single quotes are valid. The single quotes requirement is a restriction of the operating system.

## Examples

**Task** Delete files from file space abc in the proj directory.

**Command:** delete archive {abc}\proj\\*.\*

**Task** Delete a file named budget.jan.

**Command:** delete archive c:\plan\proj1\budget.jan

**Task** Delete all files archived from the c:\plan\proj1 directory with a file extension of .txt.

**Command:** delete archive c:\plan\proj1\\*.txt

**Task** Delete files archived from the c:\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 c:\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 Windows clients, however *objtype=image* is not supported on z/OS.

### Syntax



### Parameters

*options*

Table 61. Delete Backup command: Related options

Option	Where to use	Page
<i>description</i>	Command line only.	230
<i>filelist</i>	Command line only.	262
<i>fromdate</i>	Command line, and in GUI find function.	262
<i>frontime</i>	Command line, and in GUI find function.	262
<i>noprompt</i>	Command line only.	315
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line, and in GUI find function.	262
<i>pittime</i>	Command line, and in GUI find function.	262
<i>subdir</i>	Client options file (dsm.opt) or command line.	381
<i>tapeprompt</i>	Client options file (dsm.opt) or command line.	386
<i>timeformat</i>	Client options file (dsm.opt) or command line.	397
<i>todate</i>	Command line, and in GUI find function.	262
<i>totime</i>	Command line, and in GUI find function.	262

*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.

You can specify as many file specifications as available resources or other operating system limits permit. Separate file specifications with a space. You can also use the *filelist* option to process a list of files. The Tivoli Storage Manager client opens the file you specify with this option and processes the list of files within according to the specific command. See "Filelist" on page 262 for more information.

**Note:** If you indicate *filespace*, do not include a drive letter in the file specification.

### *{filespace}*

Specifies the file space (enclosed in braces) on the server that contains the file you want to delete. This is the name on the workstation drive from which the file was backed up. You can specify UNC names; drive label names are only used for removable media.

Use the *filespace* if the name has changed, or if you are deleting files backed up from another node with drive labels that are different from yours. You must specify a mixed or lowercase NTFS file space name enclosed in quotes and braces. For example, {"NTFSDrive"}.

Single quotes or quotation marks are valid in loop mode. For example, {"NTFSDrive"} and {'NTFSDrive'} are both valid. In batch mode, only single quotes are valid. The single quotes requirement is a restriction of the operating system.

### *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 file objects from file space abc in the proj directory.

**Command:** delete backup {abc}\proj\\*

**Task** Delete all active and inactive file objects named budget.jan in directory c:\plan\proj1.

**Commands:**

```
delete backup c:\plan\proj1\budget.jan -deltype=inactive
delete backup c:\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 c:\plan\proj1 directory and its subdirectories.

**Command:** delete backup c:\plan\proj1\\*.txt -deltype=inactive -subdir=yes

**Task** Delete selected active files backed up from the c:\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 c:\project\\* -pick

**Task** Delete all active and inactive versions of files and subdirectories in c:\user\myproject.

**Command:** delete backup c:\user\myproject\\* -deltype=all

**Note:** The backup versions of directory object c:\user\myproject are not deleted.

---

## Delete Filespace

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 drive from which you back up or archive files. The file space name is the same as the UNC name. Note that drive label names are only used for removable media. 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 options file (dsm.opt). The value in the client 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 options file, you must specify this option on the command line when processing NAS file systems. See “Nasnodename” on page 311 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 208 for more information.

To delete NAS file spaces using the Web client, see Chapter 4, “Backing up your data,” on page 57.

### Deleting WebSphere Application Server file spaces

Use the **delete filesystem** command to delete a Websphere Application Server file space on the Tivoli Storage Manager server.


Use the **delete group** command to delete Websphere Application Server group backups on the Tivoli Storage Manager server. See “Delete Group” on page 470 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 Windows clients.

## Syntax



## Parameters

*options*

Table 62. Delete Filespace command: Related options

Option	Where to use	Page
<i>class</i>	Command line only.	208
<i>detail</i>	Command line only.	232
<i>nasnodename</i>	Client options file (dsm.opt) or command line.	311
<i>scrolllines</i>	Client options file (dsm.opt) or command line.	358
<i>scrollprompt</i>	Client options file (dsm.opt) or command line.	359

## 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 Websphere Application Server file spaces stored on the server.

**Command:** delete filesystem

## Delete Group

Use the **delete group** command to delete a group backup on the Tivoli Storage Manager server. You can also delete Websphere Application Server 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 277 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 322 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 315 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 499 for more information.

## Supported Clients

This command is valid for all Windows 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 63. Delete Group command: Related options

Option	Where to use	Page
<i>inactive</i>	Command line only.	277
<i>noprompt</i>	Command line only.	315
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324

## Examples

**Task** Delete the current active version of the virtfs\group1 group.

**Command:**

```
delete group {virtfs}\group1
```

**Task** Delete a backup version of the virtfs\group1 group from a list of active and inactive versions.

**Command:**

```
delete group {virtfs}\group1 -inactive -pick
```

**Task** Delete the Websphere Application Server group backup from a list of active and inactive versions.

**Command:** delete group {WAS\_NODENAME}\WASGROUP -ina -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 journal service.

## Supported Clients

This command is valid for all Windows clients.

## Syntax

```
EXPire [options] filespec
```

## Parameters

*options*

Table 64. Expire command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client options file (dsm.opt) or command line.	227
<i>filelist</i>	Command line only.	262
<i>noprompt</i>	Command line only.	315
<i>numberformat</i>	Client options file (dsm.opt) or command line.	316
<i>pick</i>	Command line only.	322
<i>timeformat</i>	Client options file (dsm.opt) or command line.	397

**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 c:\home\letter1.txt`
- Task** Inactivate all files in the admin\mydir directory.  
**Command:** `expire c:\admin\mydir\*`
- Task** Inactivate all files named in the c:\avi\filelist.txt file.  
**Command:** `expire -filelist=c:\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 Windows 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.

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 239 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 153.

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 476 for more information. You can perform a full incremental backup without the journal by using the *nojournal* option. See “Nojournal” on page 314 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 559.

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 or the journal database. 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.

## Migrating to Unicode-enabled file spaces

See “Autofsrename” on page 202 for information about using the Unicode-enabled client.

## Open file support

If open file support has been configured (see “Configuring Open File Support (OFS)” on page 29, Tivoli Storage Manager performs a snapshot backup or archive of files that are locked (or “in use”) by other applications. See “Performing an image backup” on page 82, for information on how to install and configure the LVSA. To control an open file support operation with LVSA, you can specify these additional options in your dsm.opt file or as values of the *include.fs* option: *snapshotcachelocation*, *snapshotcachesize*, *snapshotfsidleretries*, *snapshotfsidlewait*, *snapshotproviderfs*, *presnapshotcmd*, *postsnapshotcmd*

To control an open file support operation with LVSA, you can specify these additional options in your dsm.opt file or as values of the *include.fs* option: *snapshotcachelocation*, *snapshotcachesize*, *snapshotfsidleretries*, *snapshotfsidlewait*, *presnapshotcmd*, *postsnapshotcmd*. See “Include options” on page 280 for more information.

### Notes:

1. You can use the *include.fs* option to set snapshot options on a per file system basis.
2. Use the *snapshotcachelocation* option to relocate the cache if necessary. You can specify a *snapshotcachelocation* for a specific drive using the *include.fs* option.
3. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with FAT, FAT32 or NTFS file systems. This support includes SAN-attached volumes that meet these requirements.
4. If the client is unable to create a snapshot, failover to non-OFS backup occurs; the same backup support that would be done if the OFS feature was not configured.
5. To enable open file support in a cluster environment all systems in the cluster should have the OFS feature configured.

For information about Tivoli Storage Manager Open File Support restrictions and issues, search for the **TSM Client v5.2 Open File Support** document under the **Storage Management** product category at the following Web site:

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

## Journal-based backup

If the *journal engine service* 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 service. Tivoli Storage Manager does not use the journaling facility inherent in Windows NTFS file systems or any other journaled file system.

The *journal engine service* 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. Performing backups on a regular basis maintains the size of the journal.

Journal-based backup can increase backup performance. With journal-based backup, the client does not scan the local file system or obtain information from 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. Tivoli Storage Manager processes, expires, and updates the resulting files according to policy constraints, such as serialization. The management-class copy frequency attribute is ignored during journal-based backup.

The journal engine service excludes specific system files (pagefile, registry, etc.) 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 service configuration file `tsmjbbd.ini` located in the Tivoli Storage Manager installation directory for specific system files that are excluded.

To support journal-based backup, you must install the journaling engine service. Install this service by using the **dsmcutil** command or the GUI Setup wizard. See “Using the **Dsmcutil** command” on page 589, for more information about using the **dsmcutil** command to install the journaling engine service. See “Journal-based backup” on page 66 for more information on how to install the journaling engine service using the GUI Setup wizard.

If the file specification on the **incremental** command is a file space, Tivoli Storage Manager processes any journal entries for that file space. Tivoli Storage Manager processes directories and wildcarded file specifications in the same way. Tivoli Storage Manager uses the domain list if you do not specify a file specification.

**Note:** Journal-based backup might not fall back to the traditional incremental backup if the policy domain of your node is changed on the server, depending on when the policy set within the domain was last updated and the date of the last incremental. 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.

When a user deletes a file with a long name, the Windows operating system might supply a short, or compressed, name to the journal engine service. After the object is deleted, the compressed name can be reused and the deletion notice might no longer identify a unique object. During a journaled incremental backup, the attempt to expire the file will fail because the compressed name is not known to the server. When this occurs, a record is placed in the journal indicating that the current directory is not exactly represented at the server. Use the *incrthreshold* option to specify what action is taken when this occurs. See “Incrthreshold” on page 288 for more information.

Beginning with Tivoli Storage Manager Version 5.3, multiple journal-based backup sessions are possible.

**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 service is not running.
- The journal service is stopped or started for any reason, including system reboot.

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).
- Attribute changes to an object require a back up of the entire object.

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 314 for more information.

### Virtual volume and virtual mount point (Windows NTFS)

If you perform an incremental backup of a file space on which a volume is mounted, Tivoli Storage Manager does not traverse the mount junction to back up all mounted data. For example, if you have a mounted volume named c:\mount on file space c:, Tivoli Storage Manager backs up only the junction point, not the data under c:\mount.

To back up all data for a file system, including virtual volumes and virtual mount points, run an incremental backup on each virtual volume. You must also perform an incremental backup on the file system.

The *exclude.dir* option is *not* valid for a mounted virtual volume if the volume is backed up as a file space. *Exclude.dir* is valid for a mount point that is backed up.

### Backing up Microsoft Dfs root

If you perform an incremental backup of Microsoft Dfs root with *dfsbackupmntpnt=yes* specified, Tivoli Storage Manager backs up only the junction points, *not* the subtree under the junctions.

If you want to traverse the Dfs tree and back up the files and subdirectories of any junction it encounters, specify the *dfsbackupmntpnt=no* option. If you want to backup both the Dfs tree structure and the data contained in the Dfs tree you must run two backups: one with *dfsbackupmntpnt=yes* and one with *dfsbackupmntpnt=no*.

This option has no effect if you are backing up individual junctions. The *exclude.dir* option behavior for Dfs junctions is same as for mounted virtual volumes. For more information, see “Dfsbackupmntpnt” on page 233.

**Note:** If a Dfs root is added or modified, Tivoli Storage Manager will not back it up. You must specify the Dfs root in the *domain* option in the client options file (dsm.opt) regardless of whether DOMAIN ALL-LOCAL is specified.

## 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. 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.

## Adaptive subfile backups

If you plan to perform a backup over a network device with limited bandwidth, such as a modem, you can reduce the network traffic by using the *subfilebackup* option. If you plan to use this option, ensure that you specify the *subfilebackup*, *subfilecachepath* and *subfilecachesize* options during the initial backup of your file. The options specify whether to perform an adaptive subfile backup along with the path and size of the data to be transferred. For information on these options, see Chapter 9, “Using processing options,” on page 165. For information about adaptive subfile backups, see “Performing a backup with limited bandwidth” on page 61.

## 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 375 for more information.

## Considerations

Tivoli Storage Manager backs up named streams on a file basis only. Backup of a named stream containing sparse file data is not supported. Tivoli Storage Manager will back up a sparse file as a regular file if the client compression is off. Enable

file compression (*compression=yes*) when backing up sparse files to minimize network transaction time and to maximize server storage space.

Microsoft disk quotas do not affect the amount of data you can back up.

Tivoli Storage Manager backs up EFS encrypted files in raw encrypted format. The files are not decrypted prior to transmission; they are always in encrypted format. The files are restored in encrypted format as well.

## Supported Clients

This command is valid for all Windows clients.

## Syntax



## Parameters

*options*

Table 65. Incremental command: Related options

Option	Where to use	Page
<i>autofsrename</i>	Client options file (dsm.opt) only.	202
<i>changingretries</i>	Client options file (dsm.opt) or command line.	207
<i>compressalways</i>	Client options file (dsm.opt) or command line.	221
<i>compression</i>	Client options file (dsm.opt) or command line.	222
<i>dironly</i>	Command line only.	235
<i>domain</i>	Client options file (dsm.opt) or command line only.	239
<i>encryptiontype</i>	Client options file (dsm.opt).	250
<i>encryptkey</i>	Client options file (dsm.opt).	251
<i>filelist</i>	Command line only.	262
<i>filesonly</i>	Command line only.	266
<i>incrbydate</i>	Command line only.	286
<i>memoryefficientbackup</i>	Client user options file (dsm.opt), server, or command line.	305
<i>nojournal</i>	Command line only.	314
<i>postsnapshotcmd</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	327
<i>preservelastaccessdate</i>	Client options file (dsm.opt) or command line.	331
<i>presnapshotcmd</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	336
<i>resetarchiveattribute</i>	Client options file (dsm.opt).	342
<i>skipntpermissions</i>	Client options file (dsm.opt) or command line.	365
<i>skipntsecuritycrc</i>	Client options file (dsm.opt) or command line.	366
<i>snapshotcachelocation</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	367
<i>snapshotcachesize</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	368

Table 65. Incremental command: Related options (continued)

<i>snapshotfsidleretries</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	369
<i>snapshotfsidlewait</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	371
<i>snapshotproviderfs</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	373
<i>snapshotproviderimage</i>	Client options file (dsm.opt) or with <i>include.image</i> option.	374
<i>snapshotroot</i>	Command line only.	375
<i>subdir</i>	Client options file (dsm.opt) or command line.	381
<i>tapeprompt</i>	Client options file (dsm.opt) or command line.	386

### *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. You can specify as many file specifications as available resources or other operating system limits permit. Separate file specifications with a space. You can also use the *filelist* option to process a list of files. The Tivoli Storage Manager client opens the file you specify with this option and processes the list of files within according to the specific command. See “Filelist” on page 262 for more information. 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 63 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 options file (dsm.opt).

**Command:** Incremental

**Task** Run an incremental backup of the c:, d:, and e: drives.

**Command:** i c: d: e:

**Task** Run an incremental backup of the \home\ngai directory and its contents on the current drive.

**Command:** i \home\ngai\

**Task** Assuming that you initiated a snapshot of the C: drive and mounted the snapshot as \\florence\c\$\snapshots\snapshot.0, 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 c:.

**Command:** dsmc inc c: -snapshotroot=\\florence\c\$\snapshots\snapshot.0

---

## 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 options file (`dsm.opt`). 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 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 165 for options that you cannot use in interactive mode.

## Supported Clients

This command is valid for all Windows 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 Windows 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:

c:\devel\project\proja  
c:\devel\project\projb  
c:\devel\project\projc

**Command:** macro backabc.mac

where backabc.mac contains the following statements:

```
selective c:\devel\project\proja\*.*  
selective c:\devel\project\projb\*.*  
selective c:\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 all Windows clients.

## 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 Windows 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-delineated output. The default is `dsmprev.txt`.

### -console

Output is written to the console, and the file.

### -traverse

Preview the current directory and subdirectories.

**Yes** Preview the current directories and subdirectories. This is the default.

**No** Preview only the current directory, not subdirectories.

**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 files to which the user has access.

## Supported Clients

This command is valid for all Windows 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
- Creation date

## Supported Clients

This command is valid for all Windows clients.

## Syntax

```
►► Query Archive [options] [filespec]
                                     {filespace}filespec
```

## Parameters

*options*

Table 66. Query Archive command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client options file (dsm.opt) or command line.	227
<i>description</i>	Command line only.	230
<i>detail</i>	Command line only.	232
<i>dironly</i>	Command line only.	235
<i>filelist</i>	Command line only.	262
<i>filesonly</i>	Command line only.	266
<i>fromdate</i>	Command line only.	267
<i>fromnode</i>	Command line only.	268
<i>fromtime</i>	Command line only.	269
<i>numberformat</i>	Client options file (dsm.opt) or command line.	316
<i>scrolllines</i>	Client options file (dsm.opt) or command line.	358
<i>scrollprompt</i>	Client options file (dsm.opt) or command line.	359
<i>subdir</i>	Client options file (dsm.opt) or command line.	381
<i>timeformat</i>	Client options file (dsm.opt) or command line.	397
<i>to date</i>	Command line only.	401
<i>to time</i>	Command line only.	402

*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.

**Note:** If you include *filespace*, do not include a drive letter in the file specification. Note that drive label names are only used for removable media.

*{filespace}*

Specifies the file space (enclosed in braces) on the server that contains the file you want to query. This is the name on the workstation drive from which the file was archived. The following example is valid for specifying a UNC name: {'\\machine\C\$'}.

Use the *filespace* if the name has changed or if you are querying files archived from another node with drive labels that are different from yours.

**Note:** You must specify a mixed or lowercase NTFS *filespace* enclosed in quotes and braces. For example, {"NTFSDrive"}. Single quotes or quotation marks are valid in loop mode. For example: {'NTFSDrive'} and {'NTFSDrive'} are both valid. In batch mode, only single quotes are valid. The single quotes requirement is a restriction of the operating system.

## Examples

- Task** Display a list of all your archived files in the c:\proj directory.  
**Command:** q ar c:\proj\\*
- Task** Display a list of archived files from your c: drive with the description "January Ledgers."  
**Command:** query archive c:\ -su=y -descr="January Ledgers"
- Task** Display a list of all your archived files in the c:\proj directory. Use the *dateformat* and *timeformat* options to reformat the dates and times.  
**Command:** q ar -date=5 -time=4 c:\proj\\*
- Task** Display a list of all your archived files in the c:\dir1 directory. Use the *detail* option to display the last modification date and the creation date of each file.  
**Command:** q ar -detail c:\dir1\\*
- Task** Display a list of archived files in the c:\proj directory containing a file extension of .dev. Use the *dateformat* and *timeformat* options.  
**Command:** q ar -date=5 -time=4 c:\proj\\*.dev
- Task** Last week you changed the label of your c: drive to **store** and archived some files. Yesterday you changed the label to **dev** and archived more files. Display a list of all the files you archived in the c:\proj directory when the label was **store**.  
**Command:** q ar {store}\proj\\*
- Task** Last week you archived files from a diskette labeled **docs**. Display a list of all the files you archived.  
**Command:** q ar {docs}\\*

---

## Query ASR

The **query ASR** command displays information about the active backup of the Automated System Recovery (ASR) files on the Tivoli Storage Manager server.

### Supported Clients

This command is valid for the Windows XP and Windows 2003 operating systems.

### Syntax

►►—Query ASR— [ *options* ] ◄◄

### Parameters

*options*

Table 67. Query ASR command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

### Examples

**Task** Display information about the active backup of the ASR object on the server.

**Command:** query asr



## 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
- Creation date

## Supported Clients

This command is valid for all Windows clients.

## Syntax

```
►► Query Backup [ options ] [ filespec ]
                                     [ {filespace}filespec ]
```

## Parameters

*options*

Table 68. Query Backup command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	204
<i>class</i>	Command line only.	208
<i>dateformat</i>	Client options file (dsm.opt) or command line.	227
<i>detail</i>	Command line only.	232
<i>dironly</i>	Command line only.	235
<i>filelist</i>	Command line only.	262
<i>filesonly</i>	Command line only.	266
<i>fromdate</i>	Command line only.	267
<i>fromnode</i>	Command line only.	268
<i>fromtime</i>	Command line only.	269
<i>inactive</i>	Command line only.	277
<i>location</i>	Command line only.	300
<i>nasnodename</i>	Client options file (dsm.opt) or command line.	311
<i>numberformat</i>	Client options file (dsm.opt) or command line.	316
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>scrolllines</i>	Client options file (dsm.opt) or command line.	358

Table 68. Query Backup command: Related options (continued)

<i>scrollprompt</i>	Client options file (dsm.opt) or command line.	359
<i>subdir</i>	Client options file (dsm.opt) or command line.	381
<i>timeformat</i>	Client options file (dsm.opt) or command line.	397
<i>today</i>	Command line only.	401
<i>totime</i>	Command line only.	402

### *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. Do not use wild cards when you query NAS file system images with *-class=nas* option.

**Note:** If you include *filespace*, do not include a drive letter in the file specification. Drive label names are only used for removable media.

### *{filespace}*

Specifies the file space (enclosed in braces) on the server that contains the file you want to query. This is the drive label or UNC name on the workstation drive from which the file was backed up. The following example is valid for specifying a UNC name: {'\\machine\C\$'}.

Use the *filespace* if the name has changed, or if you want to query files backed up from another node with drive label names that are different from yours.

**Note:** You must specify a mixed or lowercase NTFS file space name enclosed in quotes and braces. For example, {"NTFSDrive"}. Single quotes or quotation marks are valid in loop mode. For example: {"NTFSDrive"} and {'NTFSDrive'} are both valid. In batch mode, only single quotes are valid. The single quotes requirement is a restriction of the operating system.

## Examples

**Task** Query files from the abc file space proj directory.

**Command:** q b {"abc"}\proj\\*.\*

**Task** Display a list of all active and inactive backup versions that were backed up from the c:\proj directory.

**Command:** q backup -ina c:\proj\\*

**Task** Display a list of all your backups in the c:\dir1 directory. Use the *detail* option to display the last modification date and the creation date of each file.

**Command:** q backup -detail c:\dir1\\*

**Task** Display a list of all active and inactive backup versions that were backed up from the c:\proj directory. Use the *dateformat* and *timeformat* options to reformat the dates and times.

**Command:** q b -date=5 -time=4 -ina c:\proj\\*

**Task** Last week you backed up files from a diskette labeled docs. Display a list of those files.

**Command:** q b {docs}\\*

- Task** Query file system images from the **nas2** NAS file server.  
**Command:** `query backup -nasnodename=nas2 -class=nas`
- Task** Display a list of all files from your **c:** drive that are contained in the backupset name `weekly_accounting_data.3214`.  
**Command:** `query backup c:\* -subdir=yes  
-backupsetname=weekly_accounting_data.3214`
- Task** Display a list of all the files that are contained in the backupset name `weekly_accounting.3214`.  
**Command:** `query filespace -backupsetname=weekly_accounting.3214`

---

## 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 options file (dsm.opt). The value in the client options file is the default, but this value can be overridden on the command line. See “Nasnodename” on page 311 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 208 for more information.


## Query Backupset

The **query backupset** command queries a backup set from a local file or the Tivoli Storage Manager server. See “Location” on page 300 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 Windows clients. Tape support is only available on AIX, Solaris, and HP-UX.

## Syntax



## Parameters

*options*

Table 69. Query Backupset command: Related options

Option	Where to use	Page
<i>description</i>	Command line only.	230
<i>location</i>	Command line only.	300
<i>scrolllines</i>	Client options file (dsm.opt) or command line.	358
<i>scrollprompt</i>	Client options file (dsm.opt) or command line.	359

*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.

**backupsetname**

Specifies the name of the backup set.

**bset filename**

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

*LOcation*=

Specifies where Tivoli Storage Manager searches for a backup set during a query or restore operation. Use this option to locate backup sets on the server or local files. Backup sets that are generated on the server can be used locally specifying the *location* option and either the file name of the file containing the start of the backup set, or the tape device where the tape containing the start of the backup set is located.

- server** Specifies that Tivoli Storage Manager searches for the backup set on the server. This is the default.
- file** Specifies that Tivoli Storage Manager searches for the backup set on a local file.
- 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:** `query backupset backupsetname=c:\budget\weekly_budget_data.ost loc=file`

---

## Query Complusdb

The **query complusdb** command displays information about the active backup of the COM+ database system object on the Tivoli Storage Manager server.

### Supported Clients

This command is valid for Windows XP operating systems.

### Syntax

```
►►—Query COMPLusdb—┐—————►
                    └─ options ─┘
```

### Parameters

*options*

Table 70. Query Complusdb command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

### Examples

**Task** Display information about the active backup of the COM+ database system object on the server.

**Command:** query complusdb

---

## Query Eventlog

The **query eventlog** command displays information about the active backup of the event log system object on the Tivoli Storage Manager server.

### Supported Clients

This command is valid for the Windows XP Server operating system.

### Syntax

►►—Query Eventlog—┐  
                          └ options ┘◄◄

### Parameters

*options*

Table 71. Query Eventlog command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

### Examples

**Task** Display information about the active backup of the Windows XP event log system object on the server.

**Command:** query eventlog



## 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 drive at your workstation from which you back up or archive files.

A Unicode file space name might not display correctly if the server is unable to display the Unicode name. In this case, use the file space identifier (fsID) to identify these file spaces on the server. Use the **query filesystem** command with the *detail* option to determine the fsID of a file space. The fsID is also displayed in the file information dialog in the Windows and Web client GUIs.

### 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 options file (dsm.opt). The value in the client 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 options file, it must be specified on the command line when processing NAS file systems. See “Nasnodename” on page 311 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 208 for more information.

## Supported Clients

This command is valid for all Windows clients.

## Syntax

```
►► Query Filespace [file spacename] [options]
►► -BACKUPSetname=—backupsetname
```

## Parameters

### *file spacename*

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 72. Query Filespace command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	204
<i>class</i>	Command line only.	208
<i>dateformat</i>	Client options file (dsm.opt) or command line.	227
<i>detail</i>	Command line only.	232

Table 72. Query Filespace command: Related options (continued)

<i>fromnode</i>	Command line only.	268
<i>nasnodename</i>	Client options file (dsm.opt) or command line.	311
<i>scrolllines</i>	Client options file (dsm.opt) or command line.	358
<i>scrollprompt</i>	Client options file (dsm.opt) or command line.	359
<i>timeformat</i>	Client options file (dsm.opt) or command line.	397

## 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 \\florence\c\$ file space.  
**Command:** query filesystem \\florence\c\$
- Task** Display all of the file space names on the server with a file space name ending in '\$' belonging to system named *florence*.  
**Command:** query filesystem \\florence\\*\$
- Task** Query a file space from the **nas2** NAS file server.  
**Command:** query filesystem -nasnodename=nas2 -class=nas
- Task** Display your file spaces using the *backupsetname* option.  
**Command:** query filesystem -backupsetname=monthly\_accounting.2342

## Query Group

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 499 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:

```
tsm> q files
#      Last Incr Date      Type      File Space Name
-----
1      01/01/2007 00:00:00  TSMVFS    NEETAVFS
2      01/01/2007 00:00:00  TSMVFS    \NEETAVFS
3      01/01/2007 00:00:00  NTFS      \\rhyme\d$
tsm>
```

If you query a group backup without the *-inactive* option, the query displays only the latest group backup, whether it is type FULL or type DIFF:

```
tsm> q group {neetavfs}\neetag
Size      Backup Date      Mgmt Class  A/I Group
-----
345,737 B 06/05/2007 13:52:04  A123456789  A FULL NEETAVFS\NEETAG
```

## Supported Clients

This option is valid for all Windows clients.

## Syntax

►►—Query Group— *filespec* —————►  
                                  └ options ─┘

## Parameters

*options*

Table 73. Query Group command: Related options

Option	Where to use	Page
<i>fromnode</i>	Command line only.	268
<i>inactive</i>	Command line only.	277

Table 73. Query Group command: Related options (continued)

<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>showmembers</i>	Command line only.	364

*filespec*

Specifies the virtual file space name (enclosed in braces) and the group name on the server that you want to query.

## Examples

**Task** Display all the groups in the `virtfs` file space.

**Command:**

```
query group {virtfs}\*
```

**Task** Display active and inactive versions of the `virtfs\group1` filespec.

**Command:**

```
query group {virtfs}\group1 -inactive
```

**Task** Display the `virtfs\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 group {virtfs}\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. Because image backup allows you to back up only used blocks in a file system, the stored image size on the Tivoli Storage Manager server could be smaller than the volume size. For online image backups, the stored image can be larger than the file system based on the size of the cache files.
- 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 all Windows clients.

## Syntax

```
→ Query Image [ options ] [ logicalvolumename | filespace ] →
```

## Parameters

*options*

Table 74. Query Image command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	204
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>fromnode</i>	Command line only.	268
<i>inactive</i>	Command line only.	277
<i>location</i>	Command line only.	300
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>scrolllines</i>	Client options file (dsm.opt) or command line.	358
<i>scrollprompt</i>	Client options file (dsm.opt) or command line.	359
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

*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 active and inactive version of the h: image.

**Command:** q im h: -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 want 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 Windows clients.

## Syntax

►► Query INCLexcl- test pattern ◄◄

## Parameters

### *test pattern*

Use for testing the validity of patterns you want to use in your include-exclude list. When you use a test pattern with this command, the following occurs:

- The internal include-exclude list is *not* displayed
- The pattern is processed as if it had come from an include-exclude statement, including all the usual error checking
- The pattern is displayed as it would appear in the include-exclude list

If the test pattern has no errors, the compiled pattern result is the same as the test pattern.

## Examples

**Task** Display a list of include-exclude statements.

**Command:** query inclexcl

**Task** Test the validity of this pattern: ..\?x?\\*.log

**Command:** query inclexcl ..\?x?\\*.log

**Note:** Some system files are excluded explicitly by Tivoli Storage Manager. You can use the **query inclexcl** command to display a list of these files.

---

## Query 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 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 Windows clients.

## Syntax

```
▶▶—Query Mgmtclass —————▶▶  
    |  
    | options |
```

## Parameters

*options*

Table 75. Query Mgmtclass command: Related options

Option	Where to use	Page
<i>detail</i>	Command line only.	232
<i>frommode</i>	Command line only.	268

## Examples

**Task** Display default and available management classes.

**Command:** query mgmtclass



---

## Query Node

The **query node** command displays all the nodes for which an administrative user ID has authority to perform operations. The authorized administrative user ID should have at least client owner authority over the client workstation node they are using either from the command line or from the web.

When using an interactive command-line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID.

Use the *type* option to specify the type of node to filter for. Valid values are *nas*, *client*, *server*, and *any*. The default is *any*. See “Type” on page 404 for more information.

## Supported Clients

This command is valid for all Windows clients.

## Syntax

►► Query Node options ◀◀

## Parameters

*options*

Table 76. Query Node command: Related options

Option	Where to use	Page
<i>type</i>	Command line only.	404
<i>scrolllines</i>	Client options file (dsm.opt) or command line.	358
<i>scrollprompt</i>	Client options file (dsm.opt) or command line.	359

## 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 Windows 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 77. Query Options command: Related options

Option	Where to use	Page
<i>scrolllines</i>	Client options file (dsm.opt) or command line.	358
<i>scrollprompt</i>	Client options file (dsm.opt) or command line.	359

### 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 Registry

The **query registry** command displays information about the active backup of a Windows XP registry system object on the Tivoli Storage Manager server.

### Supported Clients

This command is valid for Windows XP clients.

### Syntax

```
►►—Query REgistry— [ options ]—————►►
```

### Parameters

*options*

Table 78. Query Registry command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

### Examples

**Task** Display information about the active backup of the Windows XP registry system object on the server.

**Command:** query re

---

## 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 Windows clients.

## Syntax

►►—Query Restore—◄◄

## Parameters

There are no parameters for this command.

## Examples

**Task** The following example displays the output when you use **query restore**:

```
--- Restartable Restore Information ---
Restartable Session: 1
  Start date/time: 10/17/2001 15:18:22
    Source: {"\\ers\c$"}\data\proposals\*
    Destination: - not specified by user -

Restartable Session: 2
  Start date/time: 10/17/2001 15:20:01
    Source: {"\\ers\c$"}\data\spreadsheets\*
    Destination: - not specified by user -
```

---

## 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 Windows 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 Windows clients.

## Syntax

►►—Query SEssion—◄◄

## Parameters

There are no parameters for this command.

## Examples

**Task** Display your session information.

**Command:** query session

---

## Query Sysfiles

The **query sysfiles** command displays information about the active backup of the Windows XP system and boot files on the Tivoli Storage Manager server.

### Supported Clients

This command is valid for Windows XP clients.

### Syntax

►►—Query SYSFiles—options—►►

### Parameters

*options*

Table 79. Query Sysfiles command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

### Examples

**Task** Display information about the active backup of the system and boot files system object on the server.

**Command:** query sysfiles

---

## 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.
- ENV - Environment variables.
- ERRORLOG - The Tivoli Storage Manager error log file.
- FILE - Attributes for the file name that you specify.
- FILESNOTTOBACKUP - Enumeration of Windows Registry key:

```
HKEY_LOCAL_MACHINE\  
  SYSTEM\  
    CurrentControlSet\  
      BackupRestore\  
        FilesNotToBackup
```

This key specifies those files that backup products should not back up. The **query inclexcl** command will indicate that these files are excluded per the operating system.

- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- KEYSNOTTORESTORE - Enumeration of Windows Registry key:

```
HKEY_LOCAL_MACHINE\  
  SYSTEM\  
    ControlSet001\  
      BackupRestore\  
        KeysNotToRestore
```

This key specifies those Windows Registry keys that backup products should not restore.

- MSINFO - Windows system information (output from MSINFO32.EXE).
- OPTIONS - Compiled options.
- OSINFO - Name and version of the client operating system
- POLICY - Policy set dump.
- REGISTRY - Windows Tivoli Storage Manager-related Windows Registry entries.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually dsmsched.log).
- SFP - The list of files protected by Windows System File Protection, and for each file, indicates whether that file exists. These files are backed up as part of the SYSFILES system object.
- SFP=<filename> - Indicates whether the specified file (*filename*) is protected by Windows System File Protection. For example:  

```
SFP=C:\WINNT\SYSTEM32\MSVCRT.DLL
```
- SYSTEMOBJECT - Windows system object information.
- CLUSTER - Windows 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 264 for more information.




2. Use the *console* option if you want to output the information to the console. See “Console” on page 225 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 Windows 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 80. Query Systeminfo command: Related options

Option	Where to use	Page
<i>console</i>	Command line only.	225
<i>filename</i>	Command line only.	264

## 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 dsmsoptfile errorlog  
-filename=tsminfo.txt

## Query Systemobject

The **query systemobject** command displays information about the active backup of the following valid Windows XP system objects on the Tivoli Storage Manager server, or inside a backupset when a backupset is specified.

- Active Directory (domain controller only)
- Certificate Server Database
- Cluster Database (cluster node only)
- COM+ database
- Windows Registry
- System and boot files
- System Volume
- Event logs (system, security and application)
- Removable Storage Management Database (RSM)
- Replicated file systems (FRS)
- Windows Management Instrumentation (WMI) repository

## Supported Clients

This command is valid for Windows XP clients only.

## Syntax

```
►► Query SYSTEMObject [ options ]
```

## Parameters

*options*

Table 81. Query Systemobject command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	204
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

## Examples

**Task** Display information about the active backup of all the system objects on the server.

**Command:** query systemo

**Task** Display information about the active backup of all the system objects that are contained within the backupset name `daily_backup_data.1234`.

**Command:** query systemobject -backupsetname=daily\_backup\_data.1234

---

## Query Systemstate

Use the **query systemstate** command to display information about a backup of the system state on the Tivoli Storage Manager server, or inside a backupset when a backupset is specified. The Tivoli Storage Manager Windows client does not support the current commands for querying individual system state components.

### Supported Clients

This command is valid for the Windows Server 2003 and Windows Vista clients only.

### Syntax

►►—Query SYSTEMState —————►►  
                                  └ options ─┘

### Parameters

*options*

Table 82. Query Systemstate command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	204
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>showmembers</i>	Command line only.	364
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

### Examples

**Task** Display information about the active backup of the system state on the Tivoli Storage Manager server.

**Command:** query systemstate

**Task** Display information about the active backup of the system state contained within the backupset name `daily_backup_data.1234`.

**Command:** query systemstate -backupsetname=daily\_backup\_data.1234

## Query WAS

Use the **query was** command to display backups of the WebSphere Application Server Network Deployment Manager (contains setup, application files, and configuration information) or the Application Server that match the node name and type of the Websphere Application Server 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 Websphere Application Server 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 system. See “Wastype” on page 420 for more information.

## Supported Clients

This command is valid for Windows Server 2003 clients.

## Syntax

```
►► Query WAS [ options ] -nodename
```

## Parameters

*options*

Table 83. Query WAS command: Related options

Option	Where to use	Page
<i>fromnode</i>	Command line only.	268
<i>inactive</i>	Command line only.	277
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>showmembers</i>	Command line only.	364
<i>wastype</i>	Command line only.	420

*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 Websphere Application Server backups.

## Examples

**Task** Query all Websphere Application Server 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

## Query VM

The **query VM** command can be used to verify the successful backups of the virtual machines from the backup proxy system.

## Supported Clients

This command is valid for Windows clients.

## Syntax

```
►► Query VM [ options ]
```

## Parameters

*options*

Table 84. Query VM command: Related options

Option	Where to use	Page
<i>VMLIST</i>	Command line or dsm.opt	415
<i>VMCHost</i>	Command line or dsm.opt	412
<i>VMCUSER</i>	Command line or dsm.opt	414
<i>VMCPW</i>	Command line or dsm.opt	412

## Examples

### File Level

**Command:** dsmc query vm -vmlist=vm1,vm2 -vmchost=virtserv  
-vmcuser=admin -vmcpw=xxxx

**Expected output:**

```
#      machinename      Last Incr Date      Type      File Space Name
---      -
1      vm1                01/31/2007 10:23:30  NTFS     \\VM1\C$
2      vm1                01/31/2007 10:23:30  NTFS     \\VM1\D$
3      vm2                01/31/2007 10:23:30  NTFS     \\VM2\C$
4      vm2                01/31/2007 10:23:30  NTFS     \\VM2\D$
5      vm2                01/31/2007 10:23:30  NTFS     \\VM2\E$
-----
```

---

## Query WMI

The **query wmi** command displays information about the active backup of the Windows Management Instrumentation (WMI) repository on the Tivoli Storage Manager server.

### Supported Clients

This command is valid for Windows XP only.

### Syntax

►►—Query WMI—┐  
                 └ options ┘

### Parameters

*options*

Table 85. Query WMI command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

### Examples

**Task** Query the Windows Management Instrumentation repository.

**Command:** query wmi

---

## Restart Restore

The **restart restore** command displays a list of your restartable restore sessions in the server database. You can select zero or more restartable restore sessions that you want to restart.

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 Windows 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 333.

### Notes:

1. When restoring a directory, its modification date and time is set to the date and time of the restore, not to the date and time the directory had when it was backed up. This is because Tivoli Storage Manager restores the directories first, then adds the files to the directories.
2. An error will occur if you attempt to restore a file whose name is the same as an existing file’s short name. For example, if you attempt to restore a file you specifically named ABCDEF~1.DOC into the same directory where a file named abcdefghijk.doc exists, the restore will fail because the Windows operating system equates the file named abcdefghijk.doc to a short name of ABCDEF~1.DOC. The restore function treats this as a duplicate file.

If this error should occur, perform any of the following actions to correct it:

- Restore the file with the short file name to a different location.
- Stop the restore and change the name of the existing file.
- Disable the short file name support on Windows.
- Do not use file names that would conflict with the short file naming convention; for example, do not use ABCDEF~1.DOC.

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.

See the Microsoft Knowledge Base article Q121007, *How to Disable the 8.3 Name Creation on NTFS Partitions*, for more information.

If the **restore** command is retried because of a communication failure or session loss, the transfer statistics display the bytes Tivoli Storage Manager attempted to transfer across *all* command attempts. Therefore, the statistics for bytes transferred might not match file statistics, such as those for file size.

### Virtual volume and virtual mount point (Windows NTFS)

When restoring all file system data, including virtual volumes and mount points, restore the virtual mount point before restoring any data inside the virtual volume. Otherwise, all data in the virtual volume is restored to the parent directory of the virtual volume.

**Attention:** When the client is instructed to restore the files on the mount volume, it first determines whether a mount point exists for the specified mount volume. If a mount volume exists in the specified location, the following error message appears:

```
ANS5179E Restoring a volume mount point to a non-empty directory
```

and Tivoli Storage Manager will not overwrite the mount volume to avoid potential data loss. To restore the mount volume, first delete the mount point associated with it.



You can restore a virtual mount point to a network drive; however, the restored mount point indicates the virtual volume in the network workstation, not the local client virtual volume. For best results, only restore a virtual mount point to a network drive when the network workstation already has the identical virtual mount point defined. To restore the mount volume, first delete the mount point associated with it.

To restore a mount point using the GUI or Web client, you must restore the root of the drive where the mount point is defined. If you select the mount point itself, Tivoli Storage Manager restores all mounted data, but not the mount junction.

If you use the command-line client to restore mounted data, back up the mount point in addition to backing up the mounted data. If the mount point is not backed up, you cannot use the command-line client to restore any data inside the mounted volume. For best results, back up at least one mount point for any virtual volume you want to restore.

### Restoring Microsoft Dfs junctions

To restore Microsoft Dfs junctions, you must restore Microsoft Dfs root. If you select the junction point itself, Tivoli Storage Manager restores data under junction, but not the junction itself. If you select a junction point that no longer exists under Dfs root, Tivoli Storage Manager creates a local directory under Dfs root with the same name as the junction before restoring data.

### Restoring active files

When restoring active and inactive versions of the same file using the *replace* option, only the most recently restored file is replaced.

### Performing restores

The client stores files on the Tivoli Storage Manager server using the Windows UNC (Universal Naming Convention), not the drive letter. The UNC name is the network name for the file. The system name is a part of the UNC name. For example, if your system name is STAR and you have a file named c:\doc\h2.doc, the UNC name is \\star\c\$\doc\h2.doc.

When you restore files on the same system from which they were backed up, you can use the local drive letter or the UNC name to refer to the file. For example, either of the following will restore c:\doc\h2.doc to its original location:

```
dsmc restore c:\doc\h2.doc
dsmc restore \\star\c$\doc\h2.doc
```

When you restore files on a system with a different name, then you must use the UNC name to refer to the file. This is true even if you are restoring to the same physical system, but the system name has changed since the backup occurred.

For example, if you back up c:\doc\h2.doc on system STAR and you want to restore it to system METEOR then you must use the UNC name to refer to the file. You must also specify a destination restore location. This is because the default behavior is to restore the file to its original location, which would be on system STAR. To restore the file to system METEOR, you can run either of the following on METEOR:

```
dsmc restore \\star\c$\doc\h2.doc c:\
dsmc restore \\star\c$\doc\h2.doc \\meteor\c$\
```

## Restoring from file spaces that are not Unicode-enabled

If you want to restore from file spaces that are not Unicode-enabled, you must specify the source on the server and a destination on the client. For example, you backed up your H disk named `\\your-node\h$` prior to installing the Unicode-enabled client. After the installation, you issue the following command for a selective backup:

```
sel h:\logs\*.log
```

Before the backup takes place, the server renames the file space to `\\your-node\h$_OLD`. The backup continues placing the data specified in the current operation into the Unicode-enabled file space named `\\your-node\h$`. That file space now contains only the `\logs` directory and the `*.log` files. If you want to restore a file from the (old) *renamed* file space to its original location, you must enter both the source and destination as follows:

```
restore \\your-node\h$_OLD\docs\myresume.doc h:\docs\
```

## Considerations

Tivoli Storage Manager restores named streams on a file basis only. Windows XP directories can contain named streams. Named streams attached to a directory will always be overwritten (regardless of the value of the prompt option) during a restore operation.

When restoring sparse files to a non-NTFS file system, set the Tivoli Storage Manager server communication time out value (`idletimeout`) to the maximum value of 255 to avoid client session timeout. Tivoli Storage Manager is restricted to restoring sparse files that are less than 4 gigabytes in size.

The following issues apply if more data is restored than the Microsoft disk quota allows:

- If the user performing the restore has a disk quota (e.g., belongs to the Backup Operator Group), Tivoli Storage Manager will not restore any data that exceeds the restore user's disk quota and will display a "Disk Full" message.
- If the user performing the restore does not have a disk quota (e.g., belongs to the Administrator Group), Tivoli Storage Manager will restore all data and transfer ownership of those files which exceed the original owner's disk quota to the user performing the restore (in this case, the Administrator).

If you are using subfile backup, each subfile component received from the server is counted as an object restored in the statistical summary.

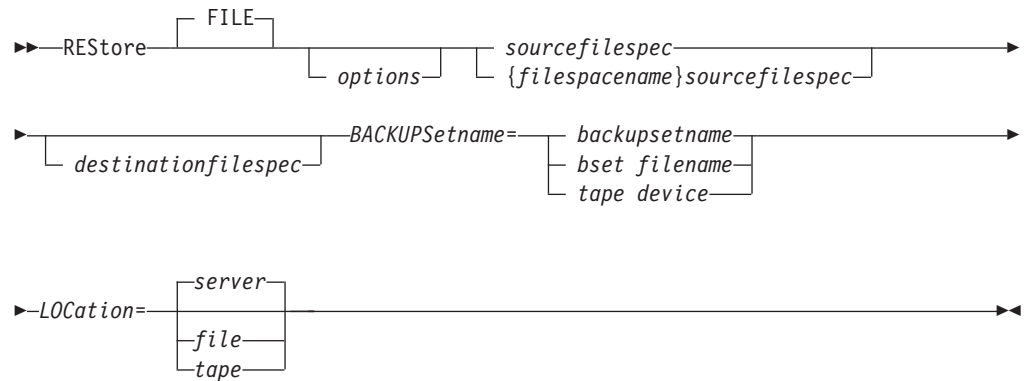
If you are using VMware Consolidated Backup, the restore will be done inside the virtual machine using the normal command syntax or client GUI interfaces, with no additional setup or command or options needed.

Tivoli Storage Manager backs up EFS encrypted files in raw encrypted format. The files are not decrypted prior to transmission; they are always in encrypted format. The files are restored in encrypted format as well.

## Supported Clients

This command is valid for all Windows 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 86. Restore command: Related options

Option	Where to use	Page
<i>asrmode</i>	Command line only.	196
<i>backupsetname</i>	Command line only.	204
<i>dateformat</i>	Client options file (dsm.opt) or command line.	227
<i>dirsonly</i>	Command line only.	235
<i>filelist</i>	Command line only.	262
<i>filesonly</i>	Command line only.	266
<i>fromdate</i>	Command line only.	267
<i>fromnode</i>	Command line only.	268
<i>frontime</i>	Command line only.	269
<i>ifnewer</i>	Command line only.	274
<i>inactive</i>	Command line only.	277
<i>latest</i>	Command line only.	298
<i>location</i>	Command line only.	300
<i>numberformat</i>	Client options file (dsm.opt) or command line.	316
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>preservepath</i>	Command line only.	333
<i>replace</i>	Client options file (dsm.opt) or command line.	340
<i>skipntpermissions</i>	Client options file (dsm.opt) or command line.	365

Table 86. Restore command: Related options (continued)

<i>skipntsecuritycrc</i>	Client options file (dsm.opt) or command line.	366
<i>subdir</i>	Client options file (dsm.opt) or command line.	381
<i>tapeprompt</i>	Client options file (dsm.opt) or command line.	386
<i>timeformat</i>	Client options file (dsm.opt) or command line.	397
<i>todate</i>	Command line only.	401
<i>totime</i>	Command line only.	402

*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.

**Note:** If you include *filespace*, do not include a drive letter in the file specification.

{*filespace*}

Specifies the file space (enclosed in braces) on the server that contains the files you want to restore. This is the name on the workstation drive from which the files were backed up.

Use the *filespace* if the drive label name has changed or if you are restoring files backed up from another node with drive labels that are different from yours. This is the drive label name or UNC name on the workstation drive from which the file was backed up. The following example is valid for specifying a UNC name: {'\\machine\C\$'}.

**Note:** You must specify a mixed or lowercase NTFS file space name enclosed in quotes and braces. For example, {"NTFSDrive"}. Single quotes or quotation marks are valid in loop mode. For example: {'NTFSDrive'} and {‘NTFSDrive’} are both valid. In batch mode, only single quotes are valid. The single quotes requirement is a restriction of the operating system.

*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.

When entering the *destinationfilespec*, consider the following:

- If the *sourcefilespec* names a single file, the *destinationfilespec* can be a file or a directory. If you are restoring a single file, you can optionally end the specification with a file name if you want to give the restored file a new name.
- If the *sourcefilespec* is wildcarded or *subdir=yes* is specified, the *destinationfilespec* must be a directory and end with a directory delimiter (\).

**Note:** If the destination path or any part of it does not exist, Tivoli Storage Manager will create it.

## Examples

**Task** Restore a single file named budget.fin.

**Command:** restore c:\devel\projecta\budget.fin

**Task** Restore a single file named budget.fin which resides in the current directory.

- Command:** restore file budget.fin
- Task** Restore files from the abc NTFS file space proj directory.  
**Command:** rest {"abc"}\proj\\*.\*
- Task** Restore all files with a file extension of .c from the c:\devel\projecta directory.  
**Command:** rest c:\devel\projecta\\*.c
- Task** Restore all files with an extension of .c from the \devel\projecta directory located in the winnt file space.  
**Command:** rest {winnt}\devel\projecta\\*.c
- Task** Restore all files with a file extension of .c from the c:\devel\projecta directory to the c:\newdevel\projectn\projecta directory. If the projectn or projectn\projecta directory does not exist, it is created.  
**Command:** restore c:\devel\projecta\\*.c c:\newdevel\projectn
- Task** Restore files in the c:\project directory. Use the *pick* and *inactive* options to select active and inactive backup versions.  
**Command:** restore c:\project\\* -pi -ina
- Task** Restore all files in the c:\mydir directory to their state as of 1:00 PM on August 17, 2002.  
**Command:** restore -pitd=8/17/2002 -pitt=13:00:00 c:\mydir\
- Task** Restore a file from the renamed file space \\your-node\h\$\_OLD to its original location. Enter both the source and destination as follows:  
**Command:** res \\your-node\h\$\_OLD\docs\myresume.doc h:\docs\
- Task** Restore files specified in the filelist to a different location.  
**Command:** res -filelist=c:\avi\restorelist.txt  
c:\NewRestoreLocation\
- Task** Restore a single file named budget.fin contained within the backupset named daily\_backup\_data.1234.  
**Command:** restore c:\projecta\budget.fin  
-backupsetname=daily\_backup\_data.1234 -location=server

---

## Restore ASR

Use the **restore asr** command to restore the Automated System Recovery (ASR) files to a specified location. By default, the client restores these files to the `adsm.sys\ASR` staging directory.

You can create an ASR recovery diskette by specifying the diskette drive as the restore location or by copying the files from the `adsm.sys\ASR` staging directory to a diskette. See “ASR preparation procedure” on page 116 for more information.

## Supported Clients

This command is valid for the Windows XP and Windows Server 2003 clients only.

## Syntax

```
►►—REStore ASR [ destinationfilespec ] [ options ]
```

## Parameters

*options*

Table 87. Restore ASR command: Related options

Option	Where to use	Page
<i>inactive</i>	Command line only.	277
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324

*destinationfilespec*

Specifies the path where you want to place the restored files. If you do not specify a destination, the client restores the files to the `adsm.sys\ASR` directory on the system drive root directory.

## Examples

**Task** Restore the ASR files from the Tivoli Storage Manager server to the `adsm.sys\ASR` directory on the system drive root directory. Use the *pick* and *inactive* options to select from a list of both active and inactive backup versions.

**Command:** `restore asr -pick -inactive`

---

## 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 300 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 {\\machinename\c$}\*
c:\destdir -backupsetname=<backupsetname>
-location=<server> -subdir=yes
```

The following is an example:

```
restore backupset {\\financeDept\d$}\*
d:\financeDeptRestDir -subdir=yes
-backupsetname=financeData1.1234 -location=server
```

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 enclosed in braces {}, followed with a terminating directory delimiter. Here are two examples:

```
restore backupset {virtfs}\* c:\newdevel\projectn\
-backupsetname=mybackupset -loc=server -subdir=yes
restore backupset {myGrpFs}\* c:\myGroupRestDir
-backupsetname=groupData.1234 -location=server -subdir=yes
```

The entire group, or groups in the virtual file space will be restored. You cannot restore a partial group by specifying a qualified source file space.

Considerations:

- 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 with 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 using the LAN.

## Supported Clients

This command is valid for all Windows clients.

## Syntax



## Parameters

*options*

Table 88. Restore Backupset command: Related options

Option	Where to use	Page
<i>dirsonly</i>	Command line only.	235
<i>filesonly</i>	Command line only.	266
<i>ifnewer</i>	Command line only.	274
<i>location</i>	Command line only.	300
<i>preservepath</i>	Command line only.	333
<i>quiet</i>	Client options file (dsm.opt) or command line.	339
<i>replace</i>	Client options file (dsm.opt) or command line.	340
<i>skipntppermissions</i>	Client options file (dsm.opt) or command line.	365
<i>subdir</i>	Client options file (dsm.opt) or command line.	381

*sourcefilespec*

Specifies that you want to restore all Windows XP system objects from the backup set. You cannot restore individual system objects from the backup set. This parameter is for the Windows XP clients only.

*destinationfilespec*

Specifies the destination path for the restored files. If you do not specify a *sourcefilespec*, you cannot specify a *destinationfilespec*. If you do not specify a destination, Tivoli Storage Manager restores the files to the original source path. Windows XP system objects can only be restored to the original source path. If you are restoring more than one file, you must end the file specification with a directory delimiter (\), otherwise, Tivoli Storage Manager assumes the last name is a file name and reports an error. If you are restoring a



single file, you can optionally end the destination file specification with a file name if you want to give the restored file a new name.

**BACKUPSetname=**

Specifies the name of the backup set on the server from which to perform a restore operation. You cannot use wildcard characters to specify the backup set name. Tivoli Storage Manager restores the backup set from the server using the LAN.

**backupsetname**

Specifies the name of the backup set on the server from which to perform a restore operation. The *location* option should be set to *server*.

**bset filename**

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

**LOCation=**

Specifies where Tivoli Storage Manager searches for a backup set during a query or restore operation. Use this option to locate backup sets on the server or local files. Backup sets that are generated on the server can be used locally specifying the *location* option and either the file name of the file containing the start of the backup set, or the tape device where the tape containing the start of the backup set is located.

**server** Specifies that Tivoli Storage Manager restores the backup set on the server. This is the default.

**file** Specifies that Tivoli Storage Manager restores the backup set on a local file.

**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:** `restore backupset backupsetname=c:\budget\  
weekly_budget_data.ost -loc=file`

**Task** Restore a backup set from the `\\.\tape0` device.

**Command:** `dsmc restore backupset \\.\tape0 -loc=tape`

**Task** Restore a group from the backup set `mybackupset` on the Tivoli Storage Manager server to the `c:\newdevel\projectn` directory.

**Command:** `dsmc restore backupset {grpFilespaceName}\*  
c:\newdevel\projectn\ -backupsetname=mybackupset -loc=server  
-subdir=yes`

In the previous example, `grpFilespaceName` is the virtual file space name under which the group objects were backed up.

---

## Restore Complusdb

The **restore complusdb** command restores a Windows XP COM+ database from the Tivoli Storage Manager server.

**Note:** You can use the **restore systemobject** command to restore all system objects at the same time.

For information on system objects that must be restored together, see “Backing up system objects (Windows XP)” on page 79.

## Supported Clients

This command is valid for Windows XP operating systems.

## Syntax

►►—REStore COMPLusdb options ◀◀

## Parameters

*options*

Table 89. Restore Complusdb command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

## Examples

**Task** Restore the COM+ database.

**Command:** restore compl

---

## Restore Eventlog

The **restore eventlog** command restores the active Windows XP event logs stored from the Tivoli Storage Manager server.

Event logs are restored to the `adsm.sys\eventlog` directory under the root directory.

**Note:** You can use the **restore systemobject** command to restore all system objects at the same time.

## Supported Clients

This command is valid for Windows XP clients only.

## Syntax

```
►►—REStore Eventlog [ options ]
```

## Parameters

*options*

Table 90. Restore Eventlog command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

## Examples

**Task** Restore the entire event log.

**Command:** rest event

## Restore Group

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 529 for more information.

## Supported Clients

This command is valid for all Windows clients.

## Syntax

```
►► restore group [ options ] sourcefilespec [ destinationfilespec ] ◀◀
```

## Parameters

*options*

Table 91. Restore Group command: Related options

Option	Where to use	Page
<i>backupsetname</i>	Command line only.	204
<i>fromdate</i>	Command line only.	267
<i>fromnode</i>	Command line only.	268
<i>fromtime</i>	Command line only.	269
<i>ifnewer</i>	Command line only.	274
<i>inactive</i>	Command line only.	277
<i>latest</i>	Command line only.	298
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>replace</i>	Client options file (dsm.opt) or command line.	340
<i>showmembers</i>	Command line only.	364
<i>skipntpermissions</i>	Client options file (dsm.opt) or command line.	365
<i>skipntsecuritycrc</i>	Client options file (dsm.opt) or command line.	366
<i>tapeprompt</i>	Client options file (dsm.opt) or command line.	386
<i>todate</i>	Command line only.	401
<i>totime</i>	Command line only.	402

*sourcefilespec*

Specifies the virtual file space name (enclosed in braces) 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 system.

**Command:**

```
restore group {virtfs}\group1
```

**Task** Display all groups within the virtfs virtual file space. Use the *showmembers* option to display a list of group members from which you can select one or more to restore.

**Command:**

```
restore group {virtfs}\* -pick -showmembers
```

**Task** Display a list of groups within the virtfs virtual file space from which you can select one or more groups to restore.

**Command:**

```
restore group {virtfs}\* -pick
```

---

## 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.

### Notes:

1. The account running the Tivoli Storage Manager client *must* have administrator authority to successfully perform any type of image restore.
2. If you use HSM for Windows or HSM for UNIX, and you restore a file system image backup and plan to run reconciliation, you must restore the files that have been backed up after the image backup. Otherwise, migrated files that have been created after the image backup expire from the HSM archive storage on the Tivoli Storage Manager server.

You can use the *verifyimage* option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log. See “Verifyimage” on page 408 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 276 for more information.

### Considerations:

- The Tivoli Storage Manager API must be installed to use the **restore image** command.
- You can restore an NTFS file system to a FAT32 volume or vice-versa.
- The destination volume to which you restore must exist and be the same size or larger than the source volume.
- The physical layout of the target volume (striped, mirrored, etc.) can differ.
- The target volume will be overwritten with data contained in the image backup.
- It is not necessary to format a target volume prior to restoring an image backup containing a file system.
- Tivoli Storage Manager requires an exclusive lock to destination volume you are restoring. The client will lock, restore, unlock, unmount, and mount the volume during a restore operation. During the restore process the destination volume will not be available to other applications.
- 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. Because image backup allows you to back up only used blocks in a file system, the stored image size on the Tivoli Storage Manager server could be smaller than the volume size. For online image backups, the stored image can be larger than the file system based on the size of the cache files.
  - 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 should run **chkdsk** to check for and repair any bad sectors or data inconsistencies (unless the restored volume is RAW).

## Supported Clients

This command is valid for Windows 32-bit platforms.





## Examples

**Task** Restore the e: drive to its original location.

**Command:** `dsmc rest image e:`

**Task** Restore the h: drive 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 h: -incremental -deletefiles`

**Task** Restore the d: drive to its original location. Use the *verifyimage* option to enable detection of bad sectors on the target volume.

**Command:** `dsmc restore image d: -verifyimage`

**Task** If bad sectors present on the target volume, use the *imagetofile* option to restore the d: drive to the e:\diskD.img file to avoid data corruption.

**Command:** `dsmc restore image d: e:\diskD.img -imagetofile`

**Task** Restore the e: drive from the backupset named `weekly_backup_data.1234` to its original location.

**Command:** `restore image e: -backupsetname=weekly_backup_data.1234 -location=server`

---

## 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 options file (*dsm.opt*). The value in the client options file is the default, but this value can be overridden on the command line. See “Nasnodename” on page 311 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 309.

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 460.

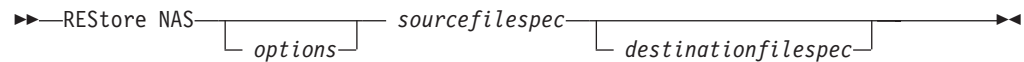
A NAS file system specification uses the following conventions:

- 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 all Windows clients.

## Syntax



## Parameters

*options*

Table 93. Restore NAS command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>mode</i>	Command line only.	307
<i>monitor</i>	Command line only.	309
<i>nasnodename</i>	Client options file (dsm.opt) or command line.	311
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

*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 Registry

Use the **restore registry** command restores the Windows XP registry.

**Note:** You can use the **restore systemobject** command to restore all system objects at the same time.

For information on system objects that must be restored together, see “Backing up system objects (Windows XP)” on page 79.

When you restore a Windows Registry key, you have a choice of whether or not to activate the key. A backup copy of the active Windows Registry and registry key is made in the `adsm.sys\registry.sav` file on the system drive root directory before the Windows Registry is restored. If Windows does not boot with the newly restored Windows Registry, you can use this saved copy to reboot.

If you cannot boot Windows, you might be able to manually restore the Windows Registry if you can boot another operating system that does not use the Windows Registry. During backup and restore, copies of the Windows Registry are saved on the workstation in the following directories (assuming that your system drive is c:):

### **c:\adsm.sys\registry**

This directory contains subdirectories that are named after the Windows workstation name. For example, if the workstation name is **tanenhaus**, the directory would be: `c:\adsm.sys\registry\tanenhaus`

**Attention:** *Restore the registry to the same workstation from which it was backed up, since the registry contains specific hardware and software information. If you try to restore to another workstation, you might not be able to boot the computer after the restore.*

For Windows XP registry restore you can simply enter:

```
RESTORE REGISTRY -ACTIVATE=NO
```

Once you have entered that command, you can then go into the `c:\adsm.sys\W2KReg` directory to find the individual registry keys you want to restore. *For the restore to take effect, you must reboot your system.*

## Supported Clients

This command is valid for Windows XP clients.

## Syntax

```
►►—REStore Registry— ACTIVATEkey  Yes  No  options
```

## Parameters

*options*

Table 94. Restore Registry command: Related options

Option	Where to use	Page
<i>activatekey</i>	Client option file ( <code>dsm.opt</code> ) or command line.	192

Table 94. Restore Registry command: Related options (continued)

<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

## Examples

**Task** Restore the entire directory.

**Command:** restore registry

---

## Restore Sysfiles

The **restore sysfiles** command restores Windows XP system and boot files. System and boot files are part of the system state data and must be restored along with the other components of that data such as the Windows Registry, cluster service information and the certificate server database.

**Note:** You can use the **restore systemobject** command to restore all system objects at the same time.

## Supported Clients

This command is valid for Windows XP clients.

## Syntax

```
►►—REStore SYSFiles [ options ]
```

## Parameters

*options*

Table 95. Restore Sysfiles command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

## Examples

**Task** Restore system and boot files.

**Command:** res sysf

## Restore Systemobject

The **restore systemobject** command restores all of the following active or inactive Windows XP system objects stored on the Tivoli Storage Manager server, or inside a backupset when a backupset is specified:

- Active Directory (domain controller only)
- Certificate Server Database
- Cluster Database (cluster node only)
- COM+ database
- Windows Registry
- System and boot files
- System Volume
- Event logs (system, security and application)
- Removable Storage Management Database (RSM)
- Replicated file systems
- Windows Management Instrumentation (WMI) repository

The **restore systemobject** command is only valid for Windows XP.

## Supported Clients

This command is valid for Windows XP clients only.

## Syntax

```
►►—REStore SYSTEMObject [ options ]
```

## Parameters

*options*

Table 96. Restore Systemobject command: Related options

Option	Where to use	Page
<i>asrmode</i>	Command line only.	196
<i>backupsetname</i>	Command line only.	204
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

## Examples

**Task** Restore all active system objects.

**Command:** restore systemobject

**Task** Restore all active system objects contained within the backupset named `daily_backup_data.1234`.

**Command:** restore systemobject  
-backupsetname=daily\_backup\_data.1234



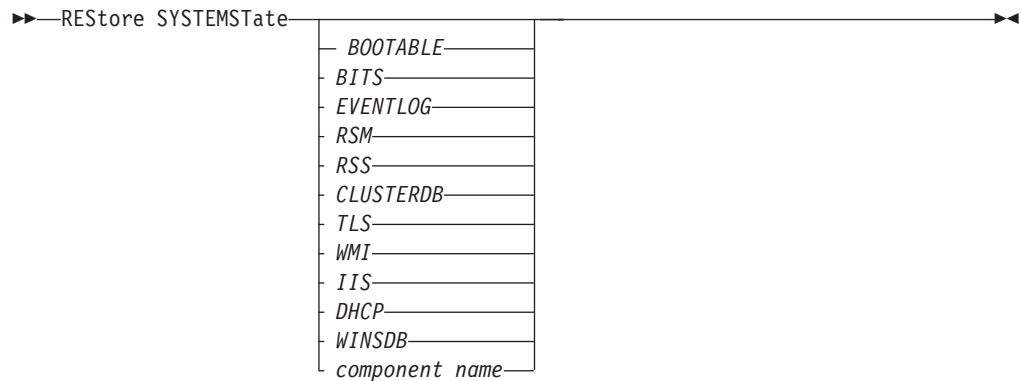
## Restore Systemstate

Use the **restore systemstate** command to restore the backup of the system state stored on the Tivoli Storage Manager server, or inside a backupset when a backupset is specified. The Tivoli Storage Manager Windows client does not support the current commands for restoring individual bootable system state components.

## Supported Clients

This command is valid for the Windows Server 2003 and Windows Vista clients only.

## Syntax



## Parameters

*options*

Table 97. Restore Systemstate command: Related options

Option	Where to use	Page
<i>asrmode</i>	Command line only.	196
<i>backupsetname</i>	Command line only.	204
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

## Examples

**Task** Restore the entire system state object.

**Command:** restore systemstate

**Task** Restore the bootable system state component.

**Command:** restore systemstate bootable

**Task** Restore the system service WMI component.

**Command:** restore systemstate wmi

|                   **Task**   Restore the system state object contained within the backupset named  
|                                   daily\_backup\_data.1234.  
|                   **Command:** restore systemstate -backupsetname=daily\_backup\_data.1234  
|                   **Note:** If a backupset is specified, it is from the server *only*.  
|

---

## Restore WAS

The **restore was** command specifies whether to restore the WebSphere Application Server 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.

**Note:** System administrator authority is required to restore the WebSphere Application Server Network Deployment Manager or the Application Server.

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 Websphere Application Server that you want to restore. The default is ND. See “Wastype” on page 420 for more information.

### WebSphere Application Server 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 Websphere Application Server on a new system:

1. Install Websphere Application Server on a system with the same name and IP address as the system on which the back up was performed.

If you backed up a Network Deployment Manager on system *ebarton*, the restore must also be performed on system *ebarton*. Install the same type of Websphere Application Server 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 Websphere Application Server to the same location in which the backup was performed. For example, if the backup was performed on an APP server installed at `c:\WebSphere\App`, the new server must also be installed at `c:\WebSphere\App`. After the installation is complete, ensure that the server is stopped.

#### Notes:

- a. You can use the *washome* option in your client options file (`dsm.opt`) to specify an override base installation path for the Application Server. See “Washome” on page 417 for more information.
  - b. You can use the *wasndhome* option in your client options file (`dsm.opt`) to specify an override base installation path for the Network Deployment manager. See “Wasndhome” on page 418 for more information.
2. Restore the Websphere Application Server 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 instance of a server. This procedure assumes that there is a preexisting Websphere Application Server installation that is running and has a backup stored on the Tivoli Storage Manager server:

1. Stop the Websphere Application Server server to be restored.
2. Restore the correct Websphere Application Server node (the 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.

For proper operation, the was node must be restored to same location and under same name.

**Attention:** Restoring data other than at the group level can corrupt your Websphere Application Server installation. Restore your data at the Network Deployment Manager node or Application Server node level only.

**Notes:**

- 1.
- 2.

## Supported Clients

This command is valid for Windows Server 2003 clients.

## Syntax

```

>> REStore WAS [ options ] --nodename <<<
  
```

## Parameters

*options*

Table 98. Restore WAS command: Related options

Option	Where to use	Page
<i>fromdate</i>	Command line only.	267
<i>fromnode</i>	Command line only.	268
<i>frontime</i>	Command line only.	269
<i>ifnewer</i>	Command line only.	274
<i>inactive</i>	Command line only.	277
<i>latest</i>	Command line only.	298
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>preservepath</i>	Command line only.	333
<i>replace</i>	Client options file (dsm.opt) or command line.	340
<i>showmembers</i>	Command line only.	364
<i>tapeprompt</i>	Client options file (dsm.opt) or command line.	386
<i>todate</i>	Command line only.	401
<i>totime</i>	Command line only.	402
<i>wastype</i>	Command line only.	420

*nodename*

Specifies the node name of the instance of Websphere Application Server 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

---

## Restore WMI

The **restore wmi** command restores a Windows XP Windows Management Instrumentation (WMI) repository on the Tivoli Storage Manager server. The WMI system object is restored to `adsm.sys\wmi\wmidbfile` file.

**Note:** You can use the **restore systemobject** command to restore all system objects at the same time.

## Supported Clients

This command is valid for Windows XP only.

## Syntax

►►—REStore WMI—┐  
                  └ options ┘

## Parameters

*options*

Table 99. Restore WMI command: Related options

Option	Where to use	Page
<i>dateformat</i>	Client option file (dsm.opt) or command line.	227
<i>inactive</i>	Command line only.	277
<i>numberformat</i>	Client option file (dsm.opt) or command line.	316
<i>pick</i>	Command line only.	322
<i>pitdate</i>	Command line only.	323
<i>pittime</i>	Command line only.	324
<i>timeformat</i>	Client option file (dsm.opt) or command line.	397

## Examples

**Task** Restore the Windows Management Instrumentation repository.

**Command:** `restore wmi`

---

## 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 190.

### Notes:

1. When retrieving a directory, its modification date and time is set to the date and time of the retrieve, not to the date and time the directory had when it was archived. This is because Tivoli Storage Manager retrieves the directories first, then adds the files to the directories.
2. An error will occur if you attempt to retrieve a file whose name is the same as the short name of an existing file. For example, if you attempt to retrieve a file you specifically named ABCDEF~1.DOC into the same directory where a file named abcdefghijk.doc exists, the retrieve will fail because the Windows operating system equates the file named abcdefghijk.doc to a short name of ABCDEF~1.DOC. The retrieve function treats this as a duplicate file.

If this error should occur, perform any of the following actions to correct it:

- Retrieve the file with the short file name you specified to a different location.
- Stop the retrieval, and change the name of the existing file.
- Disable the short file name support on Windows.
- Do not use file names that conflict with the short file naming convention. For example, do not use ABCDEF~1.DOC.

If you are using client Version 3.1.0.5 or later, the workstation name is part of the file name. Therefore, if you archive files on one workstation and you want to retrieve them to another workstation, you must specify a destination. This is true even if you are retrieving to the same physical workstation, but the workstation has a new name. For example, to retrieve the c:\doc\h2.doc file to its original directory on the workstation, named *star*, you would enter:

```
dsmc retrieve c:\doc\h2.doc \\star\c$\
```

To retrieve the file to *star*, which has now been renamed *meteor*, you would enter:

```
dsmc retrieve c:\doc\h2.doc \\meteor\c$\
```

You could also enter:

```
dsmc retrieve c:\doc\h2.doc \\star\c$\
```

You can enter the command in either of the preceding ways, because if the workstation name is not included in the specification, the local workstation is assumed (*meteor*, in this case).

## Retrieving from file spaces that are not Unicode-enabled

If you want to retrieve archives from file spaces that were renamed by the Unicode-enabled client, you must specify the source on the server and a destination on the client. For example, you archived files from your H-disk, named `\\your-node\h$` prior to installing the client. After the installation, you issue the following archive command:

```
arc h:\logs\*.log
```

Before the archive takes place, the server renames the file space to `\\your-node\h$_OLD`. The archive continues placing the data specified in the current operation into the Unicode-enabled file space named `\\your-node\h$`. That file space now contains only the `\logs` directory and the `*.log` files. If you want to retrieve a file from the (old) *renamed* file space to its original location, you must enter both the source and destination as follows:

```
retrieve \\your-node\h$_OLD\docs\myresume.doc h:\docs\
```

## Considerations

Tivoli Storage Manager retrieves named streams on a file basis only. Directories in Windows XP can contain named streams. Named streams attached to a directory will always be overwritten (regardless of the value of the *prompt* option) during the retrieve.

When retrieving sparse files to a non-NTFS file system, set the Server communication time out value (`idletimeout`) to the maximum value of 255 to avoid client session timeout.

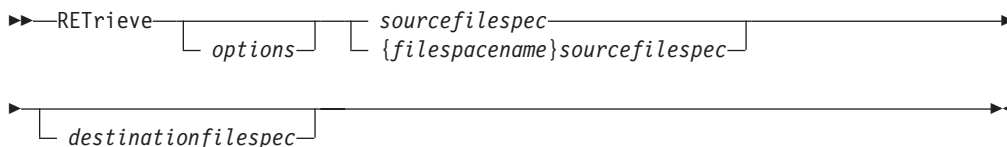
The following issues apply if more data is restored than the Microsoft disk quota allows:

- If the user performing the retrieve has a disk quota (e.g., belongs to the Backup Operator Group), Tivoli Storage Manager will not retrieve any data that exceeds the retrieve user's disk quota and will display a "Disk Full" message.
- If the user performing the retrieve does not have a disk quota (e.g., belongs to the Administrator Group), Tivoli Storage Manager will retrieve all data and transfer ownership of those files which exceed the original owner's disk quota to the user performing the retrieve (in this case, the Administrator).

## Supported Clients

This command is valid for all Windows clients.

## Syntax



## Parameters

*options*

Table 100. Retrieve command: Related options

Option	Where to use	Page
--------	--------------	------



Table 100. Retrieve command: Related options (continued)

<i>dateformat</i>	Client options file (dsm.opt) or command line.	227
<i>description</i>	Command line only.	230
<i>dironly</i>	Command line only.	235
<i>filelist</i>	Command line only.	262
<i>filesonly</i>	Command line only	266
<i>fromdate</i>	Command line only	267
<i>fromnode</i>	Command line only.	268
<i>frontime</i>	Command line only	269
<i>ifnewer</i>	Command line only	274
<i>pick</i>	Command line only.	322
<i>preservepath</i>	Command line only.	333
<i>replace</i>	Client options file (dsm.opt) or command line.	340
<i>skipntpermissions</i>	Client options file (dsm.opt) or command line	365
<i>skipntsecuritycrc</i>	Client options file (dsm.opt) or command line	366
<i>subdir</i>	Client options file (dsm.opt) or command line.	381
<i>tapeprompt</i>	Client options file (dsm.opt) or command line.	386
<i>timeformat</i>	Client options file (dsm.opt) or command line.	397
<i>todate</i>	Command line only.	401
<i>totime</i>	Command line only.	402

*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 63 for the maximum file size for retrieve processing.

**Note:** If you include *filespace*name, do not include a drive letter in the file specification.

{*filespace*name}

Specifies the file space (enclosed in braces) on the server that contains the files you want to retrieve. This name is the drive label on the workstation drive from which the files were archived.

Use the file space name if the drive label name has changed, or if you are retrieving files archived from another node with drive label names that are different from yours.

**Note:** You must specify a mixed or lowercase NTFS file space name enclosed in quotes and braces. For example, {"NTFSDrive"}. Single quotes or quotation marks are valid in loop mode. For example: {'NTFSDrive'} and {‘NTFSDrive’} are both valid. In batch mode, only single quotes are valid. The single quotes requirement is a restriction of the operating system.

*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.

When entering the *destinationfilespec* consider the following:

- If the *sourcefilespec* names a single file, the *destinationfilespec* can be a file or a directory.
- If the *sourcefilespec* is wildcarded or you specify the *subdir=yes* option, the *destinationfilespec* must be a directory and end with a directory delimiter (\).

**Note:** If the destination path or any part of it does not exist, Tivoli Storage Manager will create it.

## Examples

- Task** Retrieve a single file named budget.fin.  
**Command:** ret c:\devel\projecta\budget.fin
- Task** Retrieve all files with an extension of .c from the c:\devel\projecta directory.  
**Command:** ret c:\devel\projecta\\*.c
- Task** Retrieve all files with a file extension of .c from the \devel\projecta directory on the winnt file space.  
**Command:** ret {winnt}\devel\projecta\\*.c
- Task** Retrieve all files in the c:\devel directory.  
**Command:** ret c:\devel\\*
- Task** Retrieve files from the abc file space proj directory.  
**Command:** ret {abc}\proj\\*.\*
- Task** Retrieve all files with a file extension of .c from the c:\devel\projecta directory to the c:\newdevel\projectn\projecta directory. If the \projectn or the \projectn\projecta directory does not exist, it is created.  
**Command:** ret c:\devel\projecta\\*.c c:\newdevel\projectn\
- Task** Retrieve files in the c:\project directory. Use the *pick* option.  
**Command:** ret c:\project\\* -pick
- Task** Retrieve a file from the renamed file space \\your-node\h\$\_OLD to its original location. Enter both the source and destination as follows:  
**Command:** ret \\your-node\h\$\_OLD\docs\myresume.doc h:\docs\

---

## Schedule

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 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 options file (dsm.opt) or on the **schedule** command.

After you start the client scheduler, it continues to run and to start scheduled events until you press **Ctrl+Break**, restart the workstation, 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 Windows clients.

## Syntax

►►—Schedule—┐  
└ options ┘

## Parameters

*options*

Table 101. Schedule command: Related options

Option	Where to use	Page
<i>maxcmdretries</i>	Client options file (dsm.opt) or command line.	304
<i>password</i>	client options file (dsm.opt)	319
<i>queryschedperiod</i>	Client options file (dsm.opt) or command line.	338
<i>retryperiod</i>	Client options file (dsm.opt) or command line.	347
<i>schedlogname</i>	Client options file (dsm.opt) or command line.	352
<i>schedmode</i>	Client options file (dsm.opt) or command line.	355
<i>sessioninitiation</i>	Client options file (dsm.opt) or command line.	361
<i>tcpclientport</i>	Client options file (dsm.opt) or command line.	392

## Examples

**Task** Start the client scheduler.

**Command:** `dsmc sch -password=notell`

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 options file (`dsm.opt`). 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.

**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.

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 c:\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.

### Migrating to Unicode-enabled file spaces

See “Autofsrename” on page 202 for information about using the Unicode-enabled client.

### Open file support

If open file support has been configured (see “Configuring Open File Support (OFS)” on page 29), Tivoli Storage Manager performs a snapshot backup or archive of files that are locked (or “in use”) by other applications. See “Performing an image backup” on page 82, for information on how to install and configure the LVSA. To control an open file support operation with LVSA, you can specify these additional options in your *dsm.opt* file or as values of the *include.fs* option: *snapshotcachelocation*, *snapshotcachesize*, *snapshotfsidleretries*, *snapshotfsidlewait*, *snapshotproviderfs*, *presnapshotcmd*, *postsnapshotcmd*. See “Include options” on page 280 for more information.

#### Notes:

1. You can use the *include.fs* option to set snapshot options on a per file system basis.
2. Use the *snapshotcachelocation* option to relocate the cache if necessary. You can specify a *snapshotcachelocation* for a specific drive using the *include.fs* option.

3. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with FAT, FAT32 or NTFS file systems. This support includes SAN-attached volumes that meet these requirements.
4. If the client is unable to create a snapshot, failover to non-OFS backup occurs; the same backup support that would be done if the OFS feature was not configured.
5. To enable open file support in a cluster environment all systems in the cluster should have the OFS feature configured.

For information about Tivoli Storage Manager Open File Support restrictions and issues, search for the **TSM Client v5.2 Open File Support** document under the **Storage Management** product category at the following Web site:

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

### Adaptive subfile backups

If you perform a backup over a network device with limited bandwidth, such as with a modem, you can reduce network traffic by using the *subfilebackup* option. In order to use this feature, you must first specify the adaptive subfile backup options during a normal backup. For information about adaptive subfile backups, see “Performing a backup with limited bandwidth” on page 61. For more information about using the *subfilebackup* option, see “Subfilebackup” on page 383.

### 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 375 for more information.

### Considerations

Tivoli Storage Manager backs up named streams on a file basis only. Backup of a named stream containing sparse file data is not supported. A sparse file is backed up as a regular file if Client compression is off. Enable file compression (*compression=on*) when backing up sparse files to minimize network transaction time and to maximize server storage space.


Microsoft disk quotas do not affect the amount of data you can back up.

Tivoli Storage Manager backs up EFS encrypted files in raw encrypted format. The files are not decrypted prior to transmission; they are always in encrypted format. The files are restored in encrypted format as well.

## Supported Clients

This command is valid for all Windows clients.

## Syntax



## Parameters

*options*

Table 102. Selective command: Related options

Option	Where to use	Page
<i>changingretries</i>	Client options file (dsm.opt) or command line.	207
<i>compressalways</i>	Client options file (dsm.opt) or command line.	221
<i>compression</i>	Client options file (dsm.opt) or command line.	222
<i>dironly</i>	Command line only.	235
<i>filelist</i>	Command line only.	262
<i>filesonly</i>	Command line only.	266
<i>postsnapshotcmd</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	327
<i>preservelastaccessdate</i>	Client options file (dsm.opt) or command line.	331
<i>presnapshotcmd</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	336
<i>skipntpermissions</i>	Client options file (dsm.opt) or command line.	365
<i>skipntsecuritycrc</i>	Client options file (dsm.opt) or command line.	366
<i>snapshotcachelocation</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	367
<i>snapshotcachesize</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	368
<i>snapshotfsidleretries</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	369
<i>snapshotfsidlewait</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	371
<i>snapshotproviderfs</i>	Client options file (dsm.opt) or with <i>include.fs</i> option.	373
<i>snapshotroot</i>	Command line only.	375
<i>subdir</i>	Client options file (dsm.opt) or command line.	381
<i>tapeprompt</i>	Client options file (dsm.opt) or command line.	386

*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. You can specify as many file specifications as available resources or other operating system limits permit. Separate file specifications with a space. You can also use the *filelist* option to process a list of files. The Tivoli Storage Manager client opens the file you specify with this option and processes the list of files within according to the specific command. See “Filelist” on page 262 for more information. When backing up a file system, specify the file system with a trailing slash.

## Examples

**Task** Back up the proja.dev file in the c:\devel directory.

**Command:** `se1 c:\devel\proj_a.dev`

**Task** Back up all files in the c:\devel directory whose file names begin with proj.

**Command:** `se1 c:\devel\proj*.*`

**Task** Back up all files in the c:\devel directory whose file names begin with proj. Back up all files with a file extension of .fin in the c:\planning directory.

**Command:** `se1 c:\devel\proj* c:\planning\*.fin`

**Task** Assuming that you initiated a snapshot of the C: drive and mounted the

snapshot as \\florence\c\$\snapshots\snapshot.0, run a selective backup of the c:\dir1\sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name C:

**Command:** dsmc sel c:\dir1\sub1\\* -subdir=yes -snapshotroot=\\florence\c\$\snapshots\snapshot.0



---

## 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.

When an existing file space is renamed during Unicode conversion, any access rules that are defined for the file space remain applicable to the original file space. However, new access rules must be defined to apply to the new Unicode file space.

## Supported Clients

This command is valid for all Windows clients.

## Syntax

```
▶▶ SET Access [ Archive | Backup ] [ filespec | {filespace}filespec ] node ▶▶
```

## 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 drive. However, you cannot use a wildcard to specify all drives. 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 `d:\test\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 `d:\test`, use file spec `d:\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.

**Notes:**

1. Use the file space name if the drive label name has changed.
2. If you include *filespace*name, do not include a drive letter in the file specification.

For example:

- Your directory structure is multilevel: `d:\test\sub1\subsub1`.
- The `d:\test` directory contains the `h1.txt` and `h2.txt` files.
- The `d:\test\sub1` directory contains file `s1.htm`.
- The `d:\test\sub1\sub2` directory contains the `ss1.cpp` file.

To allow access to all files in the `d:\test\sub1\sub2` directory, enter:

```
set access backup d:\test\sub1\sub2\* * *
```

To allow access to only those files in the `d:\test` directory, enter:

```
set access backup d:\test\* * *
```

To allow access to all files in all directories *in* and *below* the `d:\test` directory, enter:

```
set access backup d:\test\* * *
set access backup d:\test\*\* * * *
```

*{filespace*name}

Specifies the file space name (enclosed in braces) on the server that contains the files to which you are giving access. This name is the drive label name on the workstation drive from which the file was backed up or archived. Use the file space name if the drive label name has changed.

*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.

## Examples

**Task** Give the user at `node_2` authority to restore all files with an extension of `.c` from the `c:\devel\proja` directory.

**Command:** `set access backup c:\devel\proja\*.c node_2`

**Task** Give the user at `node_3` authority to retrieve all files in the `c:\devel` directory, but do not permit access to files in subdirectories of `c:\devel`, such as `c:\devel\proj`.

**Command:** `set access archive c:\devel\* node_3`

**Task** Give all nodes whose names end with bldgb the authority to restore all backup versions from all directories on the d: drive. The d: drive has the file space name of project.

**Command:** set ac b {project}\\*\\* \*bldgb

---

## 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 Windows 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 227)
- *Numberformat* (See “Numberformat” on page 316)
- *Noprompt* (See “Noprompt” on page 315)
- *Subdir* (See “Subdir” on page 381)
- *Timeformat* (See “Timeformat” on page 397)

## Examples

**Task** The following example displays the verbose and statistics output from the `set event` command `set event type=hold \\user\c$\tsm521\debug\bin\winnt_unicode\dsm.opt`, with objects rebound (as opposed to archived or some other notation):

```
Rebinding--> 274 \\user\c$\tsm521\debug\
  bin\winnt_unicode\dsm.opt
Rebinding--> 290 \\user\c$\tsm521\debug\
  bin\winnt_unicode\dsm.opt

Total number of objects inspected:      2
Total number of objects archived:      0
Total number of objects updated:       0
Total number of objects rebound:      2
Total number of objects deleted:       0
Total number of objects expired:       0
Total number of objects failed:       0
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\c$\tsm521\common\winnt` 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\c$\tsm521
   |                               \common\winnt
2. | 08/01/2003 10:38:11    766 B     \\user\c$\tsm521
   |                               \common\winnt
3. | 08/05/2003 08:47:46    5.79 KB   \\user\c$\tsm521
   |                               \common\winnt
4. | 08/01/2003 10:38:11    5.79 KB   \\user\c$\tsm521
   |                               \common\winnt
5. | 08/05/2003 08:47:46   10.18 KB  \\user\c$\tsm521
   |                               \common\winnt
```

---

## Set Password

The **set password** command changes the Tivoli Storage Manager password for your workstation. If you omit the old and new passwords when you enter the **set password** command, you are prompted once for the old password and twice for the new password.

A password is *not* case-sensitive, and it can be as many as 63 characters. Valid characters are:

<b>a-z</b>	Any letter, a through z, upper or lower-case
<b>0-9</b>	Any number, 0 through 9
<b>+</b>	Plus
<b>.</b>	Period
<b>_</b>	Underscore
<b>-</b>	Hyphen
<b>&amp;</b>	Ampersand

## Supported Clients

This command is valid for all Windows 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

If Websphere Application Server security is enabled, user name and password validation for Data Protection for WebSphere Application Server is required. If you do not set the Websphere Application Server password for the security, the backup will failover to an offline backup.

**Recommendation:** Set the Websphere Application Server security password to perform consistent backups.

Use the **set waspassword** command to set the user name and password for each installation of Websphere Application Server on your system. You only need to perform this task once, unless you change your Websphere Application Server user name or password. You can only perform this task on the Tivoli Storage Manager command line.

To determine if Websphere Application Server security is enabled, enter the following command:

```
dsmc query was -wast=local
```

Tivoli Storage Manager displays the Websphere Application Server security status under the **Sec** heading.

## Supported Clients

This command is valid for Windows Server 2003 clients.

## Syntax

```
▶▶—SET WASPassword—┐  
└─ WASNode— WASType— WASUser—┘▶▶
```

## Parameters

### *WASNode*

Specifies the node name on which each installation of Websphere Application Server 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 419 for more information.

### *WASType*

Specifies the Websphere Application Server 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 420 for more information.

### *WASUser*

Specifies Websphere Application Server 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 421 for more information.

## Examples

**Task:** Set the WebSphere user name and password for each installation of Websphere Application Server on your system.

**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

This section discusses how to use the Tivoli Storage Manager scheduler.

---

### Overview of the Tivoli Storage Manager scheduler

The Tivoli Storage Manager central scheduler allows client operations to occur automatically at specified times. In order to understand scheduling with Tivoli Storage Manager, several terms need to be defined:

#### schedule definition

A definition on the Tivoli Storage Manager server which specifies critical properties of the automated activity including the type of action, the time the action should take place, and how frequently the action will take place. There are numerous other properties which can be set (see the appropriate *Tivoli Storage Manager Administrator's Reference Guide* for a detailed description of the **define schedule**.)

#### schedule association

An assignment to a specific schedule definition for a client node. Multiple schedule associations allow single schedule definitions to be used by many client nodes. Because schedule definitions are included with specific policy domains, it is only possible for nodes defined to a certain policy domain to be associated with schedules defined in that domain.

#### scheduled event

A specific occurrence of when a schedule will be executed for a node. The following conditions must be met before automatic scheduled events will take place for a client:

- A schedule definition must exist for a specific policy domain.
- A schedule association must exist for the required node which belongs to that policy domain.
- The client scheduler process must be running on the client system (see "Setting the client scheduler process to run as a background task and start automatically at boot time" on page 575 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 drives defined by the client domain option. A schedule definition using the *command* action allows an operating system command or shell script to be executed. 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 quotation marks around each file specification that contains blanks, then single quotes around all of the specifications. Examples:

```
objects="C:\Program Files\MyApp\Some file.doc"
objects="C:\Program Files\MyApp\Some file.doc" "C:\Another file.txt"
C:\noblanks.txt'
objects="'E:\My Directory With Blank Spaces\"'
```

This will ensure that C:\Program Files\MyApp\Some file.doc is treated as a single file name, as opposed to two separate files (C:\Program Files\MyApp\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 options file (dsm.opt). 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>shedlogname</i>	X		
<i>shedlogretention</i>	X		
<i>shedmode</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.opt file. 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 *shedmode* 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.exe* command might exit with return code 3, if *widget.exe* 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.exe* might look like this:

```
run 'widget.exe'  
  if lastcc == 3  
    exit 0  
  else  
    exit 1
```

See the following references for more information:

- “*Postschedulecmd/Postnschedulecmd*” on page 325
- “*Preschedulecmd/Prenschedulecmd*” on page 329
- “Client return codes” on page 151

---

## Using the client acceptor to manage scheduler services versus the traditional scheduler services

You can configure the Tivoli Storage Manager client to manage the scheduler process using the Tivoli Storage Manager client acceptor. The client acceptor daemon provides a light-weight timer which automatically starts and stops the scheduler process as needed. Alternatively, the traditional method keeps the Tivoli Storage Manager scheduler process running continuously. Generally, using the client acceptor daemon to manage the scheduler is the preferred method. These methods are compared as follows:

### client acceptor daemon-managed Services

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

### Tivoli Storage Manager traditional scheduler services

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

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

---

## Setting the client scheduler process to run as a background task and start automatically at boot time

You can configure the Tivoli Storage Manager client scheduler to run as a background system task which starts automatically when your system is started. This is true for both client acceptor daemon-managed and traditional methods of running the Tivoli Storage Manager client scheduler. When running a client acceptor daemon-managed schedule, only the client acceptor daemon process should be set to start automatically at boot time; not the scheduler process. For the traditional method, the scheduler process should be set up to start automatically at boot time.

You can configure the client acceptor daemon to run as a background system task which starts automatically when your system is started. To configure the client acceptor daemon to manage scheduled backups, you must set the *managedservices* option to manage the scheduler, or both the scheduler and Web client. The method for setting up the client acceptor daemon as a system task varies for each platform.

In order for the scheduler to start unattended, you must enable the client to store its password by setting the *passwordaccess* option to *generate*, and store the password by running a simple Tivoli Storage Manager client command such as `dsmc query session`. Note that for testing purposes, you can always start the scheduler in the foreground by running `dsmc sched` from a command prompt (without a 'managedservices' stanza set).

On Windows platforms, the scheduler and client acceptor daemon run as services. You can create and manage these services using either the setup wizard or the Tivoli Storage Manager Client Service Configuration Utility, `dsmcutil.exe`.

- To start the setup wizard, select **Utilities** → **Setup Wizard** from the Tivoli Storage Manager backup-archive GUI.
- To start the Tivoli Storage Manager Client Service Configuration Utility, open a command prompt window and issue the following command:  

```
cd /d "c:\program files\tivoli\tsm\baclient"
```

Full documentation on how to use the command-line utility is available by entering `dsmcutil help`.

The Tivoli Storage Manager client scheduler can be managed by the client acceptor daemon. When setting up Tivoli Storage Manager scheduler services to run with client acceptor daemon management, two services must be created: the scheduler service and the client acceptor daemon service. When installing the client acceptor daemon service with `dsmcutil.exe`, use the `/cadschedname:` parameter to identify which Tivoli Storage Manager scheduler service the client acceptor daemon should manage. If you use the setup wizard to install the scheduler, you can select the **Use the CAD to manage the scheduler** checkbox, which will automatically create both services and associate them.

Using the Client Service Configuration Utility, you can either of the following methods:

#### **client acceptor daemon-managed**

1. In your client options file (`dsm.opt`), either set the *managedservices* option to *schedule* or *schedule webclient*.
2. In your client options file (`dsm.opt`), set the *passwordaccess* option to *generate*.
3. Create the scheduler service:  

```
dsmcutil inst /name:"TSM Client Scheduler" /node:tsmclient1 /password:secret /autostart:no /startnow:no
```
4. Create the client acceptor daemon and associate scheduler service with the client acceptor daemon:  

```
dsmcutil inst CAD /name:"TSM Client Acceptor" /cadschedname:"TSM Client Scheduler" /node:tsmclient1 /password:secret /autostart:yes
```
5. Manually start the client acceptor daemon service:  

```
net start "TSM Client Acceptor"
```

#### **Traditional**

1. In your client options file (`dsm.opt`), set the *managedservices webclient* option.
2. In your client options file (`dsm.opt`), set the *passwordaccess* option to *generate*.
3. Create the scheduler service:

```
dsmcutil inst /name:"TSM Client Scheduler" /node:tsmclient1
/password:secret /autostart:yes
```

To increase the reliability of the Tivoli Storage Manager client scheduler service on Windows, it is recommended that you set the services to automatically recover from a failure, as follows:

- Bring up the Windows services management console (**Control Panel** → **Administrative Tools** → **Services**)
- Right-click and select **Properties** for the "TSM Client Scheduler" service.
- Click the Recovery tab.
- Define the recovery action as **Restart the service** for first, second, and subsequent failures.

If you are using the client acceptor daemon 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. The same recovery settings can also be defined to increase the reliability of the "TSM Journal Service."

---

## Changing processing options used by the scheduler service

When you configure the Tivoli Storage Manager central-scheduling services (the scheduler, the client acceptor, or the remote client agent), some of the processing options that you specify are defined in the Windows registry. The following options can also be specified in the client options file (dsm.opt).

- nodename
- httpport
- tcpserveraddress
- tcpport
- webports

When the client scheduler runs as a foreground process using the **dsmc sched** command, the options in the client options file are used. However, when the scheduler runs as a Windows service, the options in the registry are used instead. If you are using the scheduler service and change an option in the dsm.opt file, you must update the corresponding value in the registry as well.

### To update the Windows registry value:

Use the Setup wizard in the client GUI. For more information, see "Configuring the client scheduler" on page 21.

Alternatively, you can use the dsmcutil utility to change the registry value. For example: `dsmcutil update scheduler /name: <service name> /node: <new node name> /password: <new node password>`. For more information, see "Using the Dsmcutil command" on page 589.

**Note:** After updating the registry, you must restart the scheduler service for the changes to take effect. If you are using client acceptor daemon managed scheduling this is not necessary because the scheduler is restarted by the client acceptor daemon for each backup.

---

## Managing multiple schedule requirements on one system

In certain situations it is preferable to have more than one scheduled activity for each client system. Normally, you can do this by associating a node with more than one schedule definition. This is the standard method of running multiple schedules on one system. You must ensure that the schedule windows for each schedule do not overlap. A single client scheduler process is not capable of executing multiple scheduled actions simultaneously, so if there is overlap, the second schedule to start will be missed if the first schedule does not complete before the end of the startup window of the second schedule. Suppose that most of the drives on your client system must be backed up daily, and that one drive containing critical data must be backed up hourly. In this case, you would need to define two schedules to handle this requirement. To avoid conflict between the hourly and daily backup schedule, the *starttime* of each schedule needs to be varied.

In certain cases, it is necessary to run more than one scheduler process on a system. Multiple processes require a separate options file for each process and must contain the following information:

- Define a unique node name for each process
- Specify unique schedule and error logs for each process
- When running in prompted mode, you must use the *tcpclientport* option to specify a unique port for each process.

**Note:** When the scheduler runs as a service, processing options specified in the Windows registry override the same options specified in the client options file. For more information, see “Changing processing options used by the scheduler service” on page 577.

The advantages of using multiple schedule processes:

- You can run more than one scheduled backup at the same time.
- You can specify different backup criteria for each schedule started, with the Tivoli Storage Manager client option file or Tivoli Storage Manager server override options.

The disadvantages of using multiple schedule processes:

- A unique file space for each node name on the Tivoli Storage Manager server is created.
- When restoring the data, you must use the same node name associated with the backup.

You must create a separate service for each schedule process. If you are using the client acceptor daemon to manage the scheduler, a client acceptor daemon service and schedule service are required for each schedule. The following is an example of setting up two schedule processes to be managed by the client acceptor daemon:

```
dsmcutil inst /name:"TSM Client Scheduler1" /optfile:"c:\tsm\dsm.opt1"  
/node:tsmcli_sched1 /password:secret /autostart:no /startnow:no
```

```
dsmcutil inst CAD /name:"TSM Client Acceptor1" /optfile:"c:\tsm\dsm.opt1"  
/cadschedname:"TSM Client Scheduler1" /node:tsmcli_sched1 /password:secret  
/autostart:yes
```

```
dsmcutil inst /name:"TSM Client Scheduler2" /optfile:"c:\tsm\dsm.opt2"  
/node:tsmcli_sched2 /password:secret /autostart:no /startnow:no
```



```
dsmcutil inst CAD /name:"TSM Client Acceptor2" /optfile:"c:\tsm\dsm.opt2"  
/cadschedname:"TSM Client Scheduler2" /node:tsmcli_sched2 /password:secret  
/autostart:yes
```

Unique option files are required for each schedule instance, and must be identified at the time of service creation:

#### Option file #1 (c:\tsm\dsm.opt1)

```
tcps                tmserv1.storage.sanjose.ibm.com  
nodename            tsmcli_sched1  
passwordaccess      generate  
schedlogname        c:\tsm\dsm\dsmsched1.log  
errorlogname        c:\tsm\dsm\dsmerror1.log  
schedmode           prompted  
tcpclientport       1507  
domain              h:  
managedservices     schedule
```

#### Option file #2 (c:\tsm\dsm.opt2)

```
tcps                tmserv1.storage.sanjose.ibm.com  
nodename            tsmcli_sched2  
passwordaccess      generate  
schedlogname        c:\tsm\dsm\dsmsched2.log  
errorlogname        c:\tsm\dsm\dsmerror2.log  
schedmode           prompted  
tcpclientport       1508  
domain              i:  
managedservices     schedule
```

---

## Scheduled backups of Windows network drives

See “Accessing Windows network drives during a scheduled backup” on page 150 for more information.

---

## Managing different scheduling requirements for Windows system objects

It is often preferable to back up Windows system objects less frequently than other data drives. In this case you can define two schedules on the Tivoli Storage Manager server which execute at different intervals of time. For example, the first schedule can back up all local drives every day, and a second schedule will backup only system objects once a week. This can be achieved by varying the *domain* option for each schedule. In order to reset the value of the domain option between each schedule operation, using the client acceptor daemon to manage the schedule is required. See “Setting up a schedule to back up system objects (Windows XP)” on page 588 for more information.

---

## Restarting the scheduler process on a remote system

When managing a large number of Tivoli Storage Manager clients running scheduler processes, it is helpful to be able to start and stop the client service from a remote system. The Tivoli Storage Manager client for Windows provides a utility to assist with remote management of the scheduler service.

To remotely manage the client scheduler service using the dsmcutil command with the /machine: option, you must have administrative rights in the domain of the

target system. To determine whether the scheduler service is running on a remote system, check the *Current Status* field from a query. For example, enter:

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

To restart a scheduler service that is missing schedules, enter:

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

If you are using the client acceptor daemon to manage the scheduler, you might need to restart the client acceptor daemon service, or stop the scheduler service and restart the client acceptor daemon service. For example, enter:

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

---

## 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.
- By default, Tivoli Storage Manager file space naming convention is MACHINE\_NAME\VOLUME\_NAME. 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. On Windows, it is impossible to backup local volumes when the *clusternode* option is set to *yes*. 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 or the registry. 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 53 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 D, “Configuring the backup-archive client in a cluster server environment,” on page 609.

## Troubleshooting the Tivoli Storage Manager scheduler

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

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

You can limit query results to exception cases with:

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

Query results include a status field which gives a summary of the result for a specific event. By using the *format=detailed* option on the **query event** command, complete information for events displays, including the return code passed back by the Tivoli Storage Manager client. Table 103 summarizes the meaning of the event status codes which are likely to exist for a scheduled event that has already taken place:

Table 103. Scheduled client event status codes

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

## Investigating abnormal scheduled events

If a particular client node displays multiple consecutive events with a result of *missed*, the client scheduler process is not running, has terminated, or is hanging. If a scheduled event is missed, but other consecutive scheduled events for that node show a result of *completed*, investigate the server activity log and the client schedule log to determine the cause. Scheduled events with a result of *failed*, have encountered a processing error originating either from the Tivoli Storage Manager client or server. The first place to check is the server activity log during the time period in which the scheduled event was processed. The activity log shows Tivoli Storage Manager server processing errors and some client errors which have been remotely logged to the Tivoli Storage Manager server. If the explanation cannot be found in the activity log, check the client schedule log.

## Checking the server activity log

When checking the server activity log, narrow the query results down to the time period in which the scheduled event occurred. Begin the event log query at a time shortly before the start window of the scheduled event. For example, if you are investigating the following suspect event:

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

You could use one of the following queries:

```
query actlog 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 system is required for inspecting the schedule log. The schedule log file typically resides in the same directory as the Tivoli Storage Manager client error log (dsmerror.log), in a file named TSM Schedule Log. The default location for this log is /Library/Logs/tivoli/tsm. The location of the log file can be specified using the *shedlogname* option, so you might need to refer to the options file to see if the *shedlogname* option was used to relocate the log file. The schedule log can also be relocated by an option setting which is part of the schedule service definition. The dsmcutil query command can be used to check if this option has been set. When you 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 *schelogname* option to change the location and name of this file.

The schedule log continues to grow in size unless you prune it using the *schelogretention* option or specify a maximum size using the *schelogmax* 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 179.



---

## Appendix B. Using the Client Service Configuration Utility

The following Tivoli Storage Manager services can be installed during installation or by using the Tivoli Storage Manager Client Service Configuration Utility:

- Backup-Archive Scheduler Service
- Client Acceptor Service (CAD)
- Remote Client Agent Service
- TSM Journal Engine Service

For more information about using the Tivoli Storage Manager Client Service Configuration Utility to install the Client Acceptor Service and Remote Client Agent Service, see “Using the Dsmcutil command” on page 589.

---

### Installing the backup-archive scheduler service

To install the Tivoli Storage Manager scheduler, use one of the following methods:

- From the Tivoli Storage Manager GUI click **Utilities** and then click **Setup Wizard**. Select the **Help me configure the TSM Client Scheduler** option.
- If you have an account that belongs to the Administrator/Domain Administrator group, you can use the **TSM Client Service Configuration Utility** to configure client services on both local and remote Windows machines.

### Using the Client Service Configuration Utility

This section provides the following:

- An example of using the Client Service Configuration Utility with the backup-archive client to automate backups.
- An example of configuring the CAD to manage an existing scheduler service.
- An example of creating a new scheduler and associating a CAD to manage the scheduler:

This example illustrates the use of the Tivoli Storage Manager scheduler.

When the Tivoli Storage Manager backup-archive client is registered with a Tivoli Storage Manager server, the procedure involves the following steps:

1. **On the server:**
  - a. Define a schedule for the policy domain to which the backup-archive client is registered.
  - b. Associate the backup-archive client node to the defined schedule.
2. **On the backup-archive client:**
  - a. Install the Tivoli Storage Manager scheduler client as a Windows service for the backup-archive client.
  - b. Start the scheduler service installed for the backup-archive client.

### Automating backups

The example below uses the following assumptions:

- The Tivoli Storage Manager backup-archive client is registered to a Tivoli Storage Manager server with a node name of **mars** and a password of **marspswd** in policy domain **bacliwnt**.
- The event to be scheduled is a daily incremental backup of file systems on client machines. The backup begins between 9:00 and 9:15 pm.

- The backup-archive client is installed to the c:\program files\tivoli\tsm\baclient directory.
- The communication parameters in the backup-archive client options file (dsm.opt) are appropriate for the Tivoli Storage Manager server.

#### On the server:

1. Enter the following command on the server console or from an administrative client to define the schedule. The administrative client does not have to be running on the same system as the Tivoli Storage Manager server.

```
def sched bacliwnt wnt_daily_incr desc="Daily Incremental Backup"
priority=2 starttime=21:00 duration=15 durunits=minutes period=1
perunits=days dayofweek=any
```

Tivoli Storage Manager displays this message:

```
ANR2500I Schedule WNT_DAILY_INCR defined in policy domain BACLIWNT.
```

2. To associate the backup-archive client to this schedule, issue the following command:

```
define association bacliwnt wnt_daily_incr mars
```

Tivoli Storage Manager displays this message:

```
ANR2510I Node MARS associated with schedule WNT_DAILY_INCR in
policy domain BACLIWNT.
```

At this point, a schedule that performs an incremental backup is defined on the Tivoli Storage Manager server. The schedule starts around 9:00 pm. The schedule is re-executed once a day and can start on any day of the week.

**Note:** If you want to confirm that the schedule and association are set correctly, you can use the **Query Schedule** command. For additional information, see “Query Schedule” on page 511.

#### On the backup-archive client:

This example assumes that you installed the Tivoli Storage Manager backup-archive client in the c:\program files\tivoli\tsm\baclient directory. It also assumes that the options files in each of these directories are updated so that the communication parameters point to the Tivoli Storage Manager server.

1. Login using an account with administrative privileges.
2. Open a command prompt window and issue the following command:

```
cd /d "c:\program files\tivoli\tsm\baclient"
```

If the path contains a space, for example c:\program files\tivoli\tsm\baclient, enclose the name in double quotes.

3. In the window, issue the following command:

```
dsmcutil inst scheduler /name:"TSM Client Scheduler"
/node:mars /password:marspswd /clientdir:"c:\program files\
tivoli\tsm\baclient" /optfile:"c:\program files\tivoli\tsm\
baclient\dsm.opt" /autostart:yes
```

Your system is now ready to run automatic daily incremental backups. The */autostart:yes* option specifies that the scheduler service starts automatically each time the system is rebooted. You can use the */startnow:[Yes|No]* option to specify whether to start the scheduler service after executing the command; the default is *Yes*. If you specify */startnow:No* you must start the service manually using the services control panel applet on NT, or issue the following command:

```
net start "TSM Client Scheduler"
```



For additional information about these options, see “Dsmcutil options” on page 596.

4. The scheduler uses the backup-archive client options file to validate the node and password, and to contact the server for schedule information. This example assumes that the dsm.opt file is updated so that the communication parameters point to the Tivoli Storage Manager server.

If you see the following message:

```
A communications error occurred connecting to the TSM server
```

You should ensure that the options file contains entries that point to the correct Tivoli Storage Manager server. You should also ensure that the server is running. Use the dsmcutil **update** command to correct one of the parameters which was incorrectly specified during the dsmcutil install. For example, to update the client directory and options file for the specified scheduler service, enter:

```
dsmcutil update scheduler /name:"TSM Central Scheduler Service"  
/clientdir:c:\tsmptf3\baclient /optfile:c:\tsmptf3\baclient\dsm.opt
```

Then reissue the **net start "TSM Client Scheduler"** command.

#### Notes:

1. If any changes that affect the scheduler service are made to the backup-archive client options file, the scheduler service has to be restarted. If you are using CAD managed scheduling this is not necessary since the scheduler is restarted by the CAD for each backup and the changes are picked up. An example of this is the Tivoli Storage Manager server address or the schedule mode. This can be done by issuing the following commands:

```
net stop "TSM Client Scheduler"  
net start "TSM Client Scheduler"
```

2. The dsmsched.log file contains status information for the Tivoli Storage Manager scheduler service. In this example, the file is located in this path:

```
c:\program files\tivoli\tsm\baclient\dsmsched.log
```

You can override this file name by specifying the *shedlogname* option in the options file, dsm.opt.

3. Output from scheduled commands is sent to the 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 might contain 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:
```

## Configuring the CAD to manage an existing scheduler service

The example below assumes that the scheduler service name is TSM Central Scheduler Service and the CAD service name is TSM Client Acceptor, which are the default names. You can use the dsmcutil /name option to specify different names. See “Dsmcutil commands” on page 590 and “Dsmcutil options” on page 596 for more information about dsmcutil commands and options.

To configure the CAD to manage an existing scheduler service:

1. Stop the scheduler service and the CAD, as follows:  

```
dsmcutil stop /name:"tsm central scheduler service"  
dsmcutil stop /name:"tsm client acceptor"
```
2. Set the *managedservices* option to *schedule* in the client options file (dsm.opt).
3. Update the scheduler service so that it does not start automatically after a reboot:  

```
dsmcutil update /name:"tsm central scheduler service"  
/autostart:no
```
4. Associate the scheduler service with the CAD:  

```
dsmcutil update cad /name:"tsm client acceptor"  
/cadschedname:"tsm central scheduler service" /autostart:yes
```

**Note:** If this command is successful, the dsmwebcl.log file will include this message: Command will be executed in 1 minute. After one minute, the CAD will start the scheduler and you will see information regarding the next scheduled event in the dsmwebcl.log file.

### Creating a new scheduler and associating a CAD to manage the scheduler

To create a new scheduler and associate a CAD to manage the scheduler:

1. Set the *managedservices* option to *schedule* in the client options file (dsm.opt).
2. Create the scheduler service:  

```
dsmcutil install scheduler /name:"NEW_SCHEDULE_NAME" /node:yournode  
/password:xxxxx /autostart:yes /startnow:no
```
3. Create the CAD service:  

```
dsmcutil install cad /node:yournode /password:xxxxx  
/autostart:yes /startnow:no
```
4. Associate the scheduler with the CAD:  

```
dsmcutil update cad /name:"tsm client acceptor"  
/cadschedname:"NEW_SCHEDULE_NAME"
```
5. Start the CAD:  

```
dsmcutil start cad /name:"tsm client acceptor"
```

**Note:** The CAD and scheduler will start, as described above. Note that since the CAD is controlling the scheduler, you will not see the scheduler running as a service, either through the Services applet or the NET START command. To stop the scheduler, you must stop the CAD.

## Setting up a schedule to back up system objects (Windows XP)

This section describes how to set up a Tivoli Storage Manager client schedule to back up file systems on weekdays and back up the SYSTEMOBJECT domain on Saturdays. Both schedules will run in CAD managed mode.

1. To setup a schedule to back up your file systems on weekdays, enter:  

```
define schedule your_domain weekday_backup action=incremental option="-domain=  
-systemobject" starttime=20:00:00<other options here> dayofweek=weekday
```

To associate the backup-archive client to this schedule, enter:  

```
define assoc your_domain weekday_backup nodename
```

To setup a schedule to back up your SYSTEMOBJECT domain on Saturdays, enter:

```
define schedule your_domain saturday_backup action=incremental option="-domain=
systemobject" starttime=21:00:00<other options here> dayofweek=saturday
```

To associate the backup-archive client to this schedule, enter:

```
define assoc your_domain saturday_backup nodename
```

**Note:** The subtle difference between the two schedule definitions is that the weekday schedule has a minus ('-') in front of "systemobject" to explicitly remove the SYSTEMOBJECT domain.

2. The Tivoli Storage Manager client option file (dsm.opt) changes are:

```
domain c: d: <other domains>
```

3. Use the setup wizard to define a schedule service to run in CAD managed mode, update your existing schedule service to become CAD managed, or delete the current service and create a new one that is CAD managed. See "Configuring the client scheduler" on page 21 for more information.


---

## Using the Dsmcutil command

The Tivoli Storage Manager Client Service Configuration Utility, **dsmcutil**, allows Tivoli Storage Manager Client Services installation and configuration on local and remote Windows machines. You can install the following Tivoli Storage Manager services under a double-byte character set (DBCS) path:

- Backup-Archive Scheduler Service
- Client Acceptor Service (CAD)
- Remote Client Agent Service
- Journal Engine Service

The Client Service Configuration Utility must be run from an account that belongs to the Administrator/Domain Administrator group. The syntax of the utility is:



### Notes:

1. For Windows Vista, you must right click on the command prompt and select "Run as administrator".
2. Options that you specify with dsmcutil commands override those in your options file (dsm.opt). See Chapter 2, "Configuring Tivoli Storage Manager," on page 15 for more information.

The account running the utility must have the appropriate user rights for installing services and updating the Windows Registry on the target machine.

If a remote machine is specified, the account must be authorized to connect to the specified machine's Windows Registry.

**Note:** For the commands and options documented here, the minimum abbreviation is in upper case. Upper case letters denote the shortest acceptable truncation. If an item appears entirely in upper case letters, it cannot be truncated. You can type the item in any combination of upper case or lower case letters.

## Dsmcutil commands

### INSTall

Installs and configures a Tivoli Storage Manager Scheduler Service. Required options are:

- */name: servicename*
- */password: password*
- */clusternode: Yes | No* (required if running the Microsoft Cluster Server (MSCS) or Veritas Cluster Server (VCS)).
- */clustername: clustername* (required if running the MSCS or VCS).

**Recommendation:** Do not specify a clustername of more than 64 characters. If you specify more than 64 characters and you are using Veritas Storage Foundation with High Availability or a Microsoft Cluster Server configuration, you might not be able to install or start the Tivoli Storage Manager scheduler service.

The */clientdir: clientdir* option can also be used, the default is the current directory. For more information about these options, see “Dsmcutil options” on page 596.

The following files must exist in the directory specified for */clientdir: clientdir*:

- dsmcsvc.exe
- dscenu.txt
- dsm.opt
- dsmntapi.dll
- tsmutil1.dll

**Note:** If the service is being installed on a remote machine, the fully qualified client directory path should be relative to the target machine. UNC names are not allowed for the local system account. Multiple services can be installed on the same machine.

**Task** Install a scheduler service named **TSM Central Scheduler Service** on the local machine. Start the service automatically at system boot time. All required files must reside in the current directory and the client options file must point to the Tivoli Storage Manager server where node ALPHA1 is defined with password nodepw. The Tivoli Storage Manager server is contacted to verify that the specified node and password are valid. When the password is validated it is generated (encrypted) into the Windows Registry:

**Command:**

```
dsmcutil install scheduler /name:"TSM Central Scheduler Service"  
/node:ALPHA1 /password:nodepw /autostart:yes
```

**Task** Install a scheduler service named **TSM Central Scheduler Service** on remote machine PDC. Start the service automatically at system boot time. The required scheduler service files and the specified options file must reside on the remote machine in the c:\tsm\baclient directory. The password is encrypted into the Windows Registry. The Tivoli Storage Manager server is not contacted to validate the password.

**Command:**

```
dsmcutil install scheduler /name:"TSM Central Scheduler Service"  
/machine:PDC /clientdir:c:\tsm\baclient /optfile:c:\tsm\baclient\dsm.opt  
/node:PDC /validate:no /autostart:yes  
/password:nodepassword
```

**Task** Install a scheduler service named **TSM Central Scheduler Service** on

remote machine PDC. Start the service automatically at system boot time. The required scheduler service files and the specified options file must reside on the remote machine in the `c:\tsm\baclient` directory. The password is encrypted into the Windows Registry. The Tivoli Storage Manager server residing at the specified TCP/IP host and port is contacted to validate the password.

**Command:**

```
dsmcutil install scheduler /name:"TSM Central Scheduler Service"  
/machine:PDC /clientdir:c:\tsm\baclient /optfile:c:\tsm\baclient\dsm.opt  
/node:PDC /autostart:yes /password:nodepassword  
/commmethod:tcpip /commserver:alpha1.sanjose.com  
/commport:1521
```

**Task** Install the **TSM Central Scheduler Service** on one node of a MSCS (or VCS) cluster. For *group-a* from machine *node-1*, ensure that *node-1* currently owns *group-a* and then issue the command.

**Command:**

```
dsmcutil install scheduler /name:"TSM Central Scheduler Service:  
group-a" /clientdir:c:\tsm\baclient /optfile:q:\tsm\baclient\  
dsm.opt /node:mscs-cluster-group-a /password:n  
/validate:no /autostart:yes /startnow:yes  
/clusternode:yes /clustername:mscs-cluster
```

## INSTall CAD

Installs and configures the Client Acceptor Service (CAD). Required options are:

- */name:servicename*
- */node:nodename*
- */password:password*

Other valid options are:

- */optfile:optionsfile*
- */httpport:httpport*
- */webports:webports*

For more information about these options, see “Dsmcutil options” on page 596.

**Task** Install a Client Acceptor Service called TSM CAD. TSM CAD uses a node called **test** to connect to the Tivoli Storage Manager server. Use the options file `c:\tsm\baclient\dsm.opt` to connect to Tivoli Storage Manager.

**Command:**

```
dsmcutil install cad /name:"TSM CAD" /node:test  
/password:test /optfile:c:\tsm\baclient\dsm.opt
```

## INSTall Journal

Installs a journaling engine service on all Windows clients, except for clients running on Windows Server 2003 for Itanium-based Systems. A journal database is created that stores information the client uses to determine which files are eligible for backup *before* an operation starts. See “Journal-based backup” on page 66 for more information on performing journal-based backups.

If necessary, you can use the *nojournal* option with the **incremental** command to specify that you want to perform a traditional full incremental backup. See “Incremental” on page 475 for more information.

The journaling engine service is named **TSM Journal Service** and uses the configuration file `tsmjbbd.ini` from the Tivoli Storage Manager installation directory. See Appendix C, “Journal service configuration,” on page 601 for more information.

**Note:** The Journal Service is supported in a Microsoft Cluster Server environment. Multiple journal services can be installed by specifying unique pipe names using the `JournalPipe` journal config setting and client options.

There are no valid options for this command.

**Task** Install the journaling engine service **TSM Journal Service**.

**Command:**

```
dsmcutil install journal
```

### INSTAll REMOTEAgent

Installs and configures a Remote Client Agent Service. Required options are:

- */name:servicename*
- */node:nodename*
- */password:password*
- */partnername:partner service name*

Other valid options are:

- */optfile:optionsfile*

For more information about these options, see “Dsmcutil options” on page 596.

**Task** Install a Remote Client Agent Service called TSM AGENT. TSM AGENT uses a node called **test** to connect to the Tivoli Storage Manager server. The options file `c:\tsm\baclient\dsm.opt` is used to connect to Tivoli Storage Manager. The partner CAD service is TSM CAD.

**Command:**

```
dsmcutil install remoteagent /name:"TSM AGENT" /node:test  
/password:test /optfile:c:\tsm\baclient\dsm.opt /partnername:  
"TSM CAD"
```

**Note:** Both the Remote Client Agent Service and the Client Acceptor Service must be installed to run the Web client. The Client Acceptor Service must be installed *before* the Remote Client Agent Service. Use the */partnername:* option to specify the name of the partner Client Acceptor Service.

### REMOve

Remove an installed Client Service. The required option is */name:servicename*. For more information about this option, see “Dsmcutil options” on page 596.

**Task** Remove the specified scheduler service from the local machine.

**Command:**

```
dsmcutil remove /name:"TSM Central Scheduler Service"
```

**Task** Remove the journaling engine service **TSM Journal Service** from the local machine.

**Command:**

```
dsmcutil remove /name:"TSM Journal Service"
```

## UPDate

Updates Scheduler Service registry values. The required option for this command is */name:servername*, and the registry values to update. Other valid options are:

- */clientdir:clientdir*
- */optfile:optionsfile*
- */eventlogging:[Yes | No]*
- */node:nodename*
- */autostart:[Yes | No]*
- */clusternode:Yes | No* (required if running the MSCS or VCS).
- */clustername:clustername* (required if running the MSCS or VCS).

For more information about these options, see “Dsmcutil options” on page 596.

**Task** Update the client directory and options file for the specified scheduler service. All required client service files must reside in the specified directory.

**Note:** The communication options specified with the **dsmcutil** command here take precedence over those specified in the client options file.

**Command:**

```
dsmcutil update /name:"TSM Central Scheduler Service"  
/clientdir:c:\tsmptf3\baclient /optfile:c:\tsmptf3\baclient\dsm.opt
```

**Task** Update the specified scheduler service to use the TCP/IP protocol to connect to a Tivoli Storage Manager server at the specified host name on the specified port.

**Command:**

```
dsmcutil update /name:"TSM Central Scheduler Service"  
/commsserver:nt1.sanjose.com /commport:1521 /commmethod:  
tcpip
```

## UPDate CAD

Updates Client Acceptor Service registry values. The required option for this command is */name:servername*, and the registry values to update. Other valid options are:

- */node:nodename*
- */password:password*
- */optfile:optionsfile*
- */httpport:httpport*
- */webports:webports*
- */cadschedname:schedulername*

For more information about these options, see “Dsmcutil options” on page 596.

**Task** Update the Client Acceptor Service to use the specified client password and options file. All required client service files must reside in the specified directory.

**Command:**

```
dsmcutil update cad /name:"TSM CAD" /password:test  
/optfile:c:\tsm\baclient\dsm.opt
```

## UPDate REMOTEAgent

Updates Remote Client Agent Service registry values. The required option for this command is */name:servicename*, and the registry values to update. Other valid options are:

- */node:nodename*
- */password:password*
- */optfile:optionsfile*
- */partnername:partner service name*

For more information about these options, see “Dsmcutil options” on page 596.

**Task** Update a Remote Client Agent Service called TSM AGENT. TSM AGENT uses a node called **test** to connect to the Tivoli Storage Manager server. The options file `c:\tsm\baclient\dsm.opt` is used to connect to Tivoli Storage Manager. The partner CAD service is TSM CAD.

**Command:**

```
dsmcutil update remoteagent /name:"TSM AGENT" /node:test  
/password:test /optfile:c:\tsm\baclient\dsm.opt /partnername:  
"TSM CAD"
```

### Query Scheduler

Query Scheduler Service registry values. Required options are: */name:servicename*. Other valid options are:

- */machine:machinename*
- */clientdir*
- */optfile*
- */eventlogging*
- */node*
- */commmethod*
- */commport*
- */commserver*
- */errorlog*
- */schedlog*

**Note:** Do not specify a value for the non-required options. Tivoli Storage Manager returns option registry values for the scheduler service you specify.

For more information about these options, see “Dsmcutil options” on page 596.

**Task** Query registry settings for the scheduler service you specify.

**Command:**

```
dsmcutil query /name:"TSM Central Scheduler Service"
```

**Task** Query the client directory registry setting for the scheduler service you specify.

**Command:**

```
dsmcutil query /name:"TSM Central Scheduler Service"
```

### Query CAD

Queries Client Acceptor Service registry values. The required option for this command is */name:servicename*. Other valid options are:

- */machine:machinename*
- */node*
- */optfile*
- */httpport*



- */webports*
- */clientdir*
- */partnername*

**Note:** Do not specify a value for these options.

For more information about these options, see “Dsmcutil options” on page 596.

**Task** Query registry settings for the Client Acceptor Service you specify.

**Command:**

```
dsmcutil query cad /name:"TSM CAD"
```

### Query Journal

Query the journaling engine service, **TSM Journal Service** on a Windows (32-bit). There are no valid options for this command.

**Task** Query the journaling engine service, **TSM Journal Service**.

**Command:**

```
dsmcutil query journal
```

### Query REMOTEAgent

Queries Remote Client Agent Service registry values. The required option for this command is */name:servername*. Other valid options are:

- */machine:machinename*
- */node*
- */optfile*
- */partnername*
- */clientdir*

**Note:** Do not specify a value for these options.

For more information about these options, see “Dsmcutil options” on page 596.

**Task** Query registry settings for the specified Remote Client Agent Service.

**Command:**

```
dsmcutil query remoteagent /name:"TSM AGENT"
```

### List

Lists installed Client Services. There are no required options.

**Task** Locate and list the installed Tivoli Storage Manager client services on the local machine.

**Command:**

```
dsmcutil list
```

**Task** List the installed Tivoli Storage Manager client services on remote machine PDC.

**Command:**

```
dsmcutil list /MACHINE:PDC
```

### START

Use the **Start** command to start a client service. The **Start** command requires the */name: servicename* option.

**Task** Start the journaling engine service, **TSM Journal Service**.

**Command:**

```
dsmcutil start /name:"TSM Journal Service"
```

## STOP

Use the **Stop** command to stop a client service. The **Stop** command requires the */name: servicename* option.

**Task** Stop the journaling engine service, **TSM Journal Service**.

**Command:**

```
dsmcutil stop /name:"TSM Journal Service"
```

## UPDATEPW

Generate encrypted Tivoli Storage Manager registry password. The **UPDATEPW** command requires the */node: nodename* and */password: password* options. Optionally, you can use the following options:

- */validate: [Yes | No]*
- */clusternode: Yes | No* (required if running the MSCS or VCS).
- */clustername: clustername* (required if running the MSCS or VCS).
- */force: [Yes | No]*

The password is validated with the Tivoli Storage Manager server if the */validate: Yes* option is specified. The password is updated on the server if you specify the */updateonserver: Yes* option. If you specify this option, you must specify the current password with the */oldpassword* option. For more information about these options, see "Dsmcutil options."

**Task** Update the encrypted registry password for the specified node. Validate and update the password on the specified Tivoli Storage Manager server which resides on the specified TCP/IP hostname and port:

**Command:**

```
dsmcutil updatepw /node:alpha1 /commMethod:tcpip  
/commServer:alpha1.sanjose.ibm.com /commPort:1500  
/password:newpw /oldpassword:oldpw /updateonserver:yes
```

## Dsmcutil options

The following options are valid:

*/autostart: [Yes | No]*

Specifies whether the Scheduler Service starts automatically at system boot time. The default is *No*.

*/cadschedname: schedulername*

Specifies the name of the scheduler service to manage with the CAD. Use this option when the *managedservices* option is set to *schedule* in the client options file *dsm.opt*. You can specify this option only with the CAD service.

*/clientdir: clientdir*

The fully qualified directory path where the Client Service files reside. This directory should be relative to the target machine where the service is

installed. UNC names are not allowed if the local system account is set to logon. The default is the current directory.

*/clustername:clustername*

This option replaces the */group* option.

The */clustername* option specifies the cluster name to which the system belongs. You can determine the cluster name in any of the following ways:

- On MSCS, run the MSCS command, `CLUSTER /LIST`, from the command line or use the Cluster Administrator utility. When the Cluster Administrator utility starts, it will display a tree-like structure with the cluster name at the top.
- On VCS, use the VCS Cluster Manager - Java Console or open the `main.cf` file in the `%VCS_HOME%\config` directory.
- On VCS, use the following command:

```
haclus -display
```

**Recommendation:** Do not specify a clustername of more than 64 characters. If you specify more than 64 characters and you are using Veritas Storage Foundation with High Availability or a Microsoft Cluster Server configuration, you might not be able to install or start the Tivoli Storage Manager scheduler service.

This option must be used with the */clusternode:Yes* option. This option must be specified when using the `INSTALL` command in a cluster environment. It must also be specified when using the `UPDATE` command to modify the cluster settings (*/clusternode* and */clustername*).

This option can also be specified when using the `UPDATEPW` command in a cluster environment. Normally this is not required. However, if more than one scheduler service with different cluster settings are defined for a particular node, the utility cannot determine which settings are correct. In this case, the recommended action is to correct the discrepancies between the services.

Alternatively, you can specify this option with */clusternode:Yes* and */force:Yes*, to force the utility to show or update the password with the specified cluster settings.

This option is not required if */clusternode:No* is specified.

*/clusternode:Yes | No*

Specifies whether to enable support for cluster resources. The default value is *No*. You must be running the MSCS or VCS to specify */clusternode:Yes*. This option must be specified when using the `INSTALL` command in a cluster environment. This option must also be specified when using the `UPDATE` command to modify the cluster settings (*/clusternode*, */clustername*).

This option can also be specified when using the `UPDATEPW` command in a cluster environment. Normally this is not required. However, if more than one scheduler service with different cluster settings are defined for a particular node, the utility cannot determine which settings are correct. In this case, the recommended action is to correct the discrepancies between the services.

Alternatively, you can specify this option with */clustername* and */force:Yes*, to force the utility to show or update the password with the specified cluster settings. If */clusternode:No* is specified, */clustername* is not required.

*/commmethod:protocol*

Specifies client communications protocol to communicate with a Tivoli Storage Manager server. Valid protocols are: TCP/IP and Named Pipes. If you do not specify a value, the value is obtained from the client options file or set to the default client value. You can also use this option with the UPDATEPW command to specify a communication protocol to connect to a server when updating passwords.

*/commport:serverport*

Specifies the protocol specific Tivoli Storage Manager server port. For TCP/IP, this is the port on the specified hostname. If this option is not specified, the value is obtained from the client options file or set to the default client value. You can also use this option with the UPDATEPW command to specify a protocol specific server port to connect to for updating passwords.

*/commserver:servername*

Specifies the protocol specific Tivoli Storage Manager server name. Depending on the protocol used, this can be a TCP/IP hostname or a Named Pipes name. If not specified, the value is obtained from the client options file or set to the default client value.

This option can also be used with the UPDATEPW command to specify a protocol specific server name to connect to for updating passwords.

*/copyfiles*

Specifies that the service installation is copied to another location prior to installing the service. Use the */srcdir* option to specify the fully qualified source path.

*/errorlog:errorlog*

Specifies the fully qualified name of the client error log.

*/eventlogging:[Yes|No]*

Turns detailed event logging on or off for the specified scheduler service. The default is *Yes*.

*/force:[Yes|No]*

This option can also be specified when using the UPDATEPW command in a cluster environment. Normally this is not required. However, if more than one scheduler service with different cluster settings is defined for a particular node, the utility cannot determine which settings are correct. In this case, the recommended action is to correct the discrepancies between the services.

Alternatively, you can specify this option with */clusternode* and */clusternode* (if */clusternode:Yes* is specified), to force the utility to show or update the password with the specified cluster settings.

*/httpport:httpport*

Specifies a TCP/IP port address for the Web client.

*/machine:machinename*

Specifies the name of a remote machine to connect to.

*/name:servicename*

Specifies the name of the Client service. The name must be quote delimited if it contains embedded spaces.

*/node:nodename*

Specifies the Tivoli Storage Manager node name the Client Service uses

when connecting to the Tivoli Storage Manager server. Also used when displaying or updating the Tivoli Storage Manager registry password. The default is the machine name.

*/ntaccount:ntaccount*

Specifies the Windows account which the service logs in as.

*/ntdomain:ntdomain*

Specifies the Windows domain which the service logs in as.

*/ntpassword:ntpassword*

Specifies the Windows password for the account under which the service logs in.

*/oldpassword:oldpw*

Current Tivoli Storage Manager server password. Used in conjunction with the */updateonserver* option when updating a password on the server.

*/optfile:optionsfile*

The fully qualified path of the Tivoli Storage Manager options file. This is the options file the specified Client Service uses to connect to Tivoli Storage Manager. The utility also uses the file to connect to Tivoli Storage Manager to validate and update passwords. Note that although this option overrides the default option file in the current directory (*dsm.opt*), the Tivoli Storage Manager API requires that a default option file exists in the current directory. UNC names are not allowed if the local system account is set to logon. The default is the *dsm.opt* file in the */clientdir* directory.

*/partnername:partner service name*

This option is used when installing a Remote Client Agent Service to specify the partner Client Acceptor Service.

*/password:password*

The Tivoli Storage Manager password which is generated (encrypted) into the Windows Registry.

*/schedlog:schedlog*

Specifies the fully qualified name of the client schedule log.

*/srcdir:pathname*

Use this option in conjunction with the */copyfiles* option to specify the fully qualified source path to copy the service installation to another location prior to installing the service.

*/startnow:[Yes | No]*

Specifies whether *dsmcutil* starts the specified service after executing the command; the default is *Yes*. If you specify *No*, you must start the service manually using the services control panel applet, or the NET START **name of the Tivoli Storage Manager service**.

*/updateonserver:[Yes | No]*

Specifies whether the specified password is updated on the Tivoli Storage Manager server. Requires using the */oldpassword* option.

*/validate:[Yes | No]*

Specifies whether to perform validation when displaying or updating the encrypted registry password. The default is *Yes*.

*/webports: webports*

Specifies the TCP/IP port number used by the Tivoli Storage Manager Client Acceptor service and the Tivoli Storage Manager Web client agent service for communications with the Web GUI.



---

## Appendix C. Journal service configuration

Journal-Based Backup is enabled by installing and configuring the Tivoli Storage Manager Journal Service. The Tivoli Storage Manager Journal Service can be installed with the GUI Setup wizard or with the **dsmcutil** command. Basic Journal Service configuration can be done with the GUI Setup wizard, and more advanced configuration can be done by editing the Journal Service configuration file, `tsmjbbd.ini`. See Appendix B, "Using the Client Service Configuration Utility," on page 585 for more information about using the **dsmcutil** command to install the Tivoli Storage Manager Journal Service.

The primary purpose of the client GUI setup wizard is to install and configure the journal engine service and to provide basic configuration settings. You must specify more advanced journal service configuration settings in the journal configuration file `tsmjbbd.ini`. The `tsmjbbd.ini` file is a Windows style stanza based `.ini` file and should always reside in the same directory as the journal service executable module (`tsmjbbd.exe`).

Follow these steps to set up multiple journal services:

1. Create and set up a separate journal configuration file (`tsmjbbd.ini`) for each journal service to be installed. Each configuration file must specify a different "JournalPipe" value, and must also specify different drives to journal, so that the two services do not interfere with each other. Multiple journal services journaling the same drive will cause problems. The different services will attempt to write to the same journal database unless this is specifically overridden by specifying different journal directories in the different configuration files.
2. Install the multiple journal services using the `dsmcutil.exe` tool. Use distinct names for each service, and specify the `/JBBCONFIGFILE` option to identify the `tsmjbbd.ini` to be used for that particular journal instance. For example:

```
dsmcutil install journal /name:"TSM Journal Service 1"  
/JBBCONFIGFILE:c:\journalconfig\tsmjbbd1.ini  
dsmcutil install journal /name:"TSM Journal Service 2"  
/JBBCONFIGFILE:d:\journalconfig\tsmjbbd2.ini
```
3. Different backup clients (based on the distinct `dsm.opt` file used) can now connect to the desired journal service by specifying the appropriate `JournalPipe` option in the appropriate `dsm.opt`, which corresponds to the `JournalPipe` journal service setting configured in 1.

### Notes:

1. Each journal service instance is associated to *only* one Tivoli Storage Manager client node name. Changing the association requires a restart of the journal service to recognize the new association.
2. For customers already using journal-based backup, the first incremental backup after installing the Tivoli Storage Manager Version 5.5 client 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 5.5 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.
3. Network and removable file systems are not supported.

Configuration settings that you apply when the journal service is started and any changes you make while the journal service is running are applied without having to restart the service. This also applies to the journal exclude list. However, some settings for journaled file systems do not take effect until the file system is brought offline and then back online.

File systems can be brought online (added) or offline (removed) without stopping and restarting the journal service. You can bring a file system offline by removing it from the list of journaled file systems in the journal configuration file `tsmjbbd.ini`, or by shutting down the journal service. You can bring a file system back online by adding it to the list of journaled file systems in the journal configuration file `tsmjbbd.ini` or by starting (restarting) the journal service.

**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 service 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 service.

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:

*JournalPipe=pipename*

Specifies the pipe name of the journal service session manager to which backup clients initially connect, when establishing a journal-based backup session. This setting is used in conjunction with the backup client option of the same name. The default pipe name is `\\.\pipe\jn1Server`. For example, in `dsm.opt`:

```
JournalPipe \\.\pipe\jn1Server1
```

Under `tsmjbbd.ini` *[JournalSettings]* stanza:

```
JournalPipe=\\.\pipe\jn1Server1
```



**Note:** The same pipe name must be specified by the client using the *JournalPipe* option.

#### *NlsRepos*

Specifies the National Language Support repository the journal service uses for generating messages. Since the journal service is non-interactive, this only applies to messages written to the journal error log. The default value is `dscameng.txt`. For example:

```
NlsRepos=dscenu.txt
```

#### *ErrorLog*

Specifies the log file where detailed error messages generated by the journal service are written. Note that less detailed error and informational messages are written to the Windows application event log as well. The default value is `jbberror.log`. For example:

```
ErrorLog=jbberror.log
```

#### *JournalDir*

Specifies the directory where journal database files are stored and written. The default directory is the journal service installation directory. You can specify different journal locations for each file system being journaled. This is useful when running in a clustered environment because the location of the journal must be accessible by each machine in the cluster running the journal service. Typically the journal for local resources being journaled will reside in the same location and the journal for shared cluster resources (which can move from machine to machine) will be located on the shared resource to ensure that it will be accessible to both machines.

By default, this setting applies to all journaled file systems but can be overridden by an override stanza for each journal file system. If the default value is a fully qualified path (for example `c:\tsmjournal`), all journal database files will be written to the specified directory. If the default value does not specify a drive letter, (for example `\tsmjournal`) the journal database files for each journal file system will be written to the specified directory on each journal file system.

The following is an example configuration stanza:

```
[JournalSettings]
;
; Store all resources in one location unless overridden
; by an override stanza
;
JournalDir=c:\tsmjournal
;
;
[JournaledFileSystemSettings.D:\]
;
; Journal for d: only will be in location specified below
;
JournalDir=d:\tsmjournal
```

**Note:** Changes to this setting do not take effect until the journaled file systems are brought online.

---

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

The following are examples of equivalent journal exclude statements:

```
dsm.opt: tsmjbbd.ini
```

```
EXCLUDE c:\testdir\...\* c:\testdir\  
EXCLUDE.DIR c:\testdir\test* c:\testdir\test*\*
```

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]  
%SystemRoot%\System32\Config\  
%SystemDrive%\Adsm.Sys\  
%TEMP%\*  
%TMP%\*  
c:\excludedir\  
c:\dir1\excluede\file  
*.*\*.tmp
```

**Note:** The `c:\excludedir\*` statement matches the entire tree including subdirectories and files.

---

## 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 608 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:

*DirNotifyBufferSize*

Specifies the size of the buffer to record change notifications for a particular journal file system. You might need to increase this value for

journalized file systems that generate a very large volume of change activity. The buffer size is limited by memory. The default value is 1 megabyte.

#### *JournalizedFileSystems*

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

```
JournalizedFileSystems=c: d:
```

#### *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*

Specifies the size of the memory buffer receiving file system change notifications for a particular journal file system. You might need to increase this value for journalized file systems that generate a very large volume of change activity. The buffer size is limited by memory. The default value is 2 megabytes. For example:

```
NotifyBufferSize=0x00200000
```

#### *NotifyFilter*

Specifies what file system change actions generate notifications to the journal service. *NotifyFilter* applies to file changes and directory modifications. Directory name changes, such as deletions and creations, are always tracked regardless of the filter value. Multiple actions can be monitored by combining (adding) values together. The default value is 0x11F (File and Dir Name, Attrib, Size, Last Write, and security Changes). You can also use the Tivoli Storage Manager Journal Engine Wizard to specify that any or all of these actions are monitored. Supported values are:

Value type	Decimal	Hex
File Name	1	0x001
Dir Name	2	0x002
Attribute	4	0x004
File size*	8	0x008
Last Write Time*	16	0x010
Last Access Time	32	0x020
Create Time	64	0x040
Security (ACL)	256	0x100

The asterisk (\*) indicates that notification might be deferred until disk write cache is flushed. Name changes are object creations, deletions, or renames.

Example:

```
NotifyFilter=0x107
```

#### *PreserveDbOnExit setting*

This setting allows a journal to remain valid when a journalized file system

goes offline and comes back online. This is useful for preserving the journal during system reboots, cluster failovers, and resource movement.

File systems go offline when the journal service stops or when the file system is removed from the configuration file. File systems come back online when the journal service is started or when the file system is added to the configuration file.

This setting allows a journal-based backup to continue processing when the service is restarted (or the file system comes back online) without performing a full incremental backup.

**Note:** Any change activity which occurs while the journal service is not running (or the file system is offline) will not be recorded in the journal.

In a clustered environment, shared resources can move to different machines in the cluster. The journal service running on each machine in the cluster must include these shared resources in the list of journaled file systems. The journal service running on the machine which currently owns the resource will actively be journaling the shared resource while other journal services on machines in the cluster which do not own the resource must defer journaling until the resource becomes available (or is moved to that machine). The configuration settings *deferFSMonStart*, *deferRetryInterval*, and *logFSErrors* allows deferment for a file system until the file system is available and accessible.

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:

```
[JournaledFileSystemSettings.D:\]
;
; Do not delete the journal when D:\ goes offline
;
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.

The *deferFSMonStart* setting is most commonly used in a cluster environment where shared resources might move to different machines in the cluster.

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*

This setting specifies whether errors encountered while accessing a journaled file system or journal directory are logged in the `jbberor.log` and the event log.

Use the *logFSErrors* setting with the *deferFSMonStart* setting to prevent excessive *File System unavailable* messages from being logged when bringing a journaled file system online is deferred. The first error which causes the file system to be deferred will be logged. Subsequent errors will not be logged. A value of 1 indicates that the setting is on. A value of 0 indicates that the setting is off.

An example to defer journaling until the file system journal directories are valid is:

```
[JournalSettings]
;
; Place journal files in directory on each journaled file system
;
journalDir=\tsmjournal

[JournaledFileSystemSettings]
;
;journal c:, d:, and f:
;
JournaledFileSystems=c: d: d:\mountpoint f:
;
; Override stanza to defer starting journaling for f:\
; until it is a valid file system

[JournalFileSystemSettings.f:\]
;
; Keep database valid if file system goes offline
;
PreserveDBOnExit=1
;
; Defer journaling until file system and journal directory
; are valid
;
deferFSMonStart=1
;
; Attempt to start journaling every 120 seconds when deferred
;
deferRetryInterval=120
;
; Do not log excessive resource unavailable messages
;
logFsErrors=0
```

---

## Overriding stanzas

Any setting in the *JournalFileSystemSettings* stanza, except for the buffer sizes, can be overridden for a particular journaled file system by creating an override stanza.

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

**Syntax for JournalFileSystemSettings stanza:**

*[JournalFileSystemSettings.fs]*

**Syntax for stanza settings:**

*JournalFileSystemSetting=override value*

Example:

```
[JournalFileSystemSettings.C:\]  
NotifyBuffer=0x0020000  
NotifyFilter=0x107
```

---

## Appendix D. Configuring the backup-archive client in a cluster server environment

You can install the Backup-Archive client software locally on each node of a Microsoft Cluster Server (MSCS) or Veritas Cluster Server (VCS) environment cluster. Tivoli Storage Manager in a VCS environment is supported on Windows 2003.

You can also install and configure the Scheduler Service for each cluster node to manage all local disks and each cluster group containing physical disk resources.

For example, MSCS cluster **mcs-cluster** contains two nodes: *node-1* and *node-2*, and two cluster groups containing physical disk resources: *group-a* and *group-b*. In this case, an instance of the Tivoli Storage Manager Backup-Archive Scheduler Service should be installed for *node-1*, *node-2*, *group-a*, and *group-b*. This ensures that proper resources are available to the Backup-Archive client when disks move (or fail) between cluster nodes.

The *clusternode* option ensures that Tivoli Storage Manager manages backup data logically, regardless of which cluster node backs up a cluster disk resource. Use this option for Tivoli Storage Manager nodes that process cluster disk resources, and not local resources. See “Clusternode” on page 214 for more information.

**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 318 for more information.

---

### Installing the backup-archive client on the cluster nodes

Install the Backup-Archive client software on a local disk on each cluster node. The executables should reside in the same location on each local drive, for example:

```
C:\Program Files\tivoli\tsm\baclient
```

---

### Configuring the backup-archive client to process local nodes

To back up the local (non-clustered) drives and process system state information, in a Microsoft Cluster Server (MSCS) or Veritas Cluster Server (VCS) environment, the Tivoli Storage Manager Client Scheduler Service should use the following combination of the options:

```
CLUSTERNODE      NO (default)
CLUSTERDISKONLY YES (default)
```

The *clusternode* and *clusterdiskonly* options should be invoked before MSCS or VCS is started, because although the Tivoli Storage Manager Client Scheduler Service is configured to back up local drives, the scheduler depends on the cluster

service. The Tivoli Storage Manager Scheduler communicates with the cluster service (MSCS or VCS) to build the cluster disks map, get the cluster name, etc.

To ensure that the MSCS or VCS loads before the Tivoli Storage Manager Client Scheduler service, a dependency should be added (for MSCS service or VCS service) to the Tivoli Storage Manager Client Scheduler service.

Follow these steps to add the service dependency:

1. Start `regedt32.exe`
2. Locate the following subkey in the registry: `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\TSM Scheduler Service`, (where "TSM Scheduler Service" is the name of the Tivoli Storage Manager Scheduler)
3. Double-click the **DependOnService** value
4. In the **Value data** box, type the `ClusSvc` for the MSCS environment or `HAD` for the VCS environment, and then click **OK**
5. Quit the Registry Editor
6. Restart your computer

**Note:** After disabling or uninstalling the MSCS or VCS, the dependency should be removed.

You can edit your `dsm.opt` file on each local node to process local disk drives using the following options:

***nodename***

If no value is specified, the backup-archive client uses the local machine name. See "Nodename" on page 312 for more information.

***domain***

If you do not specify a value for this option, the backup-archive client processes all local drives that are not owned by the cluster. See "Domain" on page 239 for more information.

***clusternode***

Do not specify this option when processing local drives. See "Clusternode" on page 214 for more information.

You can configure the Tivoli Storage Manager Backup-Archive Scheduler Service to back up the local cluster nodes.

---

## Configuring the backup-archive client to process cluster disk resources

Ensure that the backup-archive client manages each cluster group that contains physical disk resources as a unique node. This ensures that the backup-archive client correctly manages all disk resources, regardless of which cluster node owns the resource at the time of back up.

### Step 1: Identify the cluster groups to manage

Use the Microsoft Cluster Administrator utility or VCS Configuration editor to determine which groups contain physical disk resources for the backup-archive client to process. Register a unique node name on the backup server for each group. For example, MSCS cluster **mcs-cluster** contains the following groups and resources:

- **group-a** - Contains physical disk **q**: (quorum), and physical disk **r**:



**Note:** VCS does not have quorum disk.

- **group-b** - Contains physical disk **s:**, and physical disk **t:**

In this example, the administrator registers two node names: **mcs-cluster-group-a** and **mcs-cluster-group-b**. For example, to register **mcs-cluster-group-a** the administrator can enter the following command:

```
register node mcs-cluster-group-a <password>
```

## Step 2: Configure the client options file

Configure the client options file (`dsm.opt`) for each cluster group. Locate the option file on one of the disk drives that are owned by the cluster group. For example, the option file for **mcs-cluster-group-a** should reside on either **q:** or **r:**. To configure the `dsm.opt` file for each cluster group, specify the following options:

### *nodename*

Specify a unique name. For example:

```
mcs-cluster-group-a
```

See “Nodename” on page 312 for more information about this option.

### *domain*

Specify the drive letters for the drives which are managed by the group. For example:

```
q: r:
```

See “Domain” on page 239 for more information about this option. See “Frequently asked questions” on page 620 for information on how to add a cluster drive to an existing Tivoli Storage Manager cluster scheduler service resource for backup.

### *clusternode*

Specify the *Yes* value. See “Clusternode” on page 214 for more information about this option. If you set the *clusternode* option to *yes*, Tivoli Storage Manager:

1. Checks for a cluster environment (MSCS or VCS).
2. Uses the cluster name instead node name for file space naming and encryption. This allows the use of one password file for all nodes in the cluster.
3. Builds a list of shared volumes and works only with shared volumes. Back up of local volumes not permitted if the *clusternode* option set to *yes*.

**Note:** For the VCS, cluster database processing is skipped because VCS does not have a cluster database. VCS stores all cluster configuration information in an ASCII configuration file called `main.cf`, which is located in the path pointed by `%VCS_HOME%conf/config` on each node in the cluster. If this file is corrupted, the cluster configuration will also be corrupted. Caution is recommended when handling this file. The `VCS_HOME` environment variable points to the directory where VCS is installed on the node.

### *passwordaccess*

Specify the *generate* value. See “Passwordaccess” on page 320 for more information about this option.

### *managementservices*

(Optional). Specifies whether the Tivoli Storage Manager Client Acceptor

service (CAD) manages the scheduler, the Web client, or both. The examples in this appendix assume that the CAD will manage both the Web client and the scheduler for each cluster group. To specify that the CAD manages both the Web client and the scheduler, enter the following option in the `dsm.opt` file for each cluster group:

```
managedservices webclient schedule
```

See “Managedservices” on page 302 for more information about the *managedservices* option.

#### *errorlogname*

Specify a unique error log name. See “Errorlogname” on page 254 for more information about this option.

**Note:** This is not the same errorlog file that the client uses for other operations. Ideally, this file should be stored on a cluster resource, but at the very least it should be stored in a location other than the client directory.

#### *schedlogname*

Specify a unique schedule log name. See “Schedlogname” on page 352 for more information about this option.

**Note:** This is not the same schedlog file that the client uses for other operations. Ideally, this file should be stored on a cluster resource, but at the very least it should be stored in a location other than the client directory.

## Step 3: Configure the scheduler service

Configure a Tivoli Storage Manager Backup-Archive Scheduler Service for each cluster group using one of the following:

- Tivoli Storage Manager scheduler setup wizard, by choosing “This is a cluster node”
- Client Service Configuration Utility, `dsmcutil`

Configure the scheduler service on each node in the cluster group. The scheduler service must have the same name on each node and be available for failover (moved to the other nodes in the cluster).

To install the Tivoli Storage Manager Scheduler Service for *group-a* from machine *node-1*, ensure that *node-1* currently owns *group-a* and issue the following command:

```
dsmcutil install SCHEDuler /name:"tsm scheduler service: group-a"  
/clientdir:"c:\Program Files\tivoli\tsm\baclient" /optfile:q:\tsm\  
dsm.opt /node:mscs-cluster-group-a /password:nodepassword  
/validate:yes /autostart:no /startnow:no /clusternode:yes  
/clustername:mscs-cluster
```

This installs the service on *node-1*.

**Note:** For VCS, the default value of the `/autostart` option is *no*.

#### **Notes:**

1. For more information about `dsmcutil` commands and options, see “Using the `Dsmcutil` command” on page 589.

2. See “Frequently asked questions” on page 620 for information on how to add a cluster drive to an existing Tivoli Storage Manager cluster scheduler service resource for backup.
3. See “Frequently asked questions” on page 620 for information on how to verify that a scheduler service setup in a cluster environment will work.

Using the Microsoft Cluster Administrator utility or VCS Configuration Editor, move *group-a* to *node-2*. From *node-2*, issue the same `dsmcutil` command above to install the service on *node-2*. Repeat this procedure for each cluster group.

## Step 4: Creating a generic service resource for failover

Creating a generic service resource allows the Tivoli Storage Manager services for a cluster group to correctly failover between the two nodes of the cluster.

### MSCS

To add a Generic Service resource to each cluster group managed by the backup-archive client, use the Microsoft Cluster Administrator utility, as follows:

1. Select the *group-a* cluster group under the MSCS-Cluster\Groups folder and select **File** -> **New** -> **Resource** from the dropdown menu.
2. In the New Resource dialog, enter the following information:
  - Enter a unique name in the **Name** field. For example:
 

```
TSM SCHEDULE SERVICE for GROUP-A
```
  - Enter a description in the **Description** field. For example:
 

```
TSM cluster schedule for cluster group GROUP-A
```
  - Select the *Generic Service* resource type in the **Resource Type** field.
  - Enter the group name in the **Group** field. Press Enter.
3. In the Possible Owner dialog, ensure that all cluster nodes appear as possible owners. Press Enter to continue.
4. In the Dependencies dialog add all physical disk resources as Resource Dependencies. Press Enter to continue.
5. In the Generic Service Parameters dialog, enter the client acceptor service name that you specify in **Step 3** of “Configuring the Web client in a cluster environment” on page 615, in the **Service Name** field. Leave the **Startup Parameters** field blank. Press Enter.
6. In the Registry Replication dialog, add the registry key corresponding to the Tivoli Storage Manager node name and server name. The format for this key is:

```
HKEY_LOCAL_MACHINE\SOFTWARE\IBM\ADSM\CurrentVersion\Nodes\  
<nodename>\<TSM server instance name>
```

where *nodename* is the name of your Tivoli Storage Manager node, and *server instance name* is the name of the Tivoli Storage Manager server that the node connects to. For example, if the node name is **mcs-cluster-group-a** and the Tivoli Storage Manager server name is **tsmsv1**, then you should enter the following registry key in the Registry Replication dialog:

```
SOFTWARE\IBM\ADSM\CurrentVersion\BackupClient\Nodes\mcs-cluster-group-a\tsmsv1
```

**Note:** This entry should match an existing key in the Windows Registry.

7. Select the new resource from the Microsoft Cluster Administrator utility, and click **File** and then **Bring Online** from the dropdown menu.

Repeat this procedure for each cluster group managed by Tivoli Storage Manager.

The new scheduler service is now associated with the cluster group. If the group is moved (failed) to the other nodes in the cluster, the service should correctly fail over between the cluster nodes and notify both cluster nodes of automatic password changes.

**Note:** If you manually change the Tivoli Storage Manager password, you will need to stop the scheduler service, regenerate the password, and restart the scheduler service. The password can be generated by running the command:  
`dsmc query session -optfile="q:\tsm\dsm.opt"`

**Attention:** The generic service resource should remain online during any password changes. Password corruption will occur if the generic service resource is offline while the password is changed.

## VCS

VCS provides a separate utility to create or modify your existing configuration. This utility is called the VCS Configuration Editor. By default, this utility is installed when VCS Cluster Manager is installed.

**Note:** You can also manually edit the VCS configuration file `main.cf` file; refer to the appropriate VCS documentation for more information.

To add a Generic Service resource to each cluster group managed by the backup-archive client, use the VCS Configuration Editor, as follows:

1. Open the VCS Configuration Editor. You are prompted with the **Build a new configuration or modify existing configuration** window which provides the following options:

### New Config

If you select this option you are prompted for the path for the `types.cf` file.

### Open Existing Config

If you select this option, the configuration window opens. Click on **RESOURCE GROUP** you want to modify.

2. Click on the **Edit** button and select **Add resource**. The Add Resource window opens.
3. Enter the name you want to give the resource in **Resource Name** field.
4. Select the **Resource Type** as **GenericService**. The attributes of the **GenericService** resource type appear.
5. Click the **Edit** button to modify the resource attributes.
6. Select the *ServiceName* attribute and enter the name of scheduler or CAD service that you want to make High-Availability.
7. When you are finished, close the window. The Configuration window prompts you whether to save the configuration; click **Yes**.

**Use the VCS Configuration Editor to configure the registry replication resource, as follows:**

1. Open the VCS Configuration Editor. You are prompted with the **Build a new configuration or modify existing configuration** window which provides the following options:

### New Config

If you select this option you are prompted for the path for the `types.cf` file.

### Open Existing Config

If you select this option, the configuration window opens. Click on **RESOURCE GROUP** you want to modify.

2. Click on the **Edit** button and select **Add resource**. The Add Resource window opens.

3. Enter the name you want to give the resource in **Resource Name** field.
4. Select the **Resource Type** as **RegRep**. The attributes of the **RegRep** resource type appear.
5. Click the **Edit** button to modify the resource attributes.
6. Select the *MountResName* attribute and enter the shared disk on which you want to store the registry keys.
7. When you are finished, close the window. The Configuration window prompts you whether to save the configuration; click **Yes**.

Repeat this procedure for each cluster group managed by Tivoli Storage Manager.

The new scheduler service is now associated with the cluster group. If the group is moved (failed) to the other nodes in the cluster, the service should correctly fail over between the cluster nodes and notify both cluster nodes of automatic password changes.

**Note:** If you manually change the Tivoli Storage Manager password, you will need to stop the scheduler service, regenerate the password, and restart the scheduler service. The password can be generated by running the command:

```
dsmc query session -optfile="q:\tsm\dsm.opt"
```

---

## Configuring the Web client in a cluster environment

To use the Web client in a cluster environment, you must configure the backup-archive client GUI to run in a cluster environment. See “Installing the backup-archive client on the cluster nodes” on page 609 for detailed information about installing and configuring the backup-archive client in a MSCS or VCS environment.

### Configuring the Web client to process cluster disk resources

After installing and configuring the backup-archive client in a MSCS or VCS environment, perform the following steps.

#### Step 1: Identify the cluster groups to manage

Perform the steps under **Step 1** of “Configuring the backup-archive client to process cluster disk resources” on page 610.

#### Step 2: Configure the client options file

Perform the steps under **Step 2** of “Configuring the backup-archive client to process cluster disk resources” on page 610.

In addition, specify the following options in the *dsm.opt* file for each cluster group:

#### *httpport*

Specify a unique TCP/IP port number that the Web client uses to communicate with the client acceptor service associated with the cluster group. See “Httpport” on page 273 for more information about this option.

**Note:** It is preferable to specify a different log name using the *shedlogname* option in the *dsm.opt* file for each cluster group.

#### Step 3: Install a Client Acceptor Service and Client Agent

Install a unique client acceptor service and client agent for each cluster group and generate a password file.

To install the Client Acceptor Service for *group-a* from machine *node-1*, ensure that *node-1* currently owns *group-a* and issue the following command:

```
dsmcutil install cad /name:"tsm client acceptor: group-a"  
/clientdir:"c:\Program Files\tivoli\tsm\baclient" /optfile:  
q:\tsm\dsm.opt /node:mscs-cluster-group-a /password:nodepassword  
/validate:yes /autostart:yes /startnow:no httpport:1582 /cadschedname:  
"tsm scheduler service: group-a"
```

This will install the service on *node-1*.

To install the client agent service for *group-a* from machine *node-1*, ensure that *node-1* currently owns *group-a* and issue the following command:

```
dsmcutil install remoteagent /name:"tsm client agent: group-a"  
/clientdir:"c:\Program Files\tivoli\tsm\baclient" /optfile:  
q:\tsm\dsm.opt /node:mscs-cluster-group-a /password:nodepassword  
/validate:yes /startnow:no /partnername:"tsm client acceptor: group-a"
```

This will install the remote client agent service on *node1*.

#### Notes:

1. For more information about dsmcutil commands and options, see “Using the Dsmcutil command” on page 589.
2. Do not use the */autostart:yes* option.
3. Note that because the */clusternode* and */clustername* options are not allowed in this command at this level, it is possible that the password in the Windows Registry might need to be reset. After installing the above three services for each cluster group, generate a Tivoli Storage Manager password for each cluster group node name. You will need to identify the proper dsm.opt for each cluster group node name you authenticate. For example:

```
dsmc query session -optfile="q:\tsm\dsm.opt"
```
4. See “Frequently asked questions” on page 620 for information on what to do if a generic service resource for the cluster group fails because the CAD service has been removed.

Using the Microsoft Cluster Administrator utility or VCS Configuration Editor, move *group-a* to *node-2*. From *node-2*, issue the same commands above to install the services on *node-2* and generate a password file. Repeat this procedure for each cluster group.

## Step 4: Create a network name and IP address resource

Add a network name and IP address resource for each group that is managed by the client, using the Microsoft Cluster Administrator or VCS Configuration Editor.

**MSCS:** To add an IP Address resource to each cluster group managed by Tivoli Storage Manager, use the Microsoft Cluster Administrator utility as follows:

1. Select the *group-a* folder under the MSCS-Cluster\Groups folder and select **File** → **New** → **Resource** from the dropdown menu.
2. In the New Resource dialog, enter the following information:
  - Enter a unique name in the **Name** field. For example:  
IP address for GROUP-A
  - Enter a description in the **Description** field.
  - Change resource type to *IP address* in the **Resource Type** field.
  - Enter the group name in the **Group** field. Press Enter.
3. In the Possible Owner dialog, ensure that all cluster nodes appear as possible owners. Press Enter.

4. In the Dependencies dialog add all physical disk resources as Resource Dependencies. Press Enter.
5. In the TCP/IP Address dialog, enter appropriate values for address, subnetmask, and network. Press Enter.
6. Select the new resource from the Microsoft Cluster Administrator utility, and from the dropdown menu click **File** and then **Bring Online**.

**To add a network name resource to each cluster group managed by Tivoli Storage Manager, use the Microsoft Cluster Administrator utility as follows:**

1. Select the group-a folder under the MSCS-Cluster\Groups folder and select **File** -> **New** -> **Resource** from the dropdown menu.
2. In the New Resource dialog, enter the following information:
  - Enter a unique name in the **Name** field. For example:  
Network Name for GROUP-A
  - Enter a description in the **Description** field.
  - Change resource type to *Network Name* in the **Resource Type** field.
  - Enter the group name in the **Group** field. Press Enter.
3. In the Possible Owner dialog, ensure that all cluster nodes appear as possible owners. Press Enter.
4. In the Dependencies dialog add the IP address resource and all physical disk resources as Resource Dependencies. Press Enter.
5. In the Network Name Parameters dialog, enter a network name for GROUP-A. Press Enter.
6. Select the new resource from the Microsoft Cluster Administrator utility, and from the dropdown menu click **File** and then **Bring Online**.

The IP address and network name to backup the disks in the cluster group are now resources in the same group.

Repeat this procedure for each cluster group managed by Tivoli Storage Manager.

**VCS: Add a network name and IP address resource for each group that is managed by the client, using the VCS Configuration Editor, as follows:**

1. Open the VCS Configuration Editor. You are prompted with the **Build a new configuration or modify existing configuration** window which provides the following options:
  - New Config**  
If you select this option you are prompted for the path for the types.cf file.
  - Open Existing Config**  
If you select this option, the configuration window opens. Click on RESOURCE GROUP you want to modify.
2. Click on the **Edit** button and select **Add resource**. The Add Resource window opens.
3. Enter the name you want to give the resource in **Resource Name** field.
4. Select the **Resource Type** as **IP**. The attributes of the **IP** resource type appear.
5. Click the **Edit** button to modify the resource attributes.
  - a. Select the *MACAddress* attribute and enter the MAC address of adapter you want the IP to be assigned to.
  - b. Select the *SubNetMask* attribute and enter the subnetmask.

- c. Select the *Address* attribute and enter the IP address you want to make High-Availability.
6. When you are finished, close the window. The Configuration window prompts you whether to save the configuration; click **Yes**.

### Step 5: Creating a generic service resource for failover

**MSCS:** To add a Generic Service resource to each cluster group managed by Tivoli Storage Manager, use the Microsoft Cluster Administrator utility as follows:

1. Select the *group-a* folder under the MSCS-Cluster\Groups folder and select **File** -> **New** -> **Resource** from the dropdown menu.
2. In the New Resource dialog, enter the following information:
  - Enter a unique name in the **Name** field. For example:  
TSM CLIENT ACCEPTOR SERVICE for GROUP-A
  - Enter a description in the **Description** field.
  - Change resource type to *Generic Service* in the **Resource Type** field.
  - Enter the group name in the **Group** field. Press Enter.
3. In the Possible Owner dialog, ensure that all cluster nodes appear as possible owners. Press Enter.
4. In the Dependencies dialog add all physical disk resources as Resource Dependencies. Press Enter.
5. In the Generic Service Parameters dialog, enter the service name you specified with the **dsmcutil** command, in the **Service Name** field. Leave the **Startup Parameters** field blank. Press Enter.
6. In the Registry Replication dialog, add the registry key corresponding to the Tivoli Storage Manager node name and server name. The format for this key is:

```
HKEY_LOCAL_MACHINE\SOFTWARE\IBM\ADSM\CurrentVersion\Nodes\  
<nodename>\<TSM server instance name>
```

where *nodename* is the name of your Tivoli Storage Manager node, and *server instance name* is the name of the Tivoli Storage Manager server that the node connects to. For example, if the node name is **mcs-cluster-group-a** and the Tivoli Storage Manager server name is **tsmsv1**, then you should enter the following registry key in the Registry Replication dialog:

```
HKEY_LOCAL_MACHINE\SOFTWARE\IBM\ADSM\CurrentVersion\Nodes\  
mcs-cluster-group-a\tsmsv1
```

**Note:** This entry should match an existing key in the Windows Registry.

The client acceptor service is now a resource in the same group. If the group is moved (failed) to the other nodes in the cluster, the service should correctly fail over between the cluster nodes and notify both cluster nodes of automatic password changes.

#### Notes:

1. If you manually change the password, you will need to stop the remote agent and the client acceptor services, regenerate the password, and restart the client acceptor service (do not restart the remote agent). You can generate the password by running this command:  
dsmc query session -optfile="q:\tsm\dsm.opt"
2. See "Frequently asked questions" on page 620 for information on what to do if a generic service resource for the cluster group fails because the CAD service has been removed.



**VCS:** To add a Generic Service resource to each cluster group managed by the backup-archive client, use the VCS Configuration Editor, as follows:

1. Open the VCS Configuration Editor. You are prompted with the **Build a new configuration or modify existing configuration** window which provides the following options:

**New Config**

If you select this option you are prompted for the path for the types.cf file.

**Open Existing Config**

If you select this option, the configuration window opens. Click on RESOURCE GROUP you want to modify.

2. Click on the **Edit** button and select **Add resource**. The Add Resource window opens.
3. Enter the name you want to give the resource in **Resource Name** field.
4. Select the **Resource Type** as **GenericService**. The attributes of the **GenericService** resource type appear.
5. Click the **Edit** button to modify the resource attributes.
6. Select the *ServiceName* attribute and enter the name of scheduler service that you want to make High-Availability.
7. When you are finished, close the window. The Configuration window prompts you whether to save the configuration; click **Yes**.

**Use the VCS Configuration Editor to configure the registry replication resource, as follows:**

1. Open the VCS Configuration Editor. You are prompted with the **Build a new configuration or modify existing configuration** window which provides the following options:

**New Config**

If you select this option you are prompted for the path for the types.cf file.

**Open Existing Config**

If you select this option, the configuration window opens. Click on RESOURCE GROUP you want to modify.

2. Click on the **Edit** button and select **Add resource**. The Add Resource window opens.
3. Enter the name you want to give the resource in **Resource Name** field.
4. Select the **Resource Type** as **RegRep**. The attributes of the **RegRep** resource type appear.
5. Click the **Edit** button to modify the resource attributes.
6. Select the *MountResName* attribute and enter the shared disk on which you want to store the registry keys.
7. When you are finished, close the window. The Configuration window prompts you whether to save the configuration; click **Yes**.

The client acceptor service is now a resource in the same group. If the group is moved (failed) to the other nodes in the cluster, the service should correctly fail over between the cluster nodes and notify both cluster nodes of automatic password changes.

#### Notes:

1. If you manually change the password, you will need to stop the remote agent and the client acceptor services, regenerate the password, and restart the client acceptor service (do not restart the remote agent). You can generate the password by running this command:  

```
dsmc query session -optfile="q:\tsm\dsm.opt"
```
2. See "Frequently asked questions" for information on what to do if a generic service resource for the cluster group fails because the CAD service has been removed.

#### Step 6: Start the Web client

1. Start the Tivoli Storage Manager Client Acceptor Service for each resource group on each node.
2. To start the Web client, point your browser at the IP address and httpport specified for the Resource Group. For example, if you used an IP address of 9.110.158.205 and specified an httpport value of 1583, open the web address: `http://9.110.158.205:1583`.

Alternatively, you can point your browser at the network name and httpport. For example, if you used a network name of *cluster1groupa* and specified an httpport value of 1583, open the web address: `http://cluster1groupa:1583`.

Note that the Web client connects to whichever machine currently owns the resource group. The Web client displays all of the local file spaces on that machine, but to ensure that the files are backed up with the correct node name you should only back up the files for the resource group.

When failing back to the original node after a failover scenario, ensure that the remote agent service on the original machine is stopped. The remote agent can be stopped manually, or it will stop automatically after 20 to 25 minutes of inactivity. Because the remote agent is configured for manual startup, it will not start automatically if the machine on which it was running is rebooted.

---

## Frequently asked questions

### **Q: How do you configure a shortcut for the Tivoli Storage Manager backup-archive client GUI in a cluster environment?**

**A:** To configure a Tivoli Storage Manager backup-archive client GUI icon (for example on the Windows desktop) that you can use to manage operations for a cluster resource group on a Windows cluster, perform the following steps:

1. Right-click on the desktop and select **New** → **Shortcut**.
2. In the window that appears, find the path to the `dsm.exe` executable (located by default in directory `C:\program files\tivoli\tsm\baclient\`).

**Note:** If you type the path in, instead of using the **Browse** button, the path should be enclosed in double quotes. For example:

```
"C:\Program Files\tivoli\tsm\baclient\dsm.exe"
```

3. After you enter the path and executable in the text field, add the following information after the closing double quotes (add a space between the double quotes and the following):  

```
-optfile="x:\path\to\cluster\dsm.opt"
```

This identifies the proper Tivoli Storage Manager cluster options file you want to use. This example assumes that the cluster options file is located in the folder "x:\path\to\cluster\" and has a filename dsm.opt.

4. The complete line in the text field now should look similar to the following:  
`"C:\Program Files\tivoli\tsm\baclient\dsm.exe" -optfile="x:\path\to\cluster\dsm.opt"`
5. Click **Next** and give this shortcut a meaningful name, such as **Backup-Archive GUI: Cluster Group X**.
6. Click **Finish**. A desktop icon should now be available. The properties of this icon will show the following correct Target, as noted in step 4:  
`"C:\Program Files\tivoli\tsm\baclient\dsm.exe" -optfile="x:\path\to\cluster\dsm.opt"`

**Q: How do you verify that a scheduler service setup in a cluster environment will work?**

**A:** Setting up a Tivoli Storage Manager scheduler service for a Microsoft clustered resource group can be time consuming, and can be lengthened by mistakes and errors in the syntax of the commands used to set them up. Carefully entering the commands and recording important information about the cluster setup can minimize setup time. To successfully set up a scheduler service for Microsoft cluster environments:

1. Carefully read the information in this appendix for correct syntax on setting up a scheduler service for a cluster group.
2. Ensure that the proper dsm.opt file(s) are used for the cluster. In a typical normal workstation, only one dsm.opt file is used. In a clustered environment, additional dsm.opt files are required. Each cluster group that is backed up must have its own dsm.opt file. A cluster group is any group listed under the GROUPS folder of the cluster tree within the Microsoft Cluster Administrator utility or VCS Configuration Editor.
3. Understand what the following dsmsutil.exe options mean, and when to use them.

**/clustername:<clustername>**

Specifies the name of the Microsoft cluster, where <clustername> is the name at the top level of the tree within the Microsoft Cluster Administrator utility or VCS Configuration Editor. Use this option with dsmsutil.exe, only when installing a scheduler service for a cluster group.

**Recommendation:** Do not specify a clustername of more than 64 characters. If you specify more than 256 characters and you are using Veritas Storage Foundation with High Availability or a Microsoft Cluster Server configuration, you might not be able to install or start the Tivoli Storage Manager scheduler service.

**/clusternode:yes**

Specifies that you want to enable support for cluster resources. Use this option in the dsm.opt file for each cluster group, and with dsmsutil.exe when installing a scheduler service for a cluster group.

See "Using the Dsmsutil command" on page 589 for more information about these options.

4. Common mistakes are made in typing the syntax of the dsmsutil.exe command. An easy way to prevent such syntax problems is to create a temporary text file which is accessible to the cluster group (for instance, place it on a cluster drive

belonging to that cluster group), and type the syntax into this file. When needed, cut and paste this syntax from the file to the DOS prompt and press the enter key. It guarantees the consistency of the command syntax regardless of which end computer you enter it on.

5. If the scheduler service is failing to restart once a failover of the cluster group occurs (using the MOVE GROUP option in Cluster Administrator, for example), there could be potential password synchronization issues between the two cluster machines. An easy way to verify if the passwords are the same is to browse to this registry key on each machine and compare the encrypted password value:

```
HKEY_LOCAL_MACHINE\SOFTWARE\IBM\ADSM\CurrentVersion\BackupClient\Nodes\  
<nodename>\<servername>
```

If the encrypted keys for this node do not match between the two cluster machines, there is a password mismatch on one or both of the two end machines. To correct this problem, use the dsmc.exe program to update the password manually on both end machines. For example, assume that the Y: drive is part of the cluster group that is experiencing problems being backed up with a scheduler service. The Y:\tsm directory contains the dsm.opt file for this cluster group in the Y:\tsm directory. To update the password manually, enter the following command on both end machines:

```
dsmc -optfile=Y:\tsm\dsm.opt -clusternode=yes
```

and enter the following command to receive the prompt for the node name and password:

```
dsmc q se -optfile=Y:\tsm\dsm.opt -clusternode=yes
```

Afterwards, verify that the passwords are synchronized, and restart the scheduler service to verify if the password remains consistent. If password mismatching continues, it might be due to a syntax error in the original dsmcutil.exe command used to install the scheduler service. In this case, uninstall the scheduler service (using the dsmcutil remove /name:<schedule\_name> command) and reinstall the scheduler service again (using the shared text file syntax as shown above).

**Q: How do you add a cluster drive to an existing Tivoli Storage Manager cluster scheduler service resource for backup?**

**A:** To add an additional cluster drive resource to an existing Tivoli Storage Manager client cluster scheduler service, the following components must be modified or updated to properly reflect this change:

1. The cluster drive resource, and any related resource shares, must exist and reside within the designated cluster group as defined in the Microsoft Cluster Administrator utility or VCS Configuration Editor. The designated cluster group must already contain the cluster scheduler service resource for which this new drive will be added.
2. The dsm.opt file used by the designated cluster scheduler service resource must be modified to include the additional cluster drive resource on the *domain* option statement. For example, if you want to add the R:\ drive, and the *domain* statement currently identifies cluster drives Q: and S:, update the *domain* statement in your dsm.opt file as follows:

```
domain Q: S: R:
```

3. You must modify the cluster scheduler service resource properties to include this file in the list of dependent resources needed to bring this resource online.

This ensures that the cluster drive resource being added is included in the new backups, and for backups which run after a failover occurs.

After the changes above are made, bring the cluster scheduler service resource offline, and back online. The schedule should now process this additional resource for backups.

**Q: The CAD service has been removed and now the generic service resource for the cluster group is failing. How can this be corrected?**

**A:** The CAD can be used to control the scheduler, the Web client, or both for Tivoli Storage Manager in a cluster environment. If the CAD is removed without updating the generic cluster resource, the resource will fail. To correct this:

1. Verify which scheduler service was controlled by the CAD.
2. Using the Microsoft Cluster Administrator utility or VCS Configuration Editor, go to the properties window of the service resource, select the Parameters tab, and enter the name of the correct scheduler service to use.
3. Repeat steps one and two for each cluster group that was managed by the specific CAD.
4. To test the updated service resource, initiate a failure of the resource. If the resource comes back online with no failures, the update has worked properly.

**Note:** To fully disable the CAD, remove the *managedservices* option from the cluster group *dsm.opt* file or comment it out.



---

## Appendix E. ASR supplemental information

This appendix provides supplemental information for using Microsoft Windows Automated System Recovery (ASR) to recover Windows systems with the IBM Tivoli Storage Manager backup-archive client for Windows. The scope of this appendix does not include disaster recovery planning. See “Reference information” on page 633 for additional resources.

---

### ASR for Windows XP and Windows 2003

The following sections include information about recovering Windows XP and Windows 2003 Systems with the IBM Tivoli Storage Manager backup-archive client for Windows. See “Performing a Windows XP or Windows Server 2003 system recovery” on page 114 for complete instructions to perform this task. The functions described in that section, and the information in this appendix, is also documented in the Tivoli field guide, *Using Microsoft Windows Automated System Recovery (ASR) to Recover Windows XP and Windows 2003 Systems with the IBM Tivoli Storage Manager Backup-Archive Client for Windows*, available at the following Web site: [http://www.ibm.com/software/sysmgmt/products/support/Field\\_Guides.html](http://www.ibm.com/software/sysmgmt/products/support/Field_Guides.html)

#### ASR questions and answers

**Q: What if I did not create the ASR diskette prior to needing it for ASR?**

**A:** You can also create the ASR recovery diskette for another node if you are granted access to the ASR file space by the other node. See “ASR preparation procedure” on page 116 for information.

**Q: What do I do if the smss.ASR and ntdll.ASR files are not contained in my boot drive backup and I consequently encounter an error in the ASR recovery process?**

**A:** If you have another machine running the same maintenance level of Windows as the system you want to recover, you can copy the smss.ASR and ntdll.ASR files from the \windows\repair directory of that machine to a diskette. Re-run the ASR recovery process to the point where the Tivoli Storage Manager ASR recovery script prompts you for network or backup set recovery. Use the minimized Tivoli Storage Manager command window to copy the files to the machine being recovered. For example:

```
copy a:\*.asr c:\windows\repair
```

**Q: How do I create a local backup set containing the information required for ASR recovery?**

**A:** While you can use a single command to place all the required file spaces for recovery in a single file space, it is preferable for performance reasons to use separate backup sets. For Windows 2003, use one backup set for the system drive, and one backup set for system state. For Windows XP, use one backup set for the system drive and one backup set for system objects. Use the following Tivoli Storage Manager server commands to create backup sets on the server:

**Windows XP:**

**Syntax:**

```
generate backupset <nodename> <prefix> "system object" devclass=<file based device class> nametype=unicode
```

```
generate backupset <nodename> <prefix> "\\<machinename>\<systemdrive>" devclass=<file based device class> nametype=unicode
```

**Examples:**

```
generate backupset nodea nodeabackup "system object" devc=file nametype=unicode
```

```
generate backupset nodea nodeabackup "\\machinea\c$" devc=file nametype=unicode
```

**Windows Server 2003:****Syntax:**

```
generate backupset <nodename> <prefix> "system state" devclass=<file based device class> nametype=unicode
```

```
generate backupset <nodename> <prefix> "\\<machinename>\<systemdrive>" devclass=<file based device class> nametype=unicode
```

**Examples:**

```
generate backupset nodea nodeabackup "system state" devc=file nametype=unicode
```

```
generate backupset nodea nodeabackup "\\machinea\c$" devc=file nametype=unicode
```

To identify the files associated with the backup set, enter the following command:

```
q volhistory type=backupset
```

Copy the files to a CD using CD burning software. You can rename the files for convenience.

**Q: What files are placed on the ASR recovery diskette, and what do they do?**

**A:** The following files are placed on the ASR diskette:

**asr.sif** System information file. Provides disk layout information and a list of files to be copied to the system for Tivoli Storage Manager to perform ASR recovery.

**asrnp.sif**

Plug and Play device information used by Windows during an ASR operation.

**tsmasr.cmd**

Command file that installs the Tivoli Storage Manager client and invokes the commands to restore system state and system and boot drives.

**waitforevent.exe**

Program that detects the completion of the Tivoli Storage Manager client installation.

**tsmasr.opt**

A copy of the Tivoli Storage Manager client options file that matches the Tivoli Storage Manager environment in effect during the ASR backup.

**Q: After performing the ASR restore, the user account from which the backup was performed has missing shortcuts in the start menu and has lost desktop customizations. What causes this to happen?**



**A:** The likely cause is that the user profile of this account was in use during the backup or restore operation.

**Q: Why does the `dsmc restore` command seem to run more slowly in ASR mode than when I run it on a fully installed system? The output scrolls slowly on the console. How can I make restore run faster?**

**A:** In the GUI setup mode phase of ASR, the default video driver installed by Windows does not provide support for hardware acceleration. To improve performance, you can do the following:

1. Minimize the command window after you respond to the prompts for network or local backup set restore. If you are performing local backup set restore, you must maximize the command window to respond to prompts for additional restore volumes.
2. Maximize the diagnostic command prompt, and move that window to cover most of the restore command window. This reduces the effort required by the video subsystem to draw the changing text on the restore command window.
3. Use the *quiet* option in the `tsmasr.opt` options file to reduce the amount of feedback displayed on the restore console. You would edit `tsmasr.opt` on the ASR recovery diskette prior to running ASR. The disadvantage of using the *quiet* option is that you will not see any restore progress information except for final restore statistics.

**Q: How do you use ASR in conjunction with systems that are running in a cluster?**

**A:** Restoring clustered systems using ASR with Tivoli Storage Manager is similar to the methods you would use in a non-ASR environment. There are two scenarios: (1) restoring a cluster node and (2) restoring an entire cluster when no nodes are running in the cluster.

#### **Restoring a cluster node**

This is the most common restore scenario which you would use when an individual node of a cluster fails. The local node must be restored but the shared disks have failed over to another node. The cluster database is backed up by Tivoli Storage Manager as part of **systemstate**. The **restore systemstate clusterdb** command issued during a single system ASR insures the cluster database is restored correctly. No additional steps are required.

When the node reboots, the machine should join the cluster and begin to work. Since the shared disks have failed over to another node, they will appear to be inaccessible to the node.

#### **Restoring an entire cluster**

If there are no nodes running in the cluster, and an ASR is performed on the first node, the quorum information is placed on the node disk in the repair directory. The quorum information is not copied to the quorum disk during the ASR process. Use the `clustrest.exe` utility from the Windows Resource Kit to copy the quorum information to the quorum disk following the ASR restore. Reboot the machine and the single node should resume normal cluster operations. Follow the **Restoring a cluster node** procedure described above for remaining cluster nodes that need to be recovered.

## ASR diagnostic tips

1. The tsmasr.cmd file starts a secondary command window (minimized) for debugging. *Do not* close or interrupt the main tsmasr.cmd command window; when this command shell instance is closed, the ASR process exits and the system reboots. You can use this window to retry Tivoli Storage Manager commands issued from tsmasr.cmd if there is a failure, such as unavailability of the Tivoli Storage Manager server.
2. You can use tracing. Use TRACEFLAGS WIN2K and SYSTEMOBJECT. You can modify the tsmasr.opt client options file on the floppy disk using Notepad, to include entries for TRACEFILE and TRACEFLAGS. The TRACEFILE can be directed to the floppy if it is not too large, or to a file in a directory that is accessible from the recovery console, such as %SystemRoot%\temp or %SystemRoot%\temp\~tsmasr\baclient.
3. You can modify the tsmasr.cmd file on the floppy disk using Notepad to include additional diagnostic commands.
4. To improve restore performance, you can edit the tsmasr.opt client options file and insert the *quiet* option. This significantly reduces the feedback written to the console window during restore; only final statistics are displayed. However, this can inhibit diagnosis in the event of a restore failure.
5. If an ASR recovery failure occurs and you cannot boot the system, you can reboot into Windows recovery console and examine the following files in the %SystemRoot%\temp\~tsmasr\baclient directory:

### dsmerror.log

In this file, you might find entries such as: ReadPswdFromRegistry(): RegOpenPathEx(): Win32 RC=2. These are expected when Tivoli Storage Manager attempts to read the generated password from the Windows Registry, and the password is not present.

### tsmasrdiag.txt

The presence of this file indicates that ASR processing was completed. The file contains the output from the query backup %SystemRoot%\repair\\*.asr command. Ensure that the output of this command shows entries for the ntdll.ASR and smss.ASR files. These files are generated by the Tivoli Storage Manager **backup asr** command and should be included in the incremental backup of the system drive. If these files are not present after Tivoli Storage Manager restores the system drive, the ASR restore will fail. In this case, the asr.err file will contain error text such as the following:

```
Setup was unable to start the recovery application
"C:\windows\temp\tsmasr.cmd /sifpath=C:\windows\repair\asr.sif".
The error code returned was 0x2.
```

6. The setup log from the silent installation of the Tivoli Storage Manager client is found in %SystemRoot%\repair\tsmasrsetup.log file.
7. There is limited network functionality available in ASR mode. Simple FTP is available, and basic TCP/IP connectivity in a DHCP environment. NetBios services are not available, so you cannot access files on network shares.
8. There are two lines in the tsmasr.cmd file preceded by a REM Diagnostic line. The first line would be used to circumvent a problem with the silent install. The pause would stop ASR processing so that you could use the second command window to perform diagnostics. The second line directs the output of the query backup %SystemRoot%\repair\\*.asr command to a file on the floppy disk, rather than the default (preceding) command, which directs the output to a local file.

9. The waitforevent.exe command in the tsmasr.cmd file detects the completion of the asynchronous silent installation and pauses the execution of the tsmasr.cmd file until the install is complete, or 10 minutes (600) seconds have transpired. If the timeout occurs, the Timed out message displays in the command window. If the Timed out message appears, you should examine the tsmasrsetup.log file in the Windows repair directory for installation errors. The Timed out message can also be caused by installing a Tivoli Storage Manager client package prior to Version 5.2.0 in ASR mode. You must have a Version 5.2.0 or higher client to have the required support for ASR recovery.
10. When the ASR process encounters an error during GUI mode setup, the contents of the asr.err file appear on the screen in Notepad. Unfortunately, the error codes are not documented by Microsoft. However, you can determine what program was running:

**asr\_fmt.exe**

The ASR format utility for volumes.

**asr\_pfu.exe**

The ASR protected file utility.

**asr\_ldm.exe**

The ASR utility for Logical Disk Manager.

**tsmasr.cmd**

The Tivoli Storage Manager command file for ASR recovery.

## ASR in depth

### ASR.SIF

The asr.sif file is a Unicode file created by the Tivoli Storage Manager **backup asr** command. Some portions of the file are generated by the operating system and some are created by Tivoli Storage Manager. This section examines those portions of ASR.SIF, shown in bold type, that are created by Tivoli Storage Manager in the following example. Note the line numbers are for reference in this discussion and are not part of the actual file.

```

1. ;
2. ; Microsoft Windows Automated System Recovery State Information File
3. ;

4. [VERSION]
5. Signature="$Windows NT$"
6. ASR-Version="1.0"
7. Provider="IBM Tivoli Storage Manager"

8. [SYSTEMS]
9. 1="DILEXP", "x86", "5.1", "C:\WINDOWS", 1, 0x00010100, "480 0 -60 0-10-0-5 2:00:00.0 0-4-0-1 2:00:00.0", "Pacific Standard Time", "Pacific Standard Time"

10. [BUSES]
11. 1=1,1
12. 2=1,3

13. [DISKS.MBR]
14. 1=1,1,0,0x7a4c5b3b,512,63,255,8467200
15. 2=1,2,1,0x9d7c45ca,512,63,255,12594960

16. [DISKS.GPT]

17. [PARTITIONS.MBR]
18. 1=1,0,0,,0x00,0x05,0x05,16065,3887730,0x0
19. 2=1,1,0,"\\?\Volume{76ae2548-13a2-11d7-902a-
```

```

002035c2c2a2}" ,0x00,0x06,0x06,3903795,3887730,0x0
20. 3=1,4,0,"\\??\Volume{e93d26b4-132e-11d7-aa59-
806d6172696f}" ,0x00,0x06,0x06,16128,3887667,0x0
21. 4=2,0,3,"\\??\Volume{e93d26b2-132e-11d7-aa59-
806d6172696f}" ,0x80,0x07,0x07,63,8193087,0x1000
22. 5=2,1,0,,0x00,0x05,0x05,8193150,4385745,0x0
23. 6=2,4,0,"\\??\Volume{e93d26b3-132e-11d7-aa59-
806d6172696f}" ,0x00,0x0b,0x0b,8193213,4385682,0x0

24. [PARTITIONS.GPT]

25. [COMMANDS]
26. 1=1,3000,0,"%SystemRoot%\system32\asr_fmt.exe","/restore"
27. 2=1,4990,1,"%SystemRoot%\system32\asr_pfu.exe","/restore"
28. 3=1,2000,1,"%SystemRoot%\system32\asr_ldm.exe","/restore"

29. [ASRFMT.FIXEDVOLUMES]
30. 1=1,"\\??\Volume{e93d26b3-132e-11d7-aa59-
806d6172696f}" ,"\DosDevices\D:",FAT32,"PROGRAMS",0x4000
31. 2=1,"\\??\Volume{76ae2548-13a2-11d7-902a-
002035c2c2a2}" ,"\DosDevices\F:",FAT,"GHOST2",0x8000
32. 3=1,"\\??\Volume{e93d26b4-132e-11d7-aa59-
806d6172696f}" ,"\DosDevices\E:",FAT,"GHOST1",0x8000
33. 4=1,"\\??\Volume{e93d26b2-132e-11d7-aa59-
806d6172696f}" ,"\DosDevices\C:",NTFS,"",0x1000

34. [ASRFMT.REMOVABLEMEDIA]
35. 1=1,"\\Device\CdRom0" ,"\??\Volume{e93d26b1-132e-11d7-aa59-
806d6172696f}" ,"\DosDevices\R:"
36. 2=1,"\\Device\Floppy0" ,"\??\Volume{e93d26b0-132e-11d7-aa59-
806d6172696f}" ,"\DosDevices\A:"

37. [ASRLDM.VOLUMESTATE]
38. 0, 34, 1, 125,"LDM Configuration","Microsoft0"

39. [INSTALLFILES]
40. 1=1,"TSMASR" ,"\Device\Floppy0" ,"tsmasr.opt" ,"%SystemRoot%\temp\tsmasr.opt" ,"TSM
Options",0x00000001
41. 2=1,"TSMASR" ,"\Device\Floppy0" ,"tsmasr.cmd" ,"%SystemRoot%\temp\tsmasr.cmd" ,"TSM
installation batch file",0x00000000
42. 3=1,"TSMASR" ,"\Device\Floppy0" ,"waitforevent.exe" ,"%SystemRoot%\temp\waitforevent.
exe" ,"TSM utility file",0x00000000
43. 4=1,"TSMCLI" ,"\Device\CdRom0" ,"tsmcli.exe" ,"%SystemRoot%\temp\tsmcli.exe" ,"TSM
Setup Executable",0x00000001

44. [IBM.TIVOLISTORAGEMANAGER]
45. "IBM Tivoli Storage Manager"

```

**Line 7:**

Distinguishes this ASR file from one created by other backup products, such as NTBackup.

**Line 29:**

In the [COMMANDS] stanza, Tivoli Storage Manager tells ASR to execute tsmasr.cmd.

**Lines 41-44:**

In the [INSTALLFILES] stanza, Tivoli Storage Manager tells ASR the source and destination of files to be copied during text mode setup. The "TSMASR" and "TSMCLI" entries are the volume labels of the removable media for the source files. The \Device\ entries specify where ASR will look for the removable media (floppy or CD). The next two entries on each line specify the source and destination of the copy. Next is a comment

describing the entry. Finally, there is a flag value. 0x00000001 means that ASR will always prompt for the media to be mounted. 0x00000000 means *do not prompt*.

#### Lines 45-46:

This is a vendor-provided stanza reserved for Tivoli Storage Manager use. Currently, Tivoli Storage Manager does not exploit this stanza area.

#### TSMASR.CMD

The following example of the tsmasr.cmd file is from a Windows XP machine. The line numbers are for reference in this discussion and are not part of the actual file.

```
1. @start "IBM Tivoli Storage Manager" /MIN cmd.exe
2. @echo off
3. @title IBM Tivoli Storage Manager
4. %SystemRoot%\temp\tsmcli.exe -s -e setup.exe -a /s
   /v"INSTALLDIR="%SystemRoot%\temp\~\tsmasr\" ADDLOCAL="Client\"
   TRANSFORM="C:\tsm_images\TSM_BA_CLIENT\1033.mst\" /qn /l*v
   "C:\WINDOWS\repair\tsmasrsetup.Tog\" RebootYesNo="No\" REBOOT="Suppress\"
   ALLUSERS=1 "
5. %SystemRoot%\temp\waitforevent.exe 600
6. REM Diagnostics
7. REM pause
8. pushd %SystemRoot%\temp\~\tsmasr\baclient
9. copy %SystemRoot%\temp\tsmasr.opt >nul
10. @echo Do you wish to perform Automated Service Recovery (ASR) restore from the TSM
    server or from a local backup set (from file or tape)?
11. @echo 1. Restore from the TSM server.
12. @echo 2. Restore from a local backup set (from file).
13. @echo 3. Restore from a local backup set (from tape).
14. :choice
15. set /P CH=
16. if /I "%CH%"=="1" goto NETWORK
17. if /I "%CH%"=="2" goto BACKUPSET_FILE
18. if /I "%CH%"=="3" goto BACKUPSET_TAPE
19. goto choice
20. :NETWORK
21. dsmc restore \\DILEXP\C$ C:\ -subdir=yes -preservepath=complete -replace=all -
    tapeprompt=no -nodename=DILEXPNEW -optfile=tsmasr.opt
22. dsmc restore systemobject -asrmode=yes -nodename=DILEXPNEW -computername=DILEXP -
    optfile=tsmasr.opt
23. REM Diagnostics
24. dsmc query backup \\dilexp\c$\WINDOWS\repair\*.asr -asrmode=yes -
    nodename=DILEXPNEW -optfile=tsmasr.opt >tsmasrdiag.txt
25. REM dsmc query backup \\dilexp\c$\WINDOWS\repair\*.asr -asrmode=yes -
    nodename=DILEXPNEW -optfile=tsmasr.opt >a:\tsmasrdiag.txt
26. @GOTO DONE
27. :BACKUPSET_FILE
28. dsmc restore backupset TSMASR \\DILEXP\C$ C:\ -asrmode=yes -subdir=yes -
    preservepath=complete -replace=all -tapeprompt=no -nodename=DILEXPNEW -
    location=file -optfile=tsmasr.opt
29. dsmc restore backupset TSMASR systemobject -asrmode=yes -nodename=DILEXPNEW -
    location=file -replace=all -computername=DILEXP -optfile=tsmasr.opt
30. @GOTO DONE
31. :BACKUPSET_TAPE
32. dsmc restore backupset TSMASR \\DILEXP\C$ C:\ -asrmode=yes -subdir=yes -
    preservepath=complete -replace=all -tapeprompt=no -nodename=DILEXPNEW -
    location=tape -optfile=tsmasr.opt
33. dsmc restore backupset TSMASR systemobject -asrmode=yes -nodename=DILEXPNEW -
    location=tape -replace=all -computername=DILEXP -optfile=tsmasr.opt
34. :DONE
35. @echo on
```

**Line 1:**

Starts a minimized command prompt window which you can use for diagnostic purposes during ASR restore.

**Line 4-5:**

Performs silent install of the Tivoli Storage Manager client into the %SystemRoot%\temp\~tmasr directory. The installation progress is logged to the Windows repair directory in file tmasrsetup.log. The waitforevent command detects the completion of the asynchronous silent installation and pauses the execution of the tmasr.cmd file until the install is complete, or 10 minutes (600) seconds have transpired. If the timeout occurs then Timed out message is displayed in the command window.

**Line 9:**

The original client options file (stored as tmasr.opt) is copied to the baclient directory.

**Lines 21-22:**

These Tivoli Storage Manager commands are issued to restore the system drive and system state (Windows 2003); or system objects (Windows XP) from a network-connected Tivoli Storage Manager server. Using the *asrmode* and *computername* options in the **restore systemobject**, and **restore systemstate** commands allow system objects, or system state to be restored to a different machine. The name of the machine in ASR mode is *machinename*.

**Lines 24-25:**

These lines are for diagnostic use. The presence of the tmasrdiag.txt file in the baclient directory indicates that the ASR process completed. The tmasrdiag.txt file contains the output from query backup %SystemRoot%\repair\\*.asr command. Ensure that the output of this command shows entries for the ntdll.ASR and smss.ASR files.

**Lines 28-29 and 32-33:**

These Tivoli Storage Manager commands are issued to perform local backup set restore (from tape or file) of the system drive; system state (Windows 2003); or system objects (Windows XP). Note the special TSMASR backup set designator. This designator in combination with the *asrmode=yes* option, signal the restore process to prompt the user for the first volume of the backup set.

## **ASR information in the Windows Registry**

**The ASR commands and location are found in:**

HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Windows  
NT\CurrentVersion\Asr\Commands

**The files not to be backed up by ASR are found in:**

HKEY\_LOCAL\_MACHINE\SYSTEM\ControlSet001\Control\  
BackupRestore\FilesNotToBackup

**The registry keys not to be restored during ASR recovery are found in:**

HKEY\_LOCAL\_MACHINE\SYSTEM\ControlSet001\Control\  
BackupRestore\AsrKeysNotToRestore

---

## Reference information

The following Microsoft knowledge base articles are very instructive:

- 299044: How to Install Additional Files During Automated System Recovery
- 314058: Description of the Windows XP Recovery Console
- 314470: Definition of System Partition and Boot Partition
- 314686: ASR Cannot Restore System Disk to Replacement Disk of Identical Size
- 314688: A "Logical Disk Manager ASR Utility Error" Message When You Use ASR to Restore Disks That Are in a RAID Set
- 316484: "The Files for the Recovery Diskette Could Not Be Created" Error Message Occurs When You Use Ntbackup for Automated System Recovery

Details from Microsoft about creating slipstream Windows XP SP<sup>TM</sup> 1 installation are provided in the section, *To create an integrated installation of Windows XP and the service pack*, at: [http://www.microsoft.com/WindowsXP/pro/downloads/servicepacks/sp1/spdeploy.asp#the\\_integrated\\_installation\\_fm](http://www.microsoft.com/WindowsXP/pro/downloads/servicepacks/sp1/spdeploy.asp#the_integrated_installation_fm)

Additional information regarding the creation of a bootable CD containing a Windows XP SP1 installable operating system can be found at: <http://www.windows-help.net/WindowsXP/winxp-sp1-bootcd.html>

The following publication supplements the recovery information provided in **Disaster Recovery Strategies with Tivoli Storage Management - SG-24-6844**: <http://www.redbooks.ibm.com/abstracts/sg246844.html>





---

## Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

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

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106, Japan

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:**

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

IBM Corporation  
2Z4A/101  
11400 Burnet Road  
Austin, TX 78758 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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

If you are viewing this information in softcopy form, the photographs and color illustrations not appear.

---

## Trademarks

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX	Redbooks
C Set ++	RISC System/6000
DFS	S/390
FlashCopy	SP
HACMP	System/390
i5/OS	System i
IBM	System p
IBMLink	Tivoli
Lotus	Tivoli Enterprise
Multiprise	Tivoli Enterprise Console
OpenPower	TotalStorage
OS/390	VisualAge
OS/400	WebSphere
POWER	xSeries
Power PC	z/OS

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

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

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

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



---

## Glossary

The terms in this glossary are defined as they pertain to the IBM Tivoli Storage Manager library. If you do not find the term you need, refer to the IBM Software Glossary on the Web at this address: <http://www.ibm.com/ibm/terminology/>. You can also refer to IBM Dictionary of Computing, New York: McGraw-Hill, 1994.

This glossary include terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright (ANSI). You can purchase copies from the American National Standards Institute, 11 West 42nd Street, New York, New York 10036.
- The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC2/SC1).

### A

**absolute mode.** A backup copy group mode that specifies that a file is considered for incremental backup even if the file has not changed since the last backup. See also *mode*. Contrast with *modified mode*.

**access control list (ACL).**

1. In computer security, a collection of all access rights for one object.
2. In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights. For example, an access control list is a list that is associated with a file that identifies the users who can access the file and that identifies the users' access rights to that file.

**ACL.** See *access control list*.

**active policy set.** The activated policy set that contains the policy rules currently in use by all client nodes assigned to the policy domain. See also *policy domain* and *policy set*.

**active version.** The most recent backup copy of a file stored by Tivoli Storage Manager. The active version of a file cannot be deleted until a backup process detects that the user has either replaced the file with a newer version or has deleted the file from the workstation. Contrast with *inactive version*.

**adaptive subfile backup.** Permits user to perform a backup over a network device with limited bandwidth, such as a modem. Adaptive subfile backup reduces network traffic and increases the speed of your backup. An adaptive subfile backup sends only changed portions of a file to the server during successive backup operations instead of sending the entire file.

**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.

**Automated System Recovery (ASR).** A restore feature of Windows XP Professional and Windows Server 2003 operating systems that provides a framework for saving and recovering the Windows XP or Windows Server 2003 operating state in the event of a catastrophic system or hardware failure. Tivoli Storage Manager integrates with the ASR recovery process by providing ASR recovery diskette commands that allow restoration of a system from a Tivoli Storage Manager backup.

## B

**back up.** To copy information to another location to ensure against loss of data. In IBM Tivoli Storage Manager, you can back up user files, the IBM Tivoli Storage Manager database, and storage pools. Contrast with *restore*. See also *incremental backup* and *selective backup*.

**backup-archive client.** A program that runs on a file server, PC, or workstation and provides a means for users to back up, archive, restore, and retrieve files. Contrast with *administrative client*.

**backup copy group.** A policy object containing attributes controlling the generation, destination, and expiration of backup versions of files. The backup copy group belongs to a management class.

**backup retention grace period.** The number of days Tivoli Storage Manager retains a backup version when the server is unable to rebind the file to an appropriate management class.

**backup set.** A collection of active files in your file spaces that reside on the Tivoli Storage Manager server. The Tivoli Storage Manager administrator creates the backup set and copies it onto portable media device that is supported by the Tivoli Storage Manager server and client.

**backup version.** A file that a user backed up to server storage. More than one backup version can exist in server storage, but only one backup version is the active version. See also *active version* and *inactive version*.

**binding.** The process of associating a file with a management class name. See also *rebinding*.

**boot.** To prepare a computer system for operation by loading an operating system.

## C

**cache files.** During an online image backup, the Logical Volume Snapshot Agent (LVSA) creates a snapshot of a logical volume and saves blocks immediately before they are modified during the image backup. These blocks and their logical extents are saved in the cache files.

**CAD.** Client acceptor service

**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 service that serves the Web client Java applet to the Web browsers. The program that starts the client acceptor process is called *dsmcad*. On Windows, the client acceptor is installed and run as a service.

**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 file.** A file that a client can change, containing a set of processing options that identify the server, communication method, and options for backup, archive, hierarchical storage management, and scheduling. Also called the *dsm.opt* file.

**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).

**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 **backuppool** and **archivepool**.

**domain.** See *policy domain* or *client domain*.

**drag.** Move the mouse while holding down the mouse button, thus moving the selected object.

**drag-and-drop.** Move (drag) an object on top of another object and release the mouse button, thus relocating the object.

**dsm.opt file.** See *options file*. Also called client 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 bytes.

**exclude.** To identify files in an include-exclude list that you do not want to include in a specific client operation, such as backup or archive.

**expiration.** The process in which files are identified for deletion because their expiration date or retention period is passed. Backups or archives are marked for deletion based on the criteria defined in the backup or archive copy group.

**expiring file.** A migrated or premigrated file that is marked for expiration and removal from Tivoli Storage Manager storage. If a stub file or an original copy of a premigrated file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run. It expires and is removed from Tivoli Storage Manager storage after the number of days specified with the *migfileexpiration* option have elapsed.

## F

**file server.** A dedicated computer and its peripheral storage devices connected to a local area network that stores both programs and files shared by users on the network.

**file space.** A logical space on the Tivoli Storage Manager server that contains a group of files. In Tivoli Storage Manager, users can restore, retrieve, or delete file spaces from Tivoli Storage Manager storage. A file space for systems:

- **Windows**— file spaces for removable media are identified by volume label. Fixed drive file spaces are identified by Universal Naming Convention (UNC) name.
- **UNIX** — Logical space that contains a group of files backed up or archived from the same file system, or part of a file system defined with the *virtualmountpoint* option in the client system options file.

**frequency.** A copy group attribute that specifies the minimum interval, in days, between incremental backups.

**fuzzy backup.** A backup version of a file that might not accurately reflect what is currently in the file because the file was backed up at the same time as it was being modified.

**fuzzy copy.** An archive copy of a file that might not accurately reflect what is currently in the file because Tivoli Storage Manager archived the file while the file was being modified.

## G

**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<sup>9</sup>) bytes. (2) When referring to memory capacity, 1 073 741 824 in decimal notation.

**globally unique identifier (GUID).** A 16-byte code that identifies an interface to an object across all computers and networks. The identifier is unique because it contains a time stamp and a code based on the network address that is hard-wired on the host computer's LAN interface card.

**graphical user interface (GUI).** A graphical user interface offers pictorial rather than text-based access to a computer. A graphical user interface includes:



- A combination of graphics and icons
- Use of a mouse or pointing device
- Menu bars, dropdown lists, and overlapping windows

Contrast with *command-line interface*. See *windowed interface*.

**GUI.** Graphical user interface.

**GUID.** See *globally unique identifier*.

## H

**hierarchical storage management client.** A program that runs on a workstation or file server to provide space management services. The hierarchical storage management client automatically migrates eligible files to Tivoli Storage Manager storage to maintain specific levels of free space on local file systems. Automatic recalls are made for migrated files when they are accessed. Users are also permitted to migrate and recall specific files.

**hive.** On Microsoft Windows, a registry subkey that corresponds to a set of files in the `\system32\config` directory.

**HSM.** Hierarchical Storage Management.

## I

**inactive version.** A copy of a backup file in Tivoli Storage Manager storage that either is not the most recent version, or the corresponding original file was deleted from the client file system. Inactive backup versions are eligible for expiration according to the management class assigned to the file.

**include-exclude file.** A file containing statements to determine the files to back up and the associated management classes to use for backup or archive. See *include-exclude list*.

**include-exclude list.** A list of include and exclude options that include or exclude selected files for backup. An exclude option identifies files that should not be backed up. An include option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services. The include-exclude list is defined in one or more include-exclude files or in the client options file (`dsm.opt`) 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 *options file*.

**incremental backup.** A function that permits user to back up new or changed files or directories from a client domain or from specified directories or files. These directories or files are not excluded in the include-exclude list and meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class assigned to each file. Contrast with *selective backup*.

**Internet Information Services (IIS).** A structure that stores IIS configuration settings.

**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 engine service.** The journal engine service 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

**Logical Volume Snapshot Agent (LVSA).** Allows Tivoli Storage Manager to create a snapshot of a logical volume during an online image backup.

**LVSA.** Logical Volume Snapshot Agent.

## 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

**Nagle algorithm.** The Nagle algorithm, named after engineer John Nagle, was designed to reduce LAN and other network congestion from TCP applications. TCP implementations on UNIX and Linux began using the Nagle algorithm in the 1980s, and the algorithm remains a standard feature of TCP implementations today.

**Named Pipe.** A type of interprocess communication that permits message data streams to pass between peer processes, such as between a client and a server.

**NAS 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.

**non-Unicode file space.** Non-Unicode file spaces and file names are limited to the character set of the current locale when the files were backed up.

**NTFS.** A method for managing disk storage on the Windows operating system.

## O

**offline volume backup.** During an offline volume image backup, the volume is locked so that no other system applications can access it during the operation.

**online volume backup.** During an online volume image backup, the volume is available to other system applications during the operation.

**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 Tivoli Storage Manager servers, specifies communication methods, defines scheduling options, selects backup, archive, restore, and retrieve options. Also called the client 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.

**registry.** A central database in Windows that contains information about hardware, applications, and operating system settings for each system on the network. Provides security and control over system, security, and account settings.

**reparse points.** NTFS file system objects that have a definable attribute containing user-controlled data and are used to extend functionality in the input/output (I/O) subsystem.

**restore.** A function that permits users to copy a version of a backup file from the storage pool to a workstation or file server. The backup copy in the storage pool is not affected. Contrast with *backup*.

**retention.** The amount of time, in days, that inactive backed up or archived files are retained in the storage pool before they are deleted. The following copy group attributes define retention: retain extra versions, retain only version, retain version.

**retrieve.** A function permitting users to copy an archived file from the storage pool to the workstation or file server. The archive copy in the storage pool is not affected. Contrast with *archive*.

## S

**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.

**Secure Socket Layer (SSL).** A cryptographic protocol which provides secure client/server communications. SSL exploits public key (asymmetric) encryption and it provides a mechanism for certificate validation.

**select.** Choose an item from a list or group of items.

**selective backup.** A function permitting users to back up specified files. These files are not excluded in the include-exclude list and meet the requirement for serialization in the backup copy group of the management class assigned to each file. Contrast with *incremental backup*.

**serialization.** A copy group attribute that specifies whether a file can be modified during a backup or archive operation. See *static*, *dynamic*, *shared static*, and *shared dynamic*.

**server.** A program running on a mainframe, workstation, or file server that provides shared services such as backup and archive to other various (often remote) programs (called clients).

**server-prompted scheduling mode.** A client-server communication technique where the server contacts the client node when tasks need to be done. Contrast with *client-polling scheduling mode*.

**session.** A period of time in which a user can communicate with a server to perform backup, archive, restore, or retrieve requests.

**shared dynamic.** A Tivoli Storage Manager copy group serialization mode. This mode specifies if a file changes during backup or archive and continues to change after a number of retries. The last retry commits the file to the Tivoli Storage Manager server whether or not the file changed during backup or archive. Contrast with *dynamic*, *shared static*, and *static*.

**shared static.** A copy group serialization value specifying that a file must not be modified during a backup or archive operation. Tivoli Storage Manager attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See *serialization*. Contrast with *dynamic*, *shared dynamic*, and *static*.

**share point.** A drive or directory on Windows XP and .NET whose files are available for shared access across a network. The share point name is part of a UNC name. See *Universal Naming Convention (UNC)* name.

**shift-click.** Click an item while pressing the Shift key.

**space management.** The process of keeping sufficient free storage space available on a local file system for new data and making the most efficient and economical use of distributed storage resources.

**sparse file.** A file that is handled in a way that requires much less disk space than would otherwise be needed. Sparse support allows an application to create very large files without committing disk space for every byte. For example, you could use sparse support to work with a 42GB file in which you needed to write data only to the first 64KB.

**snapshot.** An image backup type. A snapshot is a point-in-time view of a volume. When you perform an online image backup, the Logical Volume Snapshot Agent (LVSA) takes a snapshot of the volume. Any changes that occur to the volume after the snapshot begins are not included in the backup.

**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.

**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.

**system drive or partition.** On Windows, the drive or partition on which the Windows operating system is installed.

**system services.** System services refers to the following system data on the Windows Server 2003 and Windows Vista operating systems:

- Background Intelligent Transfer Service
- Event logs
- Removable Storage Management Database (RSM)
- Cluster Database
- Remote Storage Service
- Terminal Server Licensing
- Windows Management Instrumentation (WMI)
- Internet Information Services (IIS) metabase
- DHCP database
- Wins database

The list of system services components is dynamic and change depending on service pack and operating system features installed. Tivoli Storage Manager allows for the dynamic discovery and back up of these components.

**system state.** System state refers to the following system data on the Windows Server 2003 and Windows Vista operating systems:

- Active Directory
- System Volume
- Certificate Server Database
- COM+ database
- Windows Registry
- System and boot files

The list of system state components is dynamic and change depending on service pack and operating system features installed. Tivoli Storage Manager allows for the dynamic discovery and back up of these components.

## T

**target node.** A client node for which other client nodes (called agent nodes) have been granted proxy authority. The proxy authority allows the agent nodes to perform operations such as backup and restore on behalf of the target node, which owns the data being operated on.

**TCP/IP.** Transmission Control Protocol/Internet Protocol.

**timeout.** A time event involving:

- An event that happens at the end of a predetermined period of time that began at the happening of another specified event.
- A time interval allotted for certain operations to happen. For example, response to polling or addressing before system operation is interrupted and must be restarted.
- A terminal feature that logs off a user if an entry is not made within a specified period of time.

**Tivoli Storage Manager.** A client-server licensed program product that provides storage management and data access services to customers in a multivendor computer environment.

**Transmission Control Protocol/Internet Protocol (TCP/IP).** A standard set of communication protocols that supports peer-to-peer connectivity of functions for both local and wide-area networks.

## U

**Unicode.** Unicode is a universal character encoding standard that supports the interchange, processing, and display of text that is written in any of the languages of the modern world.

**Unicode-enabled file space.** Unicode file space names provide support for multi-lingual workstations without regard for the current locale.

**Universal Naming Convention (UNC) name.** A name used on Windows to access a drive or directory containing files shared across a network. The UNC name includes the machine name and a share point name that represents the shared drive or directory. See *share point*.

## 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.

**VMware Consolidated Backup.** A VMware backup solution for the ESX server in a SAN environment. Consolidated backup allows the backup of multiple virtual machines to be offloaded to a dedicated physical host. By off-loading the backup away from virtual machines and the ESX service console, it also allows backup in a LAN-free environment.

**Volume Shadowcopy Service (VSS).** An integral part of the Windows Server 2003 and Windows Vista operating systems that allows you to create shadow copy backups of volumes, exact copies of files, including all open files. Shadow copy backups ensure the following:

- Applications can continue to write data to the volume during a backup.
- Files that are open are no longer omitted during a backup.
- Backups can be performed at any time, without locking out users.

## 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*.

**WMI.** Windows Management Instrumentation repository

**workstation.** A programmable high-level workstation (usually on a network) with its own processing hardware such as a high-performance personal computer. In a local area network, a personal computer that acts as a single user or client. A workstation can also be used as a server.

**world wide name.** A unique 48 or 64 bit number assigned by a recognized naming authority (often with block assignment to a manufacturer) that identifies a connection or a set of connections to the network. Abbreviated WWN. A WWN is assigned for the life of a connection (device). Most networking technologies (e.g., Ethernet, FDDI, etc.) use a world wide name convention.





---

# Index

## Numerics

128-bit AES encryption support 63

## A

absolute mode 158

accessibility

  disability xv

  keyboard xv

  shortcut keys xv

activatekey option 192

active backup versions

  displaying 54, 106, 491

  restoring 106

Active Directory

  query for communication method and server with which to connect 405

adaptive subfile backup

  excluding files from 257

  including files for 280

  restore permissions 47

  specifying 383

  specifying client cache path 384

  specifying client cache size 385

administrative client

  allowing secure sessions within a private network 388

adsm.sys directory, adsm98.sys directory 542

archive

  a list of files 140

  assign description to 137, 140, 230

  associating local snapshot with server file space 139, 375, 434

  binding management class to 140, 161, 193

  case sensitive file name conflicts 205

  command 433

  compressing files 222

  configuring client-node proxy support 140

  copy group 154

  copy group attributes 155

  copy mode 158

  delete files after 229

  delete individual archives from server file space 142, 463

  deleting file spaces 135

  directory tree 138

  display details 232

  display the last modification date and creation date 488

  enabling cluster drive management 214

  encrypting data during 63

  encryptiontype 250

  exclude objects from 257

  files only 266, 406

  grace period retention 154

  how managed 153

  improving speed using share memory 167

  include files for 280

  information, query 488

  list of files 262

  managing file growth during compression 221

  maximum file size 63

  more than one file specification 140

archive (*continued*)

  only files; not directories 140

  overriding management class during 160

  packages 137

  primary tasks 137

  process Dfs junction as junction or directory 233

  process directories only (not files) 235

  purpose 58

  query user access 487

  regulate resources server and client use during processing 344

  retrieving using command line 143

  revoke access 462

  shared data on multiple clients under a single node name 140, 194

  specify end date and time to search for 401, 402

  specifying number of attempts to archive open files 207

  specifying whether to include subdirectories for 140, 381

  specifying whether to update last access date 331

  specifying whether to wait for tape to mount if it is required 386

  starting a Web client session 51

  summary of options 168

  suppress confirmation prompt before deleting 315

  testing include-exclude list 40

  using commands 139

archmc option 193

asnodename option 194

ASR

  preparation overview 115

  preparation procedure 116

  query 490

  restore procedure 118

  software requirements 115

asrmode option 196

Attention

  changing HFS file paths to UNIX file paths 4

  DSM\_LOG environment variable 352, 558

  generic service resource, and password changes 614

  incremental-by-date backup dates 286

  insufficient system resources on Linux 217

  running during install 9

  securing local backup sets on portable media 123

  sessioninitiation option and serveronly 361

  system hangs when using the backup registry

    command 450

  system hangs when using the backup systemobject

    command 452

  Using the asterisk (\*) and not the question mark (?) for multibyte code pages 431

auditlogging option 197

auditlogname option 200

authentication

  Tivoli Storage Manager client 47

authorization

  authorizing user to restore or retrieve your files 132

  summary of options 181

authorized user

  definition xi

autofsrename option 202

- Automated System Recovery (ASR)
  - concepts 625, 629
  - diagnostic tips 625, 628
  - questions and answers 625
  - reference information 633
  - restoring files using GUI 113
  - specifying computer name for which system state recovery is performed 224
- Automated System Recovery (ASR) files
  - back up 82, 436
  - restore 528
- automating backup services
  - displaying scheduled work 145, 147
  - establish remote connection for 151
  - mobile dial-up support 151
  - options for 145
  - overview 145, 571
  - process commands after back up 325
  - process commands before back up 329
  - starting client scheduler 23

## B

- backup
  - adaptive subfile
    - specifying 383
    - specifying client cache path 384
    - specifying client cache size 385
  - Automated System Recovery (ASR) files 82, 436
  - automating using Client Service Configuration Utility 585
  - binding management class to files 161
  - bypass processing of NTFS security information 365
  - case sensitive file name conflicts 205
  - certificate server database system object 437
  - COM+ database system object 438
  - comparison: incremental, journal-based, incremental-by-date 68
  - compressing files 222
  - configuring support for client node proxy backups 76
  - copy group 66, 154
  - copy mode 158
  - delete individual backups from server file space 75
  - delete individual files or images from server file space 465
  - deleting file spaces 135
  - display details 232
  - displaying processing status 95
  - enable LAN-free data movement 249
  - enabling cluster drive management 214
  - enabling communications for LAN-free data movement 291, 294
  - encrypting data during 63
  - encryptiontype 250
  - estimating processing time 72
  - event log system object 439
  - exclude objects from 257
  - excluding domains 73, 74, 239
  - excluding system objects 36, 80, 258
  - excluding system services components 81
  - excluding system services from 81
  - excluding system state object from 81
  - files only 266
  - files, assigning management class 159
  - fixed drives 99
  - fuzzy 158
  - grace period retention 154

- backup (*continued*)
  - group
    - delete 470
    - specify name of 271
    - specify virtual file space name 409
  - group, display all members of 364
  - how Tivoli Storage Manager manages 62
  - image 82
    - client domain 241
    - configuring online image backup 28
    - revoke access 462
    - using command line 87
    - using the GUI 86
    - using with incremental backup 84, 443
    - using with incremental mode 443
  - image, considerations 83
  - image, offline and online 83
  - improving speed using share memory 167
  - inactivate a list of 472
  - include files for 280
  - include-exclude list 62
  - include-exclude processing for 32
  - incremental
    - associating local snapshot with server file space 375, 479
    - by-date with image backup 85
    - client domain 239
    - directories, processing overview 66
    - new and changed files with modification date later than last backup 286
    - process a list of files 262
    - using command line 73
  - incremental and selective commands with snapshotroot option 79
  - incremental-by-date 68
    - directories, processing overview 68
    - using command line 74
  - incremental-by-date, using client GUI 71
  - incremental-by-date, using the client command line 73
  - incremental, specify memory-conserving algorithm for 305
  - incremental, using client GUI 71
  - incremental, using the client command line 73
  - journal-based 66, 476
  - LAN-based image 442
  - managing 153
  - managing file growth during compression 221
  - maximum file size 63
  - Microsoft Dfs trees and files 101
  - multi-session, send files contiguously to the server 215
  - NAS file systems 87
  - NAS file systems with command line 90
  - NAS file systems with web client GUI 88
  - NAS image
    - client domain 242
  - NetApp CIFS share definitions 95
  - Network Attached Storage (NAS) file systems 447
  - new or changed files 65
  - NTFS file spaces 99
  - online and offline image backup 442
  - online image
    - number of times LVSA attempts snapshot file system idle wait time before operation fails 369
    - specify amount of time before a snapshot is taken 371
    - specifying gap size of striped volumes 275
    - specifying location of the cache files 367
    - specifying the size of the cache files 368

- backup (*continued*)
  - open file support
    - configuring 29
    - specifying location of the cache files 367
    - specifying the size of the cache files 368
  - open files 97
  - overview 57, 64
  - performing traditional full incremental 66, 314, 476
  - Persistent Storage Manager 88
  - planning 57
  - pre-backup considerations 59
  - primary tasks 57
  - process Dfs junction as junction or directory 233
  - process directories only (not files) 235
  - purpose 58
  - query information 491
  - query user access 487
  - registry system object 450
  - regulate resources server and client use during
    - processing 344
  - removable media using drive label 99
  - reset Windows archive attribute 342
  - restore 105
  - revoke access 462
  - saving encryption key password 251
  - selective
    - associating local snapshot with server file space 375, 560
    - back up list of files 74
    - overview 69
    - using command line 74
  - selective, using client GUI 71
  - selective, using the client command line 73
  - shared data on multiple clients under a single node name 194
  - shared files and directories 100
  - snapshotroot option with incremental and selective commands 79
  - specify end date and time to search for 401, 402
  - specifying computer name for which system state recovery is performed in ASR mode 224
  - specifying number of attempts to backup open files 207
  - specifying one server session per file specification 215
  - specifying shared memory port for LAN-free data movement 293
  - specifying whether to compute CRC for comparison of NTFS security information 366
  - specifying whether to include subdirectories for 381
  - specifying whether to update last access date 331
  - specifying whether to wait for tape to mount if it is required 386
  - starting a Web client session 51
  - subdirectories 74
  - summary of options 168
  - system and boot files system object 451
  - system objects 79, 452
  - system services, assigning management class 81
  - system state 453
  - system state, assigning management class 37, 81
  - testing include-exclude list 40
  - using LAN-free data movement 59
  - using multiple sessions 95
  - using open file support 69
  - VMware Consolidated Backup 92
  - Websphere Application Server
    - Application Server 419, 420, 455
    - assigning management class 455

- backup (*continued*)
  - Websphere Application Server (*continued*)
    - Network Deployment Manager 419, 420, 455
    - set WebSphere node name 419
    - set WebSphere node name, user name, password 569
    - set WebSphere user name 421
  - WebSphere Application Server
    - Application Server 95
    - Network Deployment Manager 95
  - Windows Management Instrumentation repository 459
  - Windows Server 2003 system state 80
  - Windows Vista system state 80
    - with limited bandwidth 61
- backup and restore
  - NAS file servers using CIFS 91
- backup asr command 436
- backup certserverdb command 437
- backup complusdb command 438
- backup copy group
  - attributes 155
- backup eventlog command 439
- backup group command 440
- backup image command 442
- backup nas command 447
- backup registry command 450
- backup set
  - enabling GUI for local restore of 299
  - restore 105
  - restore systemobjects from 529
  - restore systemstate from 529
  - restoring in a SAN environment 529
  - specify location for restore or query 300
- backup sysfiles command 451
- backup systemobject command 452
- backup systemstate command 453
- backup VM command 457
- backup was command 455
- backup with client node proxy
  - agent node 77
  - target node 77
- Backup with client node proxy
  - overview 76
- backup wmi command 459
- backup-archive client
  - overview 1
- backup-archive scheduler service
  - install 585
- backupset
  - enabling GUI for local restore 122
  - restore 122
- backupsetname option 204
- batch mode 427
  - starting a session 49
- bottom-up processing
  - include-exclude list 41

## C

- cache files
  - specifying location of 367
  - specifying size of 368
- cancel process command 460
- cancel restore command 461
- casesensitiveaware option 205
- central scheduling
  - summary of options 179

- certificate server database
  - system object, back up 437
- changingretries option 207
- class option 208
- classic (standard) restore 109
- client
  - registering with server 30
  - setting password 30
  - specify size for the TCP/IP sliding window for client node 396
  - specifying client TCP/IP address other than the one for first server contact 391
  - specifying client TCP/IP port number other than the one for first server contact 392
- client acceptor daemon
  - manage scheduler, Web client, or both 302
- Client Acceptor service
  - configuring to manage scheduler 21
- client command options
  - overview 428
- client components 5
- client node proxy archive
  - overview 140
- client node proxy backup
  - scheduling 77, 78
- client options
  - display current settings 19, 508
  - exclude
    - exclude.archive 32
    - exclude.backup 32
    - exclude.compression 32
    - exclude.dir 32
    - exclude.file 32
    - exclude.file.backup 32
    - exclude.image 32
    - exclude.subfile 32
    - exclude.systemobject 32
  - guitreeviewafterbackup 272
  - order of processing (precedence) 183
  - overriding using command line 183
  - overview 428
  - setting in a file 18
  - using with commands 183
- client options file
  - creating and modifying 15
  - generating in shared directory 17
  - overriding using commands 183
  - required options for 15
  - specifying include-exclude options 31
- client options reference 190
- client scheduler
  - displaying scheduled work 145, 147
  - establish remote connection for 151
  - mobile dial-up support 151
  - options for 145
  - overview 145, 571
  - starting 23, 557
  - starting automatically 53
- Client Service Configuration Utility
  - commands for installing client services 590
  - configure CAD to manage existing scheduler service 587
  - create new scheduler and associate CAD to manage scheduler 587
  - options for installing client services 596
  - using to automate backups 585
- client services
  - considerations 45
- client user options
  - creating multiple files 17
- client-node proxy support 76, 140
- client-server communication
  - establishing 15
  - identify your workstation to the server 312
  - query Active Directory for communication method and server with which to connect 405
  - specify maximum disk I/O buffer size client uses when reading or writing files 237
  - specify size for the TCP/IP sliding window for client node 396
  - specifying client TCP/IP address other than the one for first server contact 391
  - specifying client TCP/IP port number other than the one for first server contact 392
  - specifying method 217
  - specifying number of kilobytes client buffers before sending transaction to server 403
  - specifying reconnection attempts after failure 219
  - specifying reconnection interval after failure 220
  - specifying size of internal TCP/IP communication buffer 389
  - specifying TCP/IP address for dsmcad 390
  - specifying TCP/IP address of Tivoli Storage Manager server 395
  - specifying TCP/IP port address of Tivoli Storage Manager server 394
  - specifying TCP/IP port address on which to establish shared memory connection 363
  - specifying the name of a named pipe 310
  - specifying whether to send small transactions to server without buffering them first 393
- clientview option 209
- closed registration
  - permissions 30
  - using 30
- cluster drives
  - enabling management of 214
- cluster resources
  - permissions 47
- clusterdiskonly command 210
- clusternode option 214
- collocatebyfilespec option 215
- COM+ database
  - system object, back up 438
  - system object, query 497
  - system object, restore 532
- command line
  - assigning description to archive 140
  - display current settings for client options 19, 508
  - displaying processing status 95
  - displaying the Euro characters in prompt 50
  - enabling 8.3 short names 245
  - enabling command history capabilities 244
  - ending a session 55
  - entering commands 428
  - general rules when entering options with commands 184
  - NAS file systems backup 90
  - options 184
  - overriding management class during archive 160
  - overview of parameters 429
  - performing image backup 87
  - performing point-in-time restore 126
  - restoring files and directories 107
  - restrictions for NAS file systems 87
  - return codes for operations 151

- command line (*continued*)
  - rules for entering commands 431
  - specifying file specification 429
  - specifying options file during session 318
  - starting a session 49
  - using wildcard characters 431
- command parameters
  - overview 429
- command processing, summary of options 181
- command session
  - ending 427
  - starting 427
- command-line prompt
  - displaying the Euro characters 50
- commands
  - archive 433
  - backup asr 436
  - backup certserverdb 437
  - backup complusdb 438
  - backup eventlog 439
  - backup group 440
  - backup image 442
  - backup nas 447
  - backup registry 450
  - backup sysfiles 451
  - backup systemobject 452
  - backup systemstate 453
  - backup VM 457
  - backup was 455
  - backup wmi 459
  - batch mode 427
  - cancel process 460
  - cancel restore 461
  - clusterdisksonly 210
  - delete access 462
  - delete archive 463
  - delete backup 465
  - delete filespace 468
  - delete group 470
  - entering 428
  - entering on command line 428
  - expire 472
  - general rules when entering options with 184
  - help 474
  - incremental 475
  - interactive (loop) mode 428
  - loop 482
  - macro 484
  - maximum file specifications permitted 430
  - monitor process 485
  - overview of parameters 429
  - preview 486
  - query access 487
  - query archive 488
  - query ASR 490
  - query backup 491
  - query backupset 495
  - query complusdb 497
  - query eventlog 498
  - query filespace 499
  - query group 501
  - query image 503
  - query inclexcl 505
  - query mgmtclass 506
  - query node 507
  - query options 508
  - query registry 509
- commands (*continued*)
  - query restore 510
  - query schedule 511
  - query session 512
  - query sysfiles 513
  - query systeminfo 514
  - query systemobject 516
  - query systemstate 517
  - query VM 519
  - query was 518
  - query wmi 520
  - restart restore 521
  - restore 522
  - restore asr 528
  - restore backupset 529
  - restore complusdb 532
  - restore eventlog 533
  - restore group 534
  - restore image 536
  - restore NAS 540
  - restore registry 542
  - restore sysfiles 544
  - restore systemobject 545
  - restore systemstate 547
  - restore was 549
  - restore wmi 552
  - retrieve 553
  - rules for entering 431
  - schedule 557
  - scheduled, enabling or disabling 150
  - selective 559
  - set access 563
  - set event 566
  - set password 568
  - set waspassword 569
  - specifying file specification 429
  - using 423
  - using in executables 151
  - using options with 183
  - using wildcard characters 431
- commmethod option 217
- commrestartduration option 219
- commrestartinterval option 220
- communication methods 7
  - summary 166
- communications
  - establishing through firewall 23
  - establishing with Secure Socket Layer (SSL) 26
- compressalways option 221
- compression
  - disabling processing 260, 283
  - enabling processing 260, 283
  - include-exclude statements 260, 283
- compression option 222
- compression processing
  - exclude from backup 260, 283
  - exclude options 260, 283
  - excluding files from 257
  - include files for 280
- compression, encryption, and subfile backup processing
  - back up 283
  - exclude from backup 283
  - exclude options 283
- computername option 224
- configuring
  - journal engine service 28
  - optional tasks 15

- configuring *(continued)*
  - required tasks 15
  - the client scheduler 21
  - the Web client 20
- console option 225
- console window
  - displaying the Euro characters 50
- copy destination attribute 158
- copy frequency attribute 156
- copy group name attribute 156
- copy groups 153
  - archive 154
  - backup 154
- copy mode attribute
  - absolute 158
  - modified 158
- copy serialization attribute 157
- copy type attribute 156
- customer support
  - contacting xv

## D

- date format
  - specifying 227
- dateformat option 227
- default domain
  - excluding domains from backup 73, 74, 239
- default management class 153
- default policy domain 153
- delete
  - file space 135
    - specifying NAS or client objects 208
- delete access command 462
- delete archive command 463
- delete backup command 465
- delete group command 470
- deleted file systems 98
- deletefiles option 229
- deleting
  - authorizations 132
    - individual archives from server file space 142, 463
    - individual backups from server file space 75, 465
- description option 230
- detail option 232
- dfsbackupmntpnt option 233
- diagnostics
  - options 183
- directories
  - adsm.sys 542
  - adsm.sys\registry.sav 542
  - adsm98.sys 542
  - adsm98.sys\registry.sav 542
  - assigning management class for 234
  - excluding 33
    - excluding from backup processing 257
    - incremental backup processing overview 66
    - processing during incremental-by-date 68
    - restoring from command line 107
    - restoring from GUI 107
    - specifying on command line 429
- directory
  - archiving 138
- dirmc option 234
- dirsonly option 235
- disability
  - accessibility xv

- disability *(continued)*
  - keyboard xv
  - shortcut keys xv
- disablenqr option 236
- Discretionary access control list (permissions)
  - back up 99
- disk recovery 134
- disk space requirements 5
- diskbuffsize option 237
- diskcachelocation option 238
- displaying
  - archive information 488
  - online help 55
  - restartable restore sessions 510
  - scheduled events 511
  - session information 512
- domain
  - back up using the GUI 73
    - include for image backup 241
    - include for incremental backup 239
    - include for NAS image backup 242
    - specifying drives in the default 73
- domain list
  - using universal naming convention to specify 100
- domain option 239
- domain.image option 241
- domain.nas option 242
- drive label
  - using to backup removable media 99
- DSM\_CONFIG environment variable 19
- DSM\_DIR environment variable 19
- DSM\_LOG environment variable 19
- dsm.opt
  - setting options in 18
- dsm.opt file
  - creating and modifying 15
    - required options for 15
    - specifying a drive specification using wildcards 38
- dsm.smp file
  - copying to dsm.opt 15
  - location 15
- dsmc command
  - using options 51
- dsmcutil utility
  - commands for installing client services 590
  - options for installing client services 596
  - overview 589
- duplicate file names
  - avoiding 105
- dynamic and shared serialization 157

## E

- EA data 134
- editor option 244
- EFS encrypted files
  - incremental backup 479
  - restoring 524
  - selective backup 560
- enable8dot3namesupport option 245
- enablearchiveretentionprotection option 247
- enablelanfree option 249
- encryption
  - multiple clients under a single node name 194
  - of file data 63
  - saving encryption key password 251

- encryption processing
  - determine encryption cipher used in current session 63
  - excluding files from 257
  - include files for 280
  - query systeminfo command 514
- encryptiontype option 250
- encryptkey option
  - encryptkey=generate 251
  - encryptkey=prompt
  - encryptkey=save 251
- enhanced query schedule 145
- enhanced query schedule command 511
- environment prerequisites 5
- environment variables 19
- error log
  - controlling the size 253
  - pruning 255
  - specifying path and file name 254
- error processing, summary of options 181
- errorlogmax option 253
- errorlogname option 254
- errorlogretention option 255
- estimate function
  - for backup/restore processing 72
- event log
  - system object, back up 439
  - system object, query 498
  - system object, restore 533
- event logging
  - scheduler 147
- event-based policy retention protection
  - archive 163
  - backup 163
- exclude options 257
  - exclude.archive 32
  - exclude.backup 32
  - exclude.compression 32
  - exclude.dir 32
  - exclude.file 32
  - exclude.file.backup 32
  - exclude.image 32
  - exclude.subfile 32
  - exclude.systemobject 32
  - processing 41
  - testing 40
  - wildcard characters 37, 39
- exclude.image option 32
- excluding files
  - remotely accessed 37
  - system files 37
  - using wildcard characters 39
  - wildcard characters 37
- executable file
  - return codes from 151
- expire command 472
- extended attributes (EA data) 134

## F

- file
  - restoring active and inactive versions 523
- file names
  - avoiding duplicate 105
- file space
  - delete 135
  - determining fsID 232
  - excluding 33

- file space (*continued*)
  - performing an image backup 442
  - specifying NAS or client objects 208
- file specification
  - maximum allowed on commands 430
- file systems
  - deleted 98
  - excluding from backup processing 257
  - image backup of 82
  - supported 64
- filelist option 262
- filename option 264
- files
  - archive
    - directory tree 138
  - archive a list of 140, 262
  - archived, overriding management class 160
  - archiving 433
  - archiving more than one file specification 140
  - assigning management classes 98
  - backing up Microsoft Dfs 101
  - backing up open 97
  - backup, how managed 62
  - binding management classes to 161
  - compressing during archive or backup 222
  - definition of changed 65
  - delete after archive 229
  - delete individual archives from server file space 142, 463
  - delete individual backups from server file space 75, 465
  - encryption 63
  - excluding groups 37, 39
  - include-exclude
    - creating in Unicode format 278
  - including groups 37, 39
  - managing growth during compression 221
  - maximum file size for operations 63
  - processing include-exclude 41
  - query archive information 488
  - query backup information 491
  - query user access 487
  - renaming file spaces that are not Unicode to
    - Unicode-enabled 202, 475, 559
  - reset Windows archive attribute after backup 342
  - restore, point-in-time 125
  - restoring files belonging to another node 133
  - restoring from command line 107
  - restoring from GUI 107
  - restoring to another workstation 134
  - retrieve archives using command line 143
  - retrieving files belonging to another node 133
  - retrieving to another workstation 134
  - revoke access 462
  - sorting list of 54
- filesonly option 266
- filesystem
  - delete 468
- firewall
  - establishing communications through 23, 273, 394
  - specify whether server or client initiates sessions
    - through 361
  - specifying TCP/IP ports for the Web client 422
  - using Web client through 422
- fixed drives
  - backing up 99
- format and language
  - summary of options 180
- fromdate option 267

- fromnode option 268
- fromtime option 269
- frsprimaryrestore option 270
- full incremental
  - comparing with journal-based, incremental-by-date 68
  - definition 64
  - description 65
  - when to use 68
- fuzzy backup 158

## G

- getting started
  - changing your password 45
  - client scheduler 45
  - command-line session 45
  - displaying online help 45
  - ending a session 45
  - GUI session 45
  - sorting file lists 45
  - Web client session 45
- graphical user interface
  - changing password 53
  - delete individual files or images from server file space 465
  - displaying active and inactive backup versions 54, 106
  - displaying online help 55
  - displaying processing status 95
  - enabling for local backup set restore 299
  - enabling local backupset 122
  - ending a session 55
  - performing image backup 86
  - restore files and directories 107
  - starting a session 48
  - using to back up objects 71
- group
  - backup, specifying full or differential 307
- group backup
  - display active and inactive objects 277
  - display all members of 364
  - overview 76
  - specify name of group 271
  - specify virtual file space name for 409
- groupname option 271
- GUI
  - ending a session 55
  - overriding management class during archive 160
  - performing point-in-time restore 126
  - starting a session 48
- guitreeviewafterbackup option 272

## H

- hardware requirements 5
- help
  - displaying online 55
  - Internet resources 55
  - online forum 55
  - service and technical support 55
- help command 474
- httpport option 273

## I

- ifnewer option 274

- image
  - restoring 120
  - using chkdsk to repair 120, 537
  - using fsck to repair 120
- image backup
  - configuring online image backup 28
  - considerations 83
  - deleting 465
  - excluding files from 257
  - file systems or logical volumes 442
  - include files for; assign management class to 280
  - incremental-by-date image backup 85
  - offline and online 83
  - perform 82
  - point-in-time restore 443
  - revoke access 462
  - specifying selective or incremental 307
  - using incremental mode 443
  - using the GUI 86
  - using with file system incremental 85, 443
  - using with incremental backup 84, 443
  - using with incremental-by-date 85
- imagegapsize option 275
- imagnetofile option 276
- inactive backup versions
  - displaying 54, 106, 491
  - restoring 106
- inactive option 277
- inexcl option 278
- include option
  - management class 159
  - processing 41
  - wildcard characters 37, 39
- include-exclude list
  - creating 31
  - query order of processing 505
  - size restriction 37, 41
  - testing 40
  - to control processing 62
- include-exclude options file
  - specifying path and file name of 278
  - Unicode-enabled file spaces 278
- include-exclude processing
  - options for 32
  - overview 32
- incrbydate option 286
- incremental backup
  - back up new and changed files with modification date later than last backup 286
  - by date 74
  - description 65
  - EFS encrypted files 479
  - new and changed files 65
  - of directories
    - processing overview 66
  - overview 64
  - performing online 479
  - specify memory-conserving algorithm for 305
  - store file space image under different file space name 479
  - using client GUI 71
  - using command line 73
  - using the client command line 73
  - using with image backup 84, 443
- incremental command 475
  - journal-based backup 476
- incremental image backup 443
- incremental option 287



- incremental-by-date
  - comparing with incremental, journal-based 68
  - description 68
  - of directories
    - processing overview 68
  - overview 64
  - using command line 73
  - when to use 68
- incremental-by-date backup
  - using client GUI 71
  - using the client command line 73
  - using with image backup 85
- incrthreshold option 288
- install
  - backup-archive scheduler service 585
- installation requirements 5
- installing
  - overview 13
- installing from server CD 8
- installing Tivoli Storage Manager
  - CD, installing from 8
  - silent installation 20
- interactive mode 428
- interactive session
  - ending 482
  - starting 50, 482
  - using 482

**J**

- Java GUI
  - configuration restrictions 49
- journal configuration file
  - how to configure 601
- journal service
  - journal configuration file settings 601
- journal-based backup 66, 476
  - comparing with incremental, incremental-by-date 68
  - excluding directories 35
  - excluding files 35
  - performing traditional full incremental, instead of 314, 476
  - specifying configuration settings 601
  - specifying how to respond to unsuccessful expire of object 288
  - when to use 68
- journalized file space
  - specifying directories with active files to expire 288
- journalpipe option 290

**K**

- keyboard
  - accessibility xv
  - disability xv
  - shortcut keys xv

**L**

- LAN-free data movement
  - enabling 59, 249
  - enabling communications for 291, 294
  - options 60
  - prerequisites 59
  - specifying shared memory port for 293
- lanfreecommmethod option 291

- lanfreeshmport option 293
- lanfreetcport option 294
- lanfreetcpserveraddress option 295
- language option 296
- last access date
  - specifying whether to update during backup or archive 66, 331
- latest option 298
- limited bandwidth
  - backup 61
- local backupset
  - enabling GUI for local restore 122
- local snapshot
  - associating local snapshot with server file space 79
- localbackupset option 299
- location option 300
- log
  - DSM\_LOG environment variable 254, 352
  - error log, controlling the size 253
  - error log, pruning 253
  - error log, specifying path and file name 253
  - errorlogname option 254
  - errorlogretention option 254
  - schedlogname option 352, 557
  - schedlogretention option 352, 557
  - specifying path and file name 254, 352, 557
- logical volume
  - image backup of 82
  - restoring 120
- Logical Volume Snapshot Agent (LVSA)
  - configuring for online image backup 28
  - configuring for open file support 29
- loop command 482
- LVSA
  - bring up an application after LVSA starts 327
  - configuring for online image backup 28
  - configuring for open file support 29
  - quiesce an application before LVSA starts 336

**M**

- macro command 484
- manageservices option 302
- management class
  - assigning 98
- management classes
  - assigning to directories 160, 234
  - assigning to files 159
  - binding archive files to 140
  - binding to files 161
  - default 153
  - displaying 155
  - displaying information about 506
  - how Tivoli Storage Manager uses 62
  - overriding during archive processing 160
  - overriding the default 159
  - processing 159
  - questions to consider 159
  - selecting for files 159
  - specifying with include option 159
  - using management class, example 159
- maxcmdretries option 304
- memory requirements 5
- memoryefficientbackup option 305
- messages
  - displaying on screen 407
  - specifying language type 296

- messages (*continued*)
  - stop displaying 339
- Microsoft Cluster Server cluster nodes
  - FAQs 620
  - installing scheduler service 609
  - installing Tivoli Storage Manager 609, 620
- Microsoft Dfs junction
  - process as junction or directory 233
  - restore 523
- Microsoft Dfs trees and files
  - backing up 101
- migrating from earlier versions of the product 1
- migration
  - Web client 4
  - Web client language files 4
- mobile dial-up support 53
- mode option 307
- modes
  - batch 427
  - interactive (loop) 428
- modified mode 158
- monitor option 309
- monitor process command 485
- multiple sessions
  - backup objects, using 95

## N

- Named Pipe communication method
  - options 167
- namedpipename option 310
- NAS
  - backing up file systems 87
  - deleting file spaces 135, 468
  - query node command 507
  - restore file systems 126, 540
  - restore NAS command 540
- NAS file servers using CIFS
  - backup and restore 91
- NAS file systems backup
  - command line 90
  - web client GUI 88
- nasnodename option 311
- NetApp
  - backing up CIFS share definitions 95
- Network Attached Storage (NAS)
  - assigning management class to file systems 280
  - excluding files from backup 257
  - monitoring backup or restore operations 309
  - specifying for query 404
  - specifying full or differential backup 307
  - specifying node name for operations 311
  - specifying whether to save Table of Contents for each file system backup 399
- Network Attached Storage (NAS) file server
  - backup file systems 87, 447
  - cancel backup and restore processes 460, 485
  - deleting file spaces 135, 468
  - display file spaces on server 499
  - display nodes for which admin ID has authority 507
  - Querying file system images belonging to 491
  - restore file systems 126, 540
- Network Data Management Protocol (NDMP) 8
- Network File System (NFS)
  - backup file systems 95
- networked file systems
  - include-exclude statements 35

- no query restore 109
- node
  - specifying type to query 404
- node name 15
- node name, setting 16
- node option, summary 167
- nodename option 134, 312
- nojournal option 314
- noprompt option 315
- NTFS file spaces
  - back up 99
- NTFS security information
  - bypass processing of 365
  - specifying whether to compute CRC for comparison of 366
- numberformat
  - specifying 316
- numberformat option 316

## O

- ODBC driver 5, 10
- offline image backup 83
- online help
  - displaying 55
  - online forum 55
  - service and technical support 55
- online image backup 83
- Open Database Connectivity (ODBC) driver 5, 10
- open file support
  - include files for 280
  - installing and configuring 29
  - number of times LVSA attempts snapshot file system idle wait time before operation fails 369
  - overview 69
  - snapshot 138
  - specify amount of time before a snapshot is taken 371
  - specifying location of the cache files 367
  - specifying the size of the cache files 368
- open registration
  - permissions 31
  - using 31
- operating system re-installation 114
- operating system requirements 6
- optfile option 318
- options
  - activatekey 192
  - archive, summary 168
  - archmc 193
  - asnodename 194
  - asrmode 196
  - auditlogging 197
  - auditlogname 200
  - authorization, summary 181
  - autofsrename 202
  - backup, summary 168
  - backupsetname 204
  - casesensitiveaware 205
  - central scheduling, summary 179
  - changingretries 207
  - class 208
  - clientview 209
  - clusternode 214
  - collocatebyfilespec 215
  - command processing, summary 181
  - commmethod 217
  - commrestartduration 219

options (*continued*)

- commrestartinterval 220
- communication, summary 166
- compressalways 221
- compression 222
- computername 224
- console 225
- dateformat 227
- deletefiles 229
- description 230
- detail 232
- dfsbackupmntpnt 233
- diagnostics 183
- dirmc 234
- dironly 235
- disablenqr 236
- diskbuffsize 237
- diskcachelocation 238
- domain 239
- domain.image 241
- domain.nas 242
- editor 244
- enable8dot3namesupport 245
- enablearchiveretentionprotection 247
- enablelanfree 249
- encryptiontype 250
- encryptkey
  - encryptkey=generate 251
  - encryptkey=prompt 251
  - encryptkey=save 251
- errorlogmax 253
- errorlogname 254
- errorlogretention 255
- exclude
  - exclude.archive 32, 257
  - exclude.backup 32, 257
  - exclude.compression 32, 257
  - exclude.dir 32, 257
  - exclude.encrypt 257
  - exclude.file 32, 257
  - exclude.file.backup 32, 257
  - exclude.fs.nas 257
  - exclude.image 32, 257
  - exclude.subfile 257
  - exclude.systemobject 32, 257
  - exclude.systemservice 257
  - wildcard characters 37, 39
- filelist 262
- filename 264
- filesonly 266
- format and language, summary 180
- fromdate 267
- fromnode 268
- fromtime 269
- frsprimaryrestore 270
- general rules when entering with commands 184
- groupname 271
- guitreeviewafterbackup 272
- httpport 273
- ifnewer 274
- imagegapsize 275
- imagetofile 276
- inactive 277
- includexl 278
- include
  - wildcard characters 37, 39
- include.archive 280

options (*continued*)

- include.backup 280
- include.compression 280
- include.encrypt 280
- include.file 280
- include.fs 280
- include.fs.nas 280
- include.image 280
- include.subfile 280
- include.systemobject 280
- include.systemstate 280
- incrbydate 286
- incremental 287
- incrthreshold 288
- journalpipe 290
- lanfreecommmethod 291
- lanfreeshmport 167, 293
- lanfreetcpsport 294
- lanfreetcpserveraddress 295
- language 296
- latest 298
- localbackupset 299
- location 300
- manageservices 302
- maxcmdretries 304
- memoryefficientbackup 305
- mode 307
- monitor 309
- namedpipename 310
- nasnodename 311
- nodename 312
- nojournall 314
- noprompt 315
- numberformat 316
- optfile 318
- order of processing (precedence) 183
- password 319
- passwordaccess 320
- pick 322
- pitdate 323
- pittime 324
- postnschedulecmd 325
- postschedulecmd 325
- postsnapshotcmd 327
- preschedulecmd 329
- preschedulecmd 329
- preservelastaccessdate 331
- preservepath 333
- presnapshotcmd 336
- querschedperiod 338
- quiet 339
- replace 340
- resetarchiveattribute 342
- resourceutilization 344
- restore and retrieve, summary 176
- retryperiod 347
- revokeremoteaccess 348
- runasservice 349
- schedcmddisabled 350
- schedlogmax 351
- schedlogname 352
- schedlogretention 353
- schedmode 355
- schedrestretrdisabled 357
- scrolllines 358
- scrollprompt 359
- sessioninitiation 361

- options (*continued*)
  - shmport 363
  - showmembers 364
  - skipntpermissions 365
  - skipntsecuritycrc 366
  - snapshotcachelocation 367
  - snapshotcachesize 368
  - snapshotfsidleretries 369
  - snapshotfsidlewait 371
  - snapshotproviderfs 373
  - snapshotproviderimage 374
  - snapshotroot 375
  - specifying in commands 183
  - srvprepostscheddisabled 378
  - srvprepostsnapdisabled 379
  - ssl 380
  - subdir 381
  - subfilebackup 383
  - subfilecachepath 384
  - subfilecachesize 385
  - tapeprompt 386
  - tcpadminport 388
  - tcpbuffsize 389
  - tcpcadaddress 390
  - tcpclientaddress 391
  - tcpclientport 392
  - tcpnodelay 393
  - tcpport 394
  - tcpserveraddress 395
  - tcpwindowsize 396
  - timeformat 397
  - toc 399
  - todate 401
  - totime 402
  - transaction processing, summary 182
  - txnbytelimit 403
  - type 404
  - usedirectory 405
  - v2archive 406
  - verbose 407
  - verifyimage 408
  - virtualfsname 409
  - virtualnodename 410
  - VMCHost 412
  - VMCPW 413
  - VMCUSER 414
  - VMLIST 415
  - wasexphome 416
  - washome 417
  - wasndhome 418
  - wasnode 419
  - wastype 420
  - wasuser 421
  - Web client, summary 183
  - webports 422
- Owner security information (SID)
  - back up 99

## P

- parameters
  - yes and no, alternatives 191
- partial incremental
  - definition 64
  - incremental by date, running 74
- password
  - changing 53, 568

- password (*continued*)
  - number of characters 54
  - setting 319
  - setting for client 30
  - specifying whether to generate automatically or set as user prompt 320
  - using 49
  - valid characters 54
- password option 319
- passwordaccess option 320
- performance
  - improving speed of backups, restores, archives, retrieves 167
  - restore operations 108
  - transaction options 182
  - transaction processing 403
- Persistent Storage Manager
  - back up 88
- pick option 322
- pitdate 323
- pittime option 324
- point-in-time restore 125
  - image backup 443
- policies, storage management 153
- policy domains
  - default policy domain 153
  - standard policy domain 153
- policy sets
  - active policy set 153
- portable media
  - restoring backupsets 122
- postschedulecmd option 325
- postsnapshotcmd option 327
- Preferences editor
  - excluding domains from back up 73
- preschedulecmd option 329
- preparation overview
  - ASR 115
- preparation procedure
  - ASR 116
- preschedulecmd option 329
- preservelastaccessdate option 331
- preservepath option 333
- Presnapshotcmd option 336
- preview command 486
- previewing
  - excluded files 486
  - included files 486
- Primary group SID
  - back up 99
- processing options
  - authorization 181
  - backup and archive 168
  - central scheduling 179
  - communication 166
  - diagnostics 183
  - error processing 181
  - format and language 180
  - node option 167
  - overview 165
  - restore and retrieve 176
  - setting 18
  - specifying in commands 183
  - transaction processing 182
  - using 165
  - Web client 183

processing time  
  estimating 72  
proxied session restrictions 77, 142

## Q

query  
  ASR 490  
  backups, establish point-in-time 323, 324  
  based on date and time of backup, archive 267, 269  
  COM+ database system object 497  
  display active and inactive objects 277  
  event log system object 498  
  files for another node 268  
  group  
    command 501  
    display members of 364  
  include-exclude list 505  
  nodes to which client has proxy authority 140  
  nodes to which client has proxy node authority 76  
  process directories only (not files) 235  
  registry system object 509  
  specify end date and time to search for backups or  
  archives 401, 402  
  specify location of backup set for 300  
  specifying amount of information that displays on  
  screen 358  
  specifying description for 230  
  specifying NAS or client objects 208  
  specifying scrolling preferences after displaying  
  information on screen 359  
  specifying whether to include subdirectories for 381  
  system and boot files system object 513  
  system information 514  
  system objects 516  
  system state 517  
  WebSphere Application Server  
    Application Server 419, 420, 518  
    Network Deployment Manager 419, 420, 518  
    Network Deployment Manager and Application  
    Server 419, 420, 518  
  Windows Management Instrumentation repository 520  
query access command 487  
query archive command 488  
query ASR command 490  
query backup command 491  
query backupset command 495  
query complusdb command 497  
query eventlog command 498  
query file space command 499  
query group command 501  
query image command 503  
query inclexcl command 505  
query mgmtclass command 506  
query node command 507  
query options command 508  
query registry command 509  
query restore command 510  
query schedule  
  enhanced 145  
query schedule command 511  
query schedule command, enhanced 511  
query session command 512  
query sysfiles command 513  
query systeminfo command 514  
  encryption processing 514  
query systemobject command 516

query systemstate command 517  
query VM command 519  
query was command 518  
query wmi command 520  
querschedperiod option 338  
quiet option 339

## R

raw logical volume  
  image backup of 82  
  restoring 120  
rebinding files to a different management class 162  
recommendations  
  default diskbuffsize setting 237  
  discontinue use of largecommbuffers xxvi, 182  
  group include-exclude options in client options file 32  
  include-exclude list 37  
  minimum include-exclude list 37  
  runasservice option 349  
  setting the WebSphere Application Server security  
  password 569  
  specifying diskbuffsize xxv  
  to simplify retrieve 58  
  traditional scheduler system resources 575  
  Using the VirtualCenter 412, 413, 414  
Recommendations  
  backing up Microsoft Dfs files  
    Dfs junction metadata 101  
  do not write on volume while backup is in progress 83  
  long user and group names 64  
  restoring WebSphere Application Server data 131  
  set up SSL in the Integrated Solutions Console (ISC)  
  environment 26  
  use the CAD to manage the client scheduler 22  
  using dynamic imagetype 83  
  using LAN-free with VMware Consolidated Backup 92  
recover  
  disk 134  
registering  
  client with server 30  
  using closed registration 30  
  using open registration 30, 31  
registry  
  system object, back up 450  
  system object, query 509  
  system object, restore 542  
remote network connection  
  establishing 53, 151  
remotely accessed files  
  excluding 37  
  UNC names 37  
removable media  
  back up 99  
repairing installations 12  
replace option 340  
resetarchiveattribute option 342  
resourceutilization option 344  
restart restore command 521  
restartable restore 109  
restartable restore sessions, display 510  
restore  
  activate Windows Registry key to update Windows  
  Registry after restore 192  
  active and inactive file versions 523  
  active version 106  
  ASR recovery mode 196

- restore (*continued*)
  - authorizing another user 132
  - Automated System Recovery (ASR) files 113, 528
  - backup set
    - supported tape devices 529
  - backups, establish point-in-time 323, 324
  - backupsets
    - overview 122
  - based on date and time of backup 267, 269
  - classic (also known as standard) 109
  - COM+ database system object 532
  - create list of backup versions to 322
  - display active and inactive objects 277
  - EFS encrypted files 524
  - estimating processing time 72
  - event log system object 533
  - files belonging to another node 133
  - files for another node 268
  - files only 266
  - from file spaces that are not Unicode-enabled 524
  - from other workstation types 134
  - from portable media
    - overview 122
  - group
    - command 534
  - GUI, displaying active and inactive versions 54
  - image 120
    - considerations 536
    - enable detection of bad sectors on target volume 408
    - to a file 276
    - using chkdsk tool to repair 120, 537
    - using fsck tool to repair 120
  - image, suppress confirmation prompt 315
  - improving speed using share memory 167
  - inactive version 106
  - large number of files 108
  - list of files 262
  - local backup set using the GUI 299
  - local backupset using the GUI 122
  - logical volume 120
  - maximum file size 63
  - Microsoft Dfs junction 523
  - Microsoft Dfs trees and files 120
  - most recent backup version 298
  - NAS file systems 126
  - NAS file systems using command line 129
  - NAS file systems using web client 126
  - NAS files and directories using web client 127
  - NetApp CIFS shares 130
  - no query 109
  - overview 105
  - point-in-time 125
  - point-in-time, using command line 126
  - point-in-time, using GUI 126
  - primary restoration of sysvol 270
  - primary tasks 105
  - process directories only (not files) 235
  - raw logical volume 120
  - registry system object 542
  - regulate resources server and client use during
    - processing 344
  - replace existing file with latest backup 274
  - restartable 109
  - sorting file list 54
  - sparse files 524
  - sparse files to a non-NTFS file system 524
  - sparse files, size restriction for 524
- restore (*continued*)
  - specify end date and time to search for backups 401, 402
  - specify location of backup set for 300
  - specifying whether to include subdirectories for 381
  - specifying whether to prompt before overwriting existing files 340
  - specifying whether to wait for tape to mount if it is required 386
  - standard (also known as classic) 109
  - starting a Web client session 51
  - summary of options 176
  - system and boot files system object 544
  - system objects 111, 545
  - system state 547
  - to different workstation 410
  - using commands 107
  - using the GUI 107
  - using universal naming convention names 106
  - virtual volumes and mount points 522
  - VMware Consolidated Backup 130
  - Websphere Application Server
    - Application Server 419, 420, 549
    - Network Deployment Manager 419, 420, 549
  - Windows Management Instrumentation repository 552
  - Windows Server 2003 system state 111
  - Windows Vista system state 111
  - workstation, to another 134
- restore asr command 528
- restore backupset command 529
- restore command 522
  - using multiple 108
- restore complusdb command 532
- restore eventlog command 533
- restore group command 534
- restore image command 536
- restore NAS command 540
- restore procedure
  - ASR 118
- restore registry command 542
- restore sysfiles command 544
- restore systemobject command 545
- restore systemstate command 547
- restore was command 549
- restore wmi command 552
- restoring
  - directories 107
  - files 107
  - files and directories 107
- restrictions
  - asnodename option 194
  - runasservice and encryptkey 349
  - runasservice and passwordaccess 349
  - runasservice and replace 349
  - within a proxied session 77, 142
- retain extra versions attribute 156
- retain only versions attribute 157
- retain versions attribute 158
- retention grace period
  - archive 154, 162
  - backup 154, 162
- retrieve
  - archive copies 142
  - archive files by name 143
  - archives using command line 143
  - authorizing another user 132
  - based on date and time of archive 267, 269
  - files belonging to another node 133

- retrieve (*continued*)
  - files for another node 268
  - files only 266
  - from other workstation types 134
  - improving speed using share memory 167
  - list of files 262
  - maximum file size 63
  - primary tasks 137
  - process directories only (not files) 235
  - replace existing file with latest archive if existing file is newer 274
  - sorting file list 54
  - specify end date and time to search for archives 401, 402
  - specifying description for 230
  - specifying whether to include subdirectories for 381
  - specifying whether to prompt before overwriting existing files 340
  - specifying whether to wait for tape to mount if it is required 386
  - starting a Web client session 51
  - summary of options 176
  - to different workstation 410
  - workstation, to another 134
- retrieve command 553
- retryperiod option 347
- return codes for operations 151
- revokeremoteaccess option 348
- runasservice option 349

## S

- SAN
  - restoring backup sets using 529
- schedcmddisabled option 350
- schedlogmax option 351
- schedlogname option 352
- schedlogretention option 353
- schedmode option 355
- schedrestretrdisabled option 357
- schedule command 557
- schedule log
  - controlling the size 351
  - specifying number of days to keep entries and whether to save pruned entries 353
  - specifying path and file name to store schedule log information 352
- scheduled (automated) backups
  - closing files before back up 97
  - displaying scheduled work 145, 147
  - establish remote connection for 151
  - mobile dial-up support 151
  - options for 145
  - overview 145, 571
  - process commands after back up 325
  - process commands before back up 329
  - restart applications after back up 97
  - starting 23
- scheduled commands
  - enabling-disabling 150
- scheduled events, displaying 511
- scheduled services 150
  - disabling scheduled commands 350
  - restrictions for NAS file systems 87
- scheduler
  - configuring 21
  - displaying scheduled work 145, 147
  - establish remote connection for 151

- scheduler (*continued*)
  - event logging 147
  - mobile dial-up support 151
  - options for 145
  - overview 145, 571
  - specify whether server or client initiates sessions through firewall 361
  - specifying number of hours between contacts to server for scheduled work 338
  - specifying number of minutes between attempts to process scheduled commands 347
  - specifying polling mode or prompted mode 355
  - specifying whether to disable execution of restore or retrieve operations 357
  - starting 23
  - troubleshooting 581
- scheduler information
  - completed work 583
  - scheduled work 583
- scheduler log
  - log file 583
  - schedlogname option 583
  - schedlogretention option 583
  - understanding the 583
- scheduler service
  - installing on Microsoft Cluster Server cluster nodes 609
  - installing on Veritas Cluster Server cluster nodes 609
- Scheduler Service Configuration Utility 23
- scheduler wizard 23
- scheduling
  - client node proxy backup 77, 78
- scrolllines option 358
- scrollprompt option 359
- Secure Socket Layer (SSL)
  - establishing communications with 26
- selective backup
  - EFS encrypted files 560
  - overview 64, 69, 74
  - using client GUI 71
  - using command line 74
  - using the client command line 73
- selective command 559
- serialization
  - copy serialization
    - dynamic 157
    - shared static 157
    - static 157
- server
  - communicating with 16
  - establishing communications through firewall 23
  - establishing communications with 15
  - establishing communications with Secure Socket Layer (SSL) 26
  - query Active Directory for communication method and server with which to connect 405
  - specifying TCP/IP address of Tivoli Storage Manager server 395
  - specifying TCP/IP port address for 394
- service and technical support 55
- session information, displaying 512
- sessioninitiation option 361
- set access command 563
  - restore-retrieve authorization 132
- set event command 566
- set password command 568
- set waspassword command 569

- setting
  - user privileges 53
- setting environment variables
  - DSM\_CONFIG 19
  - DSM\_DIR 19
  - DSM\_LOG 19
- shared dynamic serialization 157, 207
- shared files and directories
  - back up 100
- shared memory communication method 291
- Shared memory communication method
  - options 167
- Shared memory communication protocol 167
- shared static serialization 157, 207
- shmport option 363
- shortcut keys
  - accessibility xv
  - disability xv
  - keyboard xv
- showmembers option 364
- silent installation 10
- skiptntpermissions option 365
- skiptntsecuritycrc option 366
- snapshot
  - open file support 138
- snapshotcachelocation option 367
- snapshotcachesize option 368
- snapshotfsidleretries option 369
- snapshotfsidlewait option 371
- snapshotproviderfs option 373
- snapshotproviderimage option 374
- snapshotroot option 375
- software requirements 6
  - ASR 115
- sparse files
  - restore size restriction 524
  - restoring 524
  - restoring to a non-NTFS file system 524
- srvprepostscheddisabled option 378
- srvprepostsnapdisabled option 379
- ssl option 380
- standard (classic) restore 109
- standard management class
  - copy destination 158
  - copy frequency 156
  - copy group name 156
  - copy mode
    - absolute 158
    - modified 158
  - copy serialization 157
  - copy type 156
  - default values 155
  - retain extra versions 156
  - retain only version 157
  - retain versions 158
  - versions data deleted
    - active versions 156
    - inactive versions 156
  - versions data exists 156
- standard policy domain 153
- starting
  - automatically 53
  - overview 13
- starting a session
  - batch mode 49
  - interactive mode 50
- static serialization 157
- static, shared serialization 157
- storage
  - displaying restartable restore sessions 510
- Storage Agent
  - using for LAN-free data movement 59, 249
- storage area network
  - restoring backup sets using 249, 529
  - using for LAN-free data movement 59, 249
- storage management policies 153
  - assigning management classes to files 98
  - copy groups 153
  - default management class 153
  - display on backup-archive client or Web client GUI 98
  - include-exclude list 154
  - management classes 153
  - policy domains
    - default 153
    - standard 153
  - policy sets
    - active policy set 153
- subdir option 381
- subdirectories
  - archive 140
  - include in backup 74
- subfilebackup option 383
- subfilecachepath option 384
- subfilecachesize option 385
- summary of changes for version 5.3.0 xxiii
- summary of changes for version 5.3.2 xix
- support
  - gathering system information for 225, 264, 514
- Swing-enabled browser
  - necessary to run Web client 51
- System access control list (auditing information)
  - back up 99
- system and boot files
  - query 513
  - system object, back up 451
  - system object, restore 544
- system files
  - excluding, recommended 37
- system information
  - gathering 225, 264, 514
- system objects
  - assign management class 280
  - back up 79, 452
  - certificate server database
    - back up 437
  - COM+ database
    - back up 438
    - query 497
    - restore 532
  - event log
    - back up 439
    - query 498
    - restore 533
  - exclude from backup processing 36, 257, 258
  - primary restoration of sysvol 270
  - query 516
  - registry
    - back up 450
    - query 509
    - restore 542
  - restore 111, 545
  - restore from backup set 529
  - system and boot files
    - back up 451



- system objects (*continued*)
  - system and boot files (*continued*)
    - query 513
    - restore 544
  - Windows Management Instrumentation
    - back up 459
    - query 520
  - Windows Management Instrumentation repository
    - restore 552

- system recovery
  - Windows Server 2003, Windows XP 114
- system restore 114
- system services

- assigning management class 81
- components, exclude from backup 81
- exclude from backup processing 257
- system state
  - assigning management class 37, 81, 280
  - back up 80, 453
  - display active and inactive objects 277
  - query 517
  - restore 547
  - restore from backup set 529
  - restoring 111

## T

- tapeprompt option 386
- tasks
  - closed registration 30
  - open registration 30
- TCP/IP communication method
  - options 166
- tcpadminport option 388
- tcpbuffsize option 389
- tcpaddress option 390
- tcpclientaddress option 391
- tcpclientport option 392
- tcpnodelay option 393
- tcpserveraddress option 395
- tcpwindow size option 396
- time format
  - specifying 397
- timeformat option 397
- Tivoli products and documentation
  - feedback survey xv
  - ordering xv
- Tivoli Storage Manager
  - communication methods 7
  - disk space requirements 5
  - environment prerequisites 5
  - FAQs 620
  - hardware requirements 5
  - installation 9
  - installation requirements 5
  - installing on Microsoft Cluster Server cluster nodes 609, 620
  - installing on Veritas Cluster Server cluster nodes 609
  - memory requirements 5
  - migrating from earlier versions of the product 1
  - online forum 55
  - password 49
  - software requirements 6
- Tivoli Storage Manager client
  - authentication 47
- Tivoli Storage Manager password
  - using 49

- toc option 399
- today option 401
- totime option 402
- transaction processing 403
  - summary of options 182
- tsmjbbd.ini
  - configuring 601
- txnbytelimit option 403
- type option 404

## U

- UNC names
  - excluding files 37
  - remotely accessed files 37
  - restore files, using 106
- Unicode
  - renaming file spaces that are not Unicode to Unicode-enabled 202, 475, 559
  - restore from file spaces that are not Unicode-enabled 524
- universal naming convention
  - back up shared files and directories using 100
  - using to specify domain list 100
- universal naming convention names
  - restore, using 106
- upgrading
  - online image 4
  - Open File Support (OFS) 4
- upgrading, modifying, or reinstalling 12
- usedirectory option 405
- user privileges
  - setting 53

## V

- v2archive option 406
- verbose option 407
- verifyimage option 408
- Veritas Cluster Server cluster nodes
  - FAQs 620
  - installing scheduler service 609
  - installing Tivoli Storage Manager 609, 620
- versions data
  - deleted attribute 156
  - exists attribute 156
- virtualfsname option 409
- virtualnodename option 410
- VMCHost option 412
- VMCPW option 413
- VMCUSER option 414
- VMLIST option 415
- VMware Consolidated Backup
  - backing up data 92
  - restoring data 130
- Volume Shadowcopy Service (VSS)
  - configuring for online image backup 29
  - configuring for open file support 29
- VSS (see Volume Shadowcopy Service) 29

## W

- wasexphome option 416
- washome option 417
- wasndhome option 418
- wasnode option 419
- wastype option 420

- wasuser option 421
- Web client
  - configuring in cluster environment 609
  - enable to run in a Swing-enabled browser 51
  - establishing communications through firewall 23, 273
  - restrict administrator from accessing client running Web client 348
  - restrictions for NAS file systems 87
  - specifying TCP/IP port address for 273
  - starting 51
  - summary of options 183
  - supported browsers 51
  - unsupported functions 57
  - using through a firewall 422
- web client GUI
  - NAS file systems backup 88
- webports option 422
- WebSphere Application Server
  - back up WebSphere Application Server-Express 416
  - query
    - Application Server 518
    - Network Deployment Manager 518
  - restore
    - Application Server 455, 549
    - Network Deployment Manager 455, 549
  - set WebSphere node name, user name, password 569
  - specify override base install path for Application Server 417
  - specify override base install path for Network Deployment Manager 418
  - specifying full or differential backup 307
- wildcard characters
  - guidelines 431
  - include or exclude groups of files 37
  - specifying a drive specification in dsm.opt 38
  - to include or exclude groups of files 39
- Windows archive attribute
  - reset after backup 342
- Windows Management Instrumentation repository
  - system object back up 459
  - system object query 520
  - system object restore 552
- workstation
  - restore or retrieve from another type 134





Program Number: 5608-ISM  
5608-ISX  
5608-SAN

Printed in USA

SC32-0146-01



Spine information:

IBM Tivoli Storage Manager for Windows **Version 5.5**

**Backup-Archive Clients Installation and User's Guide**

