

IBM Tivoli Storage Manager for UNIX and Linux
Backup-Archive Clients
Version 7.1.8

Installation and User's Guide



IBM Tivoli Storage Manager for UNIX and Linux
Backup-Archive Clients
Version 7.1.8

Installation and User's Guide



Note:

Before you use this information and the product it supports, read the information in “Notices” on page 745.

This edition applies to version 7, release 1, modification 8 of IBM Tivoli Storage Manager (product numbers 5608-E01, 5608-E02, and 5608-E03) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 1993, 2017.**

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

Contents

Tables xi

About this publication xiii

Who should read this publication.	xiii
Publications	xiv
Conventions used in this publication	xiv
Reading syntax diagrams	xiv

What's new for Version 7.1.8. xvii

Chapter 1. Installing the Tivoli Storage Manager backup-archive clients 1

Upgrading the backup-archive client	1
Upgrade path for clients and servers	1
Additional upgrade information.	1
Automatic backup-archive client deployment	2
Client environment requirements	3
AIX client environment.	3
AIX client installable components	3
System requirements for the AIX client	3
AIX client communication methods.	4
Backup-archive client features that are available on AIX	4
HP-UX Itanium 2 client environment	4
HP-UX Itanium 2 client installable components	4
System requirements for HP-UX Itanium 2 clients	4
HP-UX Itanium 2 client communication methods.	5
Linux on Power Systems client environment.	5
Linux on Power Systems client installable components.	5
System requirements for clients on Linux on Power Systems	5
Linux on Power Systems client communication methods.	5
Linux x86_64 client environment.	6
Linux x86_64 client installable components	6
System requirements for Linux x86_64 clients	6
Linux x86_64 client communication methods	6
Linux on System z client environment.	7
Linux on System z client installable components.	7
System requirements for Linux on System z clients	7
Linux on System z client communication methods.	7
Mac OS X client environment.	7
Mac OS X client installable components	7
System requirements for Mac OS X clients	8
Mac OS X client communication methods.	8
Oracle Solaris client environment	8
Oracle Solaris client installable components	8
System requirements for Oracle Solaris clients	8
Oracle Solaris client communication methods	9

NDMP support requirements (Extended Edition only)	9
Installation requirements for backing up and archiving Tivoli Storage Manager FastBack client data	9
Install the UNIX and Linux backup-archive clients	10
Installing the AIX client	10
Uninstalling the AIX client	13
Installing the HP-UX Itanium 2 client.	14
Increasing the default limit of the data segment size	16
Uninstalling the HP-UX Itanium 2 client.	16
Installing the backup-archive client on Linux on Power Systems	17
Uninstalling the backup-archive client on Linux on Power Systems	20
Installing the backup-archive client on Linux on Power Systems (Little Endian)	22
Uninstalling the backup-archive client on Linux on Power Systems (Little Endian)	25
Installing the backup-archive client on Ubuntu Linux on Power Systems (Little Endian).	26
Uninstalling the client on Ubuntu Linux on Power Systems (Little Endian)	27
Installing the Linux x86_64 client	28
Uninstalling the Linux x86_64 client	31
Installing the Ubuntu Linux x86_64 client	33
Uninstalling the Ubuntu Linux x86_64 client	35
Installing the Linux on System z client	36
Uninstalling the Linux on System z client	39
Installing the Mac OS X client	41
Uninstalling the Mac OS X client	42
Installing the Oracle Solaris client	43
Uninstalling the Oracle Solaris client	45
Software updates	46
Installing the client management service to collect diagnostic information.	46

Chapter 2. Configure the Tivoli Storage Manager client 47

UNIX and Linux client root and authorized user tasks.	47
Enable non-root users to manage their own data	49
Enabling encryption for backup-archive client users	49
Client options file overview	50
Creating and modifying the client system-options file	51
Creating a default client-user options file	53
Creating a customized client user-options file	55
Environment variables.	55
Set language environment variables	55
Set processing environment variables.	57
Set Bourne and Korn shell variables	58
Set C shell variables	59
Set API environment variables	59
Web client configuration overview.	59

Configuring the web client on AIX, HP-UX, Linux, Mac, and Solaris systems	60
Configuring the scheduler	61
Comparison between client acceptor-managed services and traditional scheduler services	61
Configuring the client to use the client acceptor service to manage the scheduler	62
Start the client scheduler	63
Scheduling events using the command-line client	63
Configuring Tivoli Storage Manager client/server communication across a firewall	66
Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer	68
Creating a symbolic link to access the latest GSKit library	71
Certificate Authorities root certificates	72
Configure your system for journal-based backup .	73
Journal daemon configuration	73
JournalSettings stanza	75
JournalExcludeList stanza	75
JournalExcludeFilesystemSettings stanza	76
Overriding stanzas	79
Client-side data deduplication	79
Configuring the client for data deduplication .	83
Excluding files from data deduplication	85
Automated client failover configuration and use .	87
Automated client failover overview	87
Requirements for automated client failover .	88
Restrictions for automated client failover .	88
Failover capabilities of Tivoli Storage Manager components	90
Configuring the client for automated failover .	90
Determining the status of replicated client data	92
Preventing automated client failover	93
Forcing the client to fail over	93
Configuring the client to back up and archive Tivoli Storage Manager FastBack data	94
Cluster environment configuration and use	95
Overview of cluster environments	96
Active/Active: Pool cluster resources	96
Active/Passive: Fault tolerant	96
Concurrent access	96
Configuring the backup-archive client in a cluster environment	96
Enabling web client access in a Cluster Environment	102
Migrating legacy AIXIBM PowerHA SystemMirror setups	104
AIX configuration considerations prior to performing snapshot-based file backups and archives	105
Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups . . .	106
Protecting clustered-data ONTAP NetApp file server volumes	107
SnapMirror support for NetApp snapshot-assisted progressive incremental backup (snapdiff)	110
Register your workstation with a server	113
Closed registration	114
Open registration	114

Creating an include-exclude list	115
Include-exclude options	116
Exclude file spaces and directories	116
Exclude files and directories from a journal-based backup	117
Control processing with exclude statements	117
System files to exclude	119
Include and exclude files that contain wildcard characters	120
Include and exclude groups of files with wildcard characters	121
Examples using wildcards with include and exclude patterns	122
Symbolic link and alias processing	123
Determine compression and encryption processing	123
Preview include-exclude list files	124
Include and exclude option processing	125

Chapter 3. Getting started 129

Configuring security settings to connect the client to the Tivoli Storage Manager server	129
Configuring by using the default security settings (fast path)	129
Configuring without automatic certificate distribution	131
Secure password storage	133
Tivoli Storage Manager client authentication . .	134
Starting a Java GUI session	135
Tivoli Storage Manager password	136
Setup wizard	137
Starting a command-line session	137
Using batch mode	137
Issuing a series of commands by using interactive mode	138
Specifying input strings that contain blank spaces or quotation marks	138
Starting: Additional considerations	139
Using the web client in the new security environment	139
Starting a web client session	140
User privileges	140
Start the client scheduler automatically	141
Changing your password	141
Sorting file lists using the backup-archive client GUI	143
Displaying online help	144
Ending a session	144
Online forums	145

Chapter 4. Backing up your data . . . 147

Planning your backups	147
Which files are backed up	147
When to back up and when to archive files . . .	148
Pre-backup considerations (UNIX and Linux) .	149
LAN-free data movement	149
LAN-free prerequisites	149
LAN-free data movement options	149
Incremental backups on memory-constrained systems	150

Incremental backups on systems with a large number of files	150
Include-exclude options to control processing	152
Data encryption during backup or archive operations	152
File system and ACL support	153
Maximum file size for operations.	157
Long user and group names	157
Mac OS X volume names	158
Mac OS X volume naming precautions	158
Mac OS X volume naming precautions on dual boot systems	158
Mac OS X Unicode enablement	159
Mac OS X Time Machine backup disk	159
Performing an incremental, selective, or incremental-by-date backup (UNIX and Linux)	160
Full and partial incremental backup	160
Journal-based backup on AIX and Linux	162
Incremental-by-date backup	164
Comparing incremental-by-date, journal-based, and NetApp snapshot difference to full incremental and partial incremental backups	165
Snapshot differential backup with an HTTPS connection	167
Running a snapshot differential backup with an HTTPS connection	167
Selective backup	168
Solaris global zone and non-global zones backups	168
Saving access permissions	169
Setting a virtual mount point	169
Backing up data using the Java GUI.	169
Backing up data using the command line	170
Deleting backup data.	173
Deleting file spaces	174
Backing up files from one or more file spaces for a group backup (UNIX and Linux).	174
Backing up data with client-node proxy support (UNIX and Linux).	175
Enabling multiple node operations from the GUI	176
Setting up encryption	176
Scheduling backups with client-node proxy support	177
Examples of how to schedule a backup of an IBM PowerHA SystemMirror cluster.	178
Scheduling a backup of a GPFS file system	179
Associate a local snapshot with a server file space (UNIX and Linux).	180
Image backup	181
Performing prerequisite tasks before creating an image backup	182
Volume device type support for an image backup	183
Utilizing image backups to perform file system incremental backups	184
Method 1: Using image backups with file system incremental backups	185
Method 2: Using image backups with incremental-by-date image backups	185
Comparing methods 1 and 2	186

Performing an image backup using the GUI	187
Performing an image backup using the command line	188
Snapshot-based file backup and archive and snapshot-based image backup	188
Protecting Btrfs file systems	189
Backing up and restoring Btrfs file systems	190
Backing up and restoring Btrfs subvolumes	191
Back up NAS file systems using Network Data Management Protocol	192
Backing up NAS file systems with the web client GUI using NDMP protocol.	193
Back up NAS file systems using the command line.	195
Backup network file systems	197
Back up NFS file systems with the global namespace feature.	197
Back up AIX workload partition file systems	198
Backing up Solaris Zettabyte file systems	199
AIX JFS2 encrypted file system backup.	200
Back up AIX JFS2 extended attributes	201
Backing up VMware virtual machines	201
Preparing the environment for full backups of VMware virtual machines	203
Creating full backups for VMware virtual machines.	204
Parallel backups of virtual machines.	205
Back up and archive Tivoli Storage Manager FastBack data	209
Display backup processing status.	210
Backup (UNIX and Linux): Additional considerations	212
Stored files	213
Special file systems	214
NFS or virtual mount points	214
Management classes	214
Back up symbolic links	214
Examples: Incremental or selective backup of symbolic links	215
Incremental backup of a domain only	216
Hard links	216
Sparse files	217
NFS hard and soft mounts	217
Deleted file systems	218
Opened files.	218
Wildcard characters	219

Chapter 5. Restoring your data 221

Restoring an image	221
Restoring an image using the GUI	222
Restoring an image using the command line	224
Restore data from a backup set	224
Restore backup sets: considerations and restrictions	226
Backup set restore.	227
Restoring backup sets using the GUI	228
Backup set restores using the client command-line interface	229
Restoring or retrieving data during a failover.	229
Restore an image to file	231
Manage GPFS file system data with storage pools	231

Restoring data to a point in time	233
Restore AIX encrypted files.	234
Restore AIX workload partition file systems	235
Restore NAS file systems	236
Restoring NAS file systems using the web client	236
Options and commands to restore NAS file	
systems from the command line	237
Restore active or inactive backups	238
Restoring data using the GUI	239
Command line restore examples	240
Examples: Command line restores for large	
amounts of data	241
Standard query restore, no-query restore, and	
restartable restore	242
Standard query restore process	242
No-query restore process	242
Restartable restore process	243
Restoring Solaris Zettabyte (ZFS) file systems	243
Additional restore tasks	244
Authorizing another user to restore or retrieve	
your files.	244
Restoring or retrieving files from another client	
node	245
Restore or retrieve files to another workstation	246
Restoring a disk in case of disk loss	246
Deleting file spaces	247
Enable SELinux to restore files on the Red Hat	
Enterprise Linux 5 client	248

Chapter 6. Archive and retrieve your data (UNIX and Linux). 249

Archive files.	249
Archiving data with the GUI	249
Archive data examples by using the command	
line.	250
Associate a local snapshot with a server file	
space	251
Archiving data with client node proxy	251
Deleting archive data.	253
Advanced archive tasks	254
Access permissions	254
Archive and retrieve symbolic links	254
Hard links	255
Retrieve archives	255
Retrieving data with the GUI	256
Retrieve data examples by using the command	
line.	256
Archive management classes	257

Chapter 7. Tivoli Storage Manager scheduler overview. 259

Examples: Blank spaces in file names in schedule	
definitions	260
Preferential start times for certain nodes	260
Scheduler processing options	261
Evaluate schedule return codes in schedule	
scripts.	262
Return codes from preschedulecmd and	
postschedulecmd scripts.	263

Client-acceptor scheduler services versus the	
traditional scheduler services	264
Setting the client scheduler process to run as a	
background task and start automatically at startup .	264
Examples: Display information about scheduled	
work	266
Display information about completed work	268
Specify scheduling options	269
Scheduler options for commands.	269
Enable or disable scheduled commands	269
Manage multiple schedule requirements on one	
system	270

Chapter 8. Client return codes 273

Chapter 9. Storage management policies 275

Policy domains and policy sets	275
Management classes and copy groups	276
Display information about management classes	
and copy groups	277
Copy group name attribute.	277
Copy type attribute	277
Copy frequency attribute	278
Versions data exists attribute	278
Versions data deleted attribute	278
Retain extra versions attribute.	278
Retain only version attribute	278
Copy serialization attribute.	279
Copy mode parameter	279
Copy destination attribute	280
Retain versions attribute.	280
Deduplicate data attribute	280
Select a management class for files	280
Assign a management class to files	281
Override the management class for archived files	282
Select a management class for directories	282
Bind management classes to files	283
Rebind backup versions of files	284
Retention grace period	284
Event-based policy retention protection.	285
Archive files on a data retention server.	285

Chapter 10. Processing options 287

Processing options overview	287
Communication options	288
TCP/IP options	288
Shared memory options	289
Server options	289
Backup and archive processing options.	290
Restore and retrieve processing options.	299
Scheduling options	302
Format and language options	304
Command processing options	304
Authorization options	304
Error processing options.	305
Transaction processing options	306
Web client options.	306
Diagnostics options	307
Using options with commands	307

Entering options with a command	308	I	Encryptkey	384
Initial command-line-only options	313		Errorlogmax	385
Client options that can be set by the Tivoli Storage Manager server	314		Errorlogname	387
Client options reference	316		Errorlogretention	388
Absolute	316		Exclude options	389
Afmskipuncachedfiles	317		Controlling symbolic link and alias processing	392
Archmc	318		Controlling compression processing	393
Archsymlinkasfile	319		Processing NAS file systems	393
Asnodename	319		Virtual machine exclude options	394
Session settings and schedules for a proxy operation	320		Fbbranch	396
Auditlogging	321		Fbclientname	396
Auditlogname	323		Fbpolicyname	398
Autodeploy	325		Fbreposlocation	399
Autofsrename	326		Fbserver	400
Automount	328		Fbvolumename	401
Backmc	329		Filelist	402
Backupsetname	329		Filename	405
Basesnapshotname	330		Filesonly	406
Cadlistenonport	332		Followsymbolic	407
Changingretries	333		Forcefailover	408
Class	333		Fromdate	409
Collocatebyfilespec	334		Fromnode	410
Commmethod	335		Fromowner	411
Commrestartduration	337		Fromtime	412
Commrestartinterval	337		Groupname	412
Compressalways	338		Groups (deprecated)	413
Compression	339		Guitreeviewafterbackup	413
Console	341		Host	414
Createnewbase	342		Httpport	414
Datacenter	343		Hsmreparsetag	414
Datastore	344		Ieobjtype	416
Dateformat	344		Ifnewer	417
Dedupcachepath	348		Imagegapsize	417
Dedupcachesize	348		Imagetofile	418
Deduplication	349		Inactive	419
Defaultserver	350		Incllexcl	420
Deletefiles	351		Considerations for Unicode-enabled clients	421
Description	351		Include options	421
Detail	353		Controlling symbolic link and alias processing	426
Diffsnapshot	354		Compression and encryption backup processing	427
Diffsnapshotname	355		Processing NAS file systems	427
Dirmc	356		Virtual machine include options	428
Dirsonly	357		Incrbydate	436
Disablenqr	358		Incremental	438
Diskbuffsize	359		Instrlogmax	438
Diskcachelocation	359		Instrlogname	439
Domain	361		Lanfreecommmethod	441
Domain.image	366		Lanfreeshmport	442
Domain.nas	366		Lanfreetcppport	443
Domain.vcd	368		Lanfreessl	444
Domain.vmfull	369	I	Lanfreetcpserversaddress	444
Dontload	374		Latest	445
Dynamicimage	375		Localbackupset	446
Efsdecrypt	376		Makesparsefile	447
Enablearchiveretentionprotection	377		Manageservices	448
Enablededupcache	378		Maxcmdretries	450
Enableinstrumentation	380		Mbobjrefreshthresh	450
Enablelanfree	382		Mbpctrefreshthresh	451
Encryptiontype	383			

Memoryefficientbackup	452	Snapshotroot	528
Mode	454	Srvoptsetencryptiondisabled	530
Monitor	457	Srvprepostscheddisabled	531
Myreplicationserver	458	Srvprepostsnapdisabled	532
Nasnodename	460	Ssl	533
Nfstimeout	461	Sslacceptcertfromserv	534
Nodename	462	Ssldisablelegacytls	535
Nojournal	463	Sslfipsmode	536
Noprompt	464	Sslrequired	536
Nrtablepath	464	Stagingdirectory	538
Numberformat	465	Subdir	539
Optfile	467	Tapeprompt	541
Password	468	Tcpadminport	542
Passwordaccess	470	Tcpbuffsize	543
Passworddir	471	Tcpcadaddress	544
Pick	472	Tcpclientaddress	545
Pitdate	473	Tcpclientport	545
Pittime	474	Tcpnodelay	546
Postschedulecmd/Postnschedulecmd	475	Tcpport	547
Postsnapshotcmd	476	Tcpserveraddress	548
Preschedulecmd/Prenschedulecmd	478	Tcpwindowsize	548
Preserveaccessdate	479	Timeformat	549
Preservepath	480	Toc	552
Presnapshotcmd	483	Todate	554
Querschedperiod	484	Totime	554
Querysummary	485	Txnbytelimit	555
Quiet	486	Type	556
Quotesareliteral	487	Updatectime	557
Removeoperandlimit	488	Useexistingbase	558
Replace	489	Usereplicationfailover	559
Replserverguid	490	Users (deprecated)	560
Replservername	492	V2archive	560
Replsslport	493	Vappmc	561
Repltcpport	495	Vcdhost	562
Repltcpserveraddress	497	Vcdpw	562
Resourceutilization	499	Vcduser	563
Regulating backup and archive sessions	499	Verbose	564
Regulating restore sessions	500	Verifyimage	565
Multiple client session considerations	500	Virtualfsname	565
Retryperiod	501	Virtualmountpoint	566
Revokeremoteaccess	502	Virtualnodename	567
Schedcmddisabled	503	Vmbackdir	568
Schedcmdexception	504	Vmbackupmailboxhistory	569
Schedlogmax	504	Vmbackuptype	570
Schedlogname	506	Vmchost	570
Schedlogretention	507	Vmcpw	571
Schedmode	509	Vmctlmc	572
Schedrestretrdisabled	510	Vmcuser	573
Scrolllines	511	Vmdatastorethreshold	574
Scrollprompt	512	Vmdefaultdvportgroup	575
Servername	513	Vmdefaultdvswitch	576
Sessioninitiation	514	Vmdefaultnetwork	577
Shmport	516	Vmenabletemplatebackups	578
Showmembers	517	Vmlimitperdatastore	579
Skipacl	517	Vmlimitperhost	580
Skipaclupdatecheck	518	Vmmaxparallel	581
Snapdiff	519	Vmmaxvirtualdisks	582
Snapdiffhttps	524	Vmmc	584
Snapshotcachesize	525	Vmnoprmdisks	584
Snapshotproviderfs	526	Vmnovrmdisks	585
Snapshotproviderimage	527	Vmpreferdagpassive	586

Vmprocessvmwithindependent	587	Query Backup	668
Vmprocessvmwithprdm	588	Query NAS file system images	670
Vmskipctlcompression	589	Query Backupset	670
Vmskipmaxvirtualdisks	590	Query Backupset without the backupsetname	
Vmskipmaxvmdks.	591	parameter	672
Vmtagdatamover.	591	Query Filespace	674
Vmverifyifaction	594	Query NAS file spaces	676
Vmverifyiflatest	595	Query Group	676
Vmvstortransport	597	Query Image	678
Vmtimeout	598	Query Inclexcl	679
Webports	599	Query Mgmtclass	680
Wildcardsareliteral	600	Query Node	681
		Query Options	682
Chapter 11. Using commands	603	Query Restore	683
Start and end a client command session	606	Query Schedule	684
Process commands in batch mode	606	Query Session	684
Process commands in interactive mode	607	Query Systeminfo	686
Enter client command names, options, and		Query VAPP	687
parameters	607	Query VM	691
Command name	608	Restart Restore.	692
Options	608	Restore	693
Options in interactive mode	608	Restore from file spaces that are not	
Parameters	608	Unicode-enabled	698
File specification syntax	609	Restore Backupset	698
Wildcard characters	610	Restore backup sets: considerations and	
Client commands reference	611	restrictions	701
Archive	611	Restore backup sets in a SAN environment	702
Archive FastBack	614	Restore Backupset without the backupsetname	
Backup FastBack	616	parameter	702
Backup Group	618	Restore Group	705
Backup Image	620	Restore Image	707
Static, dynamic, and snapshot image backup	624	Restore NAS	710
Utilizing image backup to perform file system		Restore VAPP	712
incremental backup	624	Restore VM	715
Backup NAS	626	Retrieve	720
Backup VAPP	628	Retrieve archives from file spaces that are not	
Backup VM	634	Unicode-enabled	722
Cancel Process	641	Schedule	723
Cancel Restore	642	Selective	725
Delete Access	642	Associate a local snapshot with a server file	
Delete Archive	643	space	728
Delete Backup	645	Set Access	728
Delete Filespace	648	Set Event.	731
Delete Group	649	Set Netappsvm	733
Expire.	651	Set Password	734
Help	652	Set Vmtags	740
Incremental	654		
Journal-based backup.	657	Appendix. Accessibility features for	
Incremental-by-Date	659	the Tivoli Storage Manager product	
Associate a local snapshot with a server file		family.	743
space	660		
Loop	660	Notices	745
Macro	661		
Monitor Process	662	Glossary	749
Preview Archive	663		
Preview Backup	664	Index	751
Query Access	665		
Query Archive	665		

Tables

1. Upgrading the client from different server versions	2	40. Backup and restore capabilities for VMware virtual machines on Linux platforms.	201
2. AIX client communication methods	4	41. Client command line informational messages	210
3. Supported features on AIX	4	42. Backup set GUI restore restrictions	225
4. HP-UX Itanium 2 client communication methods	5	43. Backup set command-line restore restrictions	225
5. Linux on Power Systems client communication methods	6	44. Sample WPAR restore commands with dsm.opt file	236
6. Linux on Intel x86_64 client communication methods	6	45. NAS options and commands	238
7. Linux on System z client communication methods	7	46. Command-line restore examples	240
8. Mac OS X client communication methods	8	47. Command line archive examples	250
9. Oracle Solaris client communication methods	9	48. Symbolic link management table for archive and retrieve	255
10. HP-UX Itanium 2 client: Language codes for installation packages	15	49. Command line examples of retrieving archives	257
11. Package names, contents, and default directory	18	50. Sample classic query schedule output	267
12. Language pack identifiers	21	51. Sample enhanced query schedule output	268
13. Package names, contents, and default directory	22	52. Client return codes and their meanings	273
14. Package names, contents, and default directory	26	53. Default attribute values in the standard management class	277
15. Package names, contents, and default directory	29	54. TCP/IP options	288
16. Language pack identifiers	32	55. Shared memory communication options	289
17. Package names, contents, and default directory	33	56. Sample client system-options file	290
18. Language pack identifiers	36	57. Backup and archive processing options	290
19. Package names, contents, and default directory	37	58. Restore and retrieve processing options	299
20. Language pack identifiers	40	59. Scheduling options	302
21. Installation package names and descriptions	43	60. Format and language options	304
22. Tasks for root users and authorized users	47	61. Command processing options	304
23. Mac OS X authorization tools and associated Tivoli Storage Manager applications	49	62. Authorization options	304
24. Client acceptor-managed services versus traditional scheduler services.	62	63. Error processing options	305
25. Data deduplication settings: Client and server	82	64. Transaction processing options	306
26. Options for excluding file spaces and directories.	116	65. Web client options	306
27. Options for controlling processing using include and exclude statements	117	66. Diagnostics options	307
28. Wildcard and other special characters	121	67. Client command options	309
29. Using wildcard characters with include and exclude patterns	122	68. Options that are valid on the initial command line only	314
30. Options for controlling symbolic link and alias processing	123	69. Options that can be set by the Tivoli Storage Manager server	315
31. Options for controlling compression and encryption	124	70. Sample time format settings in the locale configuration (t_fmt line)	347
32. Working with your files using the backup-archive client GUI	143	71. Sample date format settings in the locale configuration (d_fmt line)	347
33. Supported file systems and ACL support	153	72. Interaction of domain definitions from several sources.	365
34. Maximum file size	157	73. Other optional parameters	425
35. Command-line backup examples	170	74. Incremental command: Related options	521
36. Volume device-type support for an image backup.	183	75. Effects of server and client SSL settings on success or failure of login attempts	537
37. Comparing incremental image backup methods	186	76. Sample time format settings in the locale configuration (t_fmt line)	552
38. LVM1 and LVM2 image operation comparisons	188	77. Sample date format settings in the locale configuration (d_fmt line)	552
39. NAS options and commands	195	78. Commands	603
		79. Wildcard characters	611
		80. Archive command: Related options	613
		81. Archive FastBack command: Related options	615
		82. Backup FastBack command: Related options	617

83. Backup Group command: Related options	620	101. Query Node command: Related options	682
84. Backup Image command: Related options	622	102. Query Options command: Related options	683
85. Backup NAS command: Related options	627	103. Query Systeminfo command: Related options	687
86. Backup VAPP command: Related options	630	104. Query VAPP command: Related options	688
87. Delete Archive command: Related options	644	105. Query VM command: Related options for	
88. Delete Backup command: Related options	647	VMware virtual machine queries.	691
89. Delete Filespace command: Related options	649	106. Restore command: Related options	696
90. Delete Group command: Related options	650	107. Restore Backupset command: Related options	700
91. Expire command: Related options.	652	108. Restore Group command: Related options	705
92. Incremental command: Related options	655	109. Restore Image command: Related options	708
93. Query Archive command: Related options	666	110. Restore NAS command: Related options	711
94. Query Backup command: Related options	669	111. Restore VAPP command: Related options	714
95. Query Backupset command: Related options	672	112. Restore VM command: Related options when	
96. Query Backupset command: Related options	673	restoring VMware virtual machines	718
97. Query Filespace command: Related options	674	113. Retrieve command: Related options	721
98. Query Group command: Related options	677	114. Schedule command: Related options	724
99. Query Image command: Related options	678	115. Selective command: Related options	726
100. Query Mgmtclass command: Related options	681		

About this publication

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 enables 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:

- A server program that acts as a backup and archive server for distributed workstations and file servers.
The server program also supplies hierarchical storage management (HSM) services, and enables systems to perform as a migration server.
- An administrative client program that you can access from a web browser or from the command line. The program enables a Tivoli Storage Manager administrator to control and monitor server activities, define storage management policies for backup, archive, and space management services, and set up schedules to perform those services at regular intervals.
- An application programming interface (API) that you can use to enhance an existing application with storage management services. When an application is registered with a server as a client node, the application can back up, restore, archive, and retrieve objects from storage.
- A web backup-archive client that enables an authorized administrator, help desk person, or other users to perform backup, restore, archive, and retrieve services by using a web browser on a remote system.

Associated with Tivoli Storage Manager, but sold separately, are the IBM Tivoli Storage Manager for Space Management and IBM Tivoli Storage Manager HSM for Windows client programs. These products automatically migrate eligible files to storage to maintain specific levels of free space on local file systems and automatically recall migrated files when they are accessed. It also enables users to migrate and recall specific files.

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

Related concepts:

“Planning your backups” on page 147

“What’s new for Version 7.1.8” on page xvii

Chapter 1, “Installing the Tivoli Storage Manager backup-archive clients,” on page 1

Who should read this publication

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

Publications

The Tivoli Storage Manager product family includes IBM Tivoli Storage FlashCopy® Manager, IBM Tivoli Storage Manager for Space Management, IBM Tivoli Storage Manager for Databases, and several other storage management products from IBM.

To view IBM product documentation, see IBM Knowledge Center.

Conventions used in this publication

This publication uses the following typographical conventions:

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

Reading syntax diagrams

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

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

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

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

Symbols

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

- * Asterisk
- { } Braces

- : Colon
- , Comma
- = Equal Sign
- - Hyphen
- () Parentheses
- . Period
- Space
- " quotation mark
- ' single quotation mark


Variables

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


►► cmd_name—*<var_name>*—————►◄

Repetition

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

►►  repeat—————►◄

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

►►  repeat—————►◄

Notes:

- 1 Specify *repeat* up to 5 times.

Required choices

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

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

►► cmd_name —————►◄

Optional choices

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



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.



Repeatable choices

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

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



Defaults

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

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



What's new for Version 7.1.8

IBM Tivoli Storage Manager Version 7.1.8 introduces new features and updates.

For a list of new features and updates in this release and previous Version 7.1 releases, see Backup-archive client updates.

New and changed information in this product documentation is indicated by a vertical bar (|) to the left of the change.

Protect your client with enhanced security settings

Several changes are introduced in the backup-archive client to work with the Tivoli Storage Manager V7.1.8 server, which provides enhancements to improve the security between client and server communications.

After the Tivoli Storage Manager server is upgraded to V7.1.8 and configured with the improved security protocol, and the backup-archive client is upgraded to V7.1.8, the security settings for the client must be configured to work with the security enhancements on the server. For more information about the configuring the client for different security scenarios, see “Configuring security settings to connect the client to the Tivoli Storage Manager server” on page 129.

The following changes are available in this release:

New `sslacceptcertfromserv` option, and changes to three existing SSL-related options

To simplify the process of distributing server certificates, Tivoli Storage Manager now includes a new `sslacceptcertfromserv` option to control whether the backup-archive client or the API application accepts and trusts the Tivoli Storage Manager server's Secure Sockets Layer (SSL) public certificate the first time they connect.

In addition, three SSL-related options `ssl`, `ssldisablelegacytls`, and `sslrequired` have changed. Their operation varies depending on whether operation is with Tivoli Storage Manager server V7.1.8 or V8.1.2, or with Tivoli Storage Manager server V7.1.7 and earlier levels, and V8.1.1 and earlier V8 levels.

For a description of these option settings and important version information, see:

- “`sslacceptcertfromserv`” on page 534
- “`ssl`” on page 533
- “`ssldisablelegacytls`” on page 535
- “`sslrequired`” on page 536

The Tivoli Storage Manager password location is changed

Beginning in V7.1.8, the IBM Global Security Kit (GSKit) keystores are used to store all Tivoli Storage Manager passwords. When you upgrade to Tivoli Storage Manager V7.1.8 backup-archive client, the existing passwords are automatically migrated to new locations on the client system.

For more information about the new password location and related considerations, see “Secure password storage” on page 133.

A new way of importing server certificates is available with the **dsmscert** utility. For information about importing server certificates, see “Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer” on page 68.

Because of the use of the new secure password storage, new procedures for granting password access to non-administrative users are available.

For UNIX and Linux clients, see “Enable non-root users to manage their own data” on page 49.

Automatically deploy backup-archive client updates

The Tivoli Storage Manager server administrator can use server commands to schedule updates for one or more backup-archive clients. The updates can be fix packs or new releases. This feature was available in previous releases of Tivoli Storage Manager, but an improved procedure is available for V7.1.8.

For more information, see “Automatic backup-archive client deployment” on page 2.

Deprecated functions

The following functions are deprecated in this release:

lanfreessl and replsslport options

Two SSL-related options, **lanfreessl** and **replsslport**, are deprecated and unavailable if you are connecting to a Tivoli Storage Manager server V7.1.8 and later or V8.1.2 and later. These options are still valid and available if you are connecting to a Tivoli Storage Manager server V7.1.7 and earlier levels, and V8.1.1 and earlier V8 levels.

Trusted communications agent (TCA)

The trusted communications agent (TCA), previously used by non-root users in V8.1.0 and V7.1.6 and older clients, is no longer available. Root users can use other methods to allow non-root users to manage their files.

For more information, see:

- “The trusted communications agent is no longer available” on page 134
- “Enable non-root users to manage their own data” on page 49

Tivoli Storage Manager web client

You can no longer use the web client to connect to the Tivoli Storage Manager V7.1.8 and later server or V8.1.2 and later server. However, you can still use the web client to connect to Tivoli Storage Manager server V7.1.7 and earlier levels, or V8.1.1 and earlier V8 levels.

For more information, see “Using the web client in the new security environment” on page 139.

NDMP command-line and web client operations

If you are connecting to the Tivoli Storage Manager V7.1.8 and later server or V8.1.2 and later server, you can no longer use the client command-line interface or web client to back up or restore NAS file servers using Network Data Management Protocol (NDMP). Alternatively, you can use

Tivoli Storage Manager server commands with the administrative command-line client (**dsmdmc**) to restore NDMP data.

For more information, see the NDMP information in “Using the web client in the new security environment” on page 139.

Open registration

If the client is connecting to the Tivoli Storage Manager V7.1.8 and later server or V8.1.2 and later server, the open registration feature is no longer available. To connect to the Tivoli Storage Manager server, you must use closed registration. For more information, see “Closed registration” on page 114.

Open registration is still available if the client is connecting to Tivoli Storage Manager server V7.1.7 and earlier levels, or V8.1.1 and earlier V8 levels.

For more information about security improvements on the server, see *Protect your storage environment with an improved security protocol*.

An administrative user ID is no longer created by default with the REGISTER NODE server command

With Tivoli Storage Manager V7.1.8, the **REGISTER NODE** server command does not automatically create an administrative user ID that matches the node name. This product update is designed to optimize user authentication to a Lightweight Directory Access Protocol (LDAP) server.

This product update does not affect existing client nodes, but can affect the process of registering new client nodes, including but not limited to nodes for Tivoli Storage Manager backup-archive clients. In some cases, you might have to create an administrative user ID when you register a node. You can create the administrative user ID by issuing the **REGISTER NODE** command and specifying the **USERID** parameter. For information about the types of clients that are affected, see technote 7048963.

If you plan to use the web client, you must manually create an administrative user ID when you register a new node. For more information, see “Register your workstation with a server” on page 113.

Maintenance updates

The documentation is updated to include information from APARs.

Related information:

“About this publication” on page xiii

Chapter 1. Installing the Tivoli Storage Manager backup-archive clients

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

You can maintain backup versions of your files that you can restore if the original files are damaged or lost. You can also archive infrequently used files, preserve them in their current state, and retrieve them when necessary.

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

Related concepts:

“What’s new for Version 7.1.8” on page xvii

“Planning your backups” on page 147

Upgrading the backup-archive client

The following sections explain what you need to do if you are upgrading to IBM Tivoli Storage Manager backup-archive client Version 7.1.8 from a previous version.

Upgrade path for clients and servers

Tivoli Storage Manager clients and servers can be upgraded at different times. The combination of servers and clients that you deploy must be compatible with each other.

To prevent disruption of your backup and archive activities while you upgrade from one release to another, follow the compatibility guidelines for Tivoli Storage Manager clients and servers in technote 1053218.

For information about upgrading your current AIX® IBM PowerHA® SystemMirror® setups, see “Migrating legacy AIX IBM PowerHA SystemMirror setups” on page 104.

Additional upgrade information

When you upgrade the backup-archive client, there is additional information to consider before you use the new client software.

Be aware of the following information when you upgrade a backup-archive client:

- If you are upgrading from the Tivoli Storage Manager Version 7.1.2 or earlier backup-archive client on the HP-UX, Linux, or Oracle Solaris operating system, you must uninstall any previously installed language packages before you proceed with the upgrade.
- For Mac users, updates to the Mac OS X client contained in Tivoli Storage Manager V6.3, or newer versions, require you to consider the following items:
 - When you use the Mac OS X client that is provided in this release, ensure that the `dsm.sys` and `dsm.opt` files are encoded by using Unicode (UTF-8). UTF-8 encoding enables the use of characters from any language in the

options files. If your `dsm.sys` or `dsm.opt` files were previously encoded as MacRoman (or anything other than UTF-8), open them in an editor like TextEdit and save them with UTF-8 encoding, and without the `.txt` extension. Your include-exclude lists can be encoded as either UTF-8 or UTF-16. For more information about using Unicode, see “Considerations for Unicode-enabled clients” on page 421.

- Tivoli Storage Manager server file spaces that were created by Mac OS 9 clients cannot be managed by the Mac OS X client that was provided in Tivoli Storage Manager V6.3. Use `q file node f=d` on the server to list files stored for a node. Any Mac-platform files that do not start with a slash (/) were probably created by an older Mac client. You cannot restore or otherwise manage these files by using the Mac OS X client that is provided in this release. You can manage these files, but you must use a Mac client that is installed on a Tivoli Storage Manager V6.2.2 (or older) node.
- For a list of new and changed messages since the previous Tivoli Storage Manager release, see the `client_message.chg` file in the client package.

Automatic backup-archive client deployment

The Tivoli Storage Manager server administrator can automatically deploy a backup-archive client to update workstations where the backup-archive client is already installed.

The Tivoli Storage Manager server can be configured to automatically upgrade backup-archive clients on client workstations. The existing backup-archive clients must be at Version 6.4.3 or later.

The procedure for automatically deploying client upgrades depends on the version of the Tivoli Storage Manager server that you are upgrading the client from. The following table shows the client upgrade procedures for different versions of the server.

Table 1. Upgrading the client from different server versions

Server version	Procedure
V8.1.3 or later	Use the Tivoli Storage Manager Operations Center. For more information, see <i>Scheduling client updates</i> .
V7.1.8 or V8.1.2	See technote 2004596.
V7.1.6 and earlier servers, or V8.1.1 and earlier V8 servers	See technote 1673299.

Restrictions: The following restrictions apply to automatic client deployment:

- The Windows cluster services environment is not supported.
- Automatic deployment of AIX 6.1 clients is not supported.
- Only the backup-archive client can be deployed from the Tivoli Storage Manager server. Other related products such as Tivoli Storage Manager for Space Management, Tivoli Storage Manager HSM for Windows, Tivoli Storage Manager for Virtual Environments, and other Data Protection products are not supported. If a deployment of an unsupported product is attempted, the deployment process stops with a failure message.
- Do not schedule automatic client deployments to systems that have any of the following applications installed on them:
 - IBM Tivoli Storage Manager for Virtual Environments

- IBM Tivoli Storage Manager for Databases
- IBM Tivoli Storage Manager for Mail
- IBM Tivoli Storage Manager for Enterprise Resource Planning

Related reference:

“Autodeploy” on page 325

Client environment requirements

Each of the Tivoli Storage Manager clients has hardware and software requirements.

The following list shows the location of the environment prerequisites for each supported platform.

- “AIX client environment”
- “HP-UX Itanium 2 client environment” on page 4
- “Linux on Power Systems client environment” on page 5
- “Linux x86_64 client environment” on page 6
- “Linux on System z client environment” on page 7
- “Mac OS X client environment” on page 7
- “Oracle Solaris client environment” on page 8
- “NDMP support requirements (Extended Edition only)” on page 9

For current information about the client environment prerequisites for all of the supported backup-archive client platforms, see technote 1243309.

AIX client environment

This section contains client environment information, backup-archive client components, and hardware and software requirements for the AIX platform.

AIX client installable components

The backup-archive client is comprised of several installable components.

The installable components for the AIX client are as follows:

- Backup-archive command line client
- Administrative client
- Backup-archive client graphical user interface, which uses Oracle Java™ technology
- Backup-archive web client
- Tivoli Storage Manager 64-bit API

The API can be separately installed. The other components are all installed when you install the AIX package (`tivoli.tsm.client.api.64bit`).

System requirements for the AIX client

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

For software and hardware requirements for all supported versions of AIX clients, including the most recent fix packs, see technote 1052226.

AIX client communication methods

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

You can use the following communication methods with the Tivoli Storage Manager Version 7.1.8 AIX client:

Table 2. AIX client communication methods

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

Backup-archive client features that are available on AIX

This topic lists the features that are supported on AIX.

Table 3. Supported features on AIX

Features	Supported on AIX?
Backup-archive command-line and GUI	yes
Journal-based backup	yes
NetApp Snapshot Difference (snapdiff option)	yes
LAN-free operations	yes
Online image backup	yes
Offline image backup	yes

HP-UX Itanium 2 client environment

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

HP-UX Itanium 2 client installable components

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

You can install the following components with Tivoli Storage Manager Version 7.1.8:

- Backup-archive command-line client
- Administrative client
- Backup-archive Java graphical user interface (GUI)
- Web backup-archive client
- Tivoli Storage Manager API (64-bit)

System requirements for HP-UX Itanium 2 clients

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

For software and hardware requirements for all supported versions of HP-UX Itanium 2 clients, including the most recent fix packs, see technote 1197146.

HP-UX Itanium 2 client communication methods

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

You can use the following communication methods with the HP-UX Itanium 2 backup-archive client:

Table 4. HP-UX Itanium 2 client communication methods

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

Linux on Power Systems client environment

This section contains client environment information, backup-archive client components, and hardware and software requirements for the Linux on Power Systems™ client platforms.

Linux on Power Systems client installable components

The backup-archive client command-line, Java GUI, web backup-archive, and API comprise the Linux on Power Systems backup-archive client installable components.

You can install the following components with Tivoli Storage Manager Version 7.1.8:

- Backup-archive client
- Administrative client
- Backup-archive Java graphical user interface (GUI)
- Web backup-archive client
- Tivoli Storage Manager API (64-bit)

System requirements for clients on Linux on Power Systems

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

For software and hardware requirements for all supported versions of clients on Linux on Power Systems, including the most recent fix packs, see technote 1169963.

Linux on Power Systems client communication methods

Backup-archive clients on Linux on Power Systems can use either TCP/IP or shared memory as the communications method for client-server communications.

Table 5 on page 6 lists the available Linux on Power Systems client communications methods, and the Tivoli Storage Manager server operating systems that you can use them with.

Table 5. Linux on Power Systems client communication methods

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

Linux x86_64 client environment

This section contains client environment information, backup-archive client components, and hardware and software requirements for the Linux on Intel (Linux x86_64) platform.

Linux x86_64 client installable components

The backup-archive client command-line, Java GUI, web backup-archive, administrative client, and the API comprise the Linux on Intel (Linux x86_64) backup-archive client installable components.

You can install the following components with Tivoli Storage Manager Version 7.1.8:

- Backup-archive client
- Administrative client
- Backup-archive Java graphical user interface (GUI)
- Web backup-archive client
- Tivoli Storage Manager API

System requirements for Linux x86_64 clients

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

For software and hardware requirements for all supported versions of Linux x86_64 clients, including the most recent fix packs, see technote 1052223.

Linux x86_64 client communication methods

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

You can use the following communication methods with the Tivoli Storage Manager Version 7.1.8 Linux on Intel (Linux x86_64) client:

Table 6. Linux on Intel x86_64 client communication methods

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

Linux on System z client environment

This section contains client environment information, backup-archive client components, and hardware and software requirements for the Linux on System z[®] platform.

Linux on System z client installable components

The backup-archive client command-line, administrative client, web backup-archive client, and API comprise the Linux on System z backup-archive client installable components.

You can install the following components with Tivoli Storage Manager Version 7.1.8:

- Backup-archive client
- Administrative client
- Web backup-archive client
- Tivoli Storage Manager API

System requirements for Linux on System z clients

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

For software and hardware requirements for all supported versions of Linux System z clients, including the most recent fix packs, see technote 1066436.

Linux on System z client communication methods

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

You can use the following communication methods with the Tivoli Storage Manager Version 7.1.8 Linux on System z client:

Table 7. Linux on System z client communication methods

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

Mac OS X client environment

This section contains client environment information, backup-archive client components, and hardware and software requirements for the Mac OS X client.

Mac OS X client installable components

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

You can install the following components with Tivoli Storage Manager Version 7.1.8:

- Backup-archive client
- Administrative client

- Web backup-archive client
- Tivoli Storage Manager API
- Backup-archive Java graphical user interface (GUI)

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

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

System requirements for Mac OS X clients

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

For software and hardware requirements for all supported versions of Mac OS X clients, including the most recent fix packs, see technote 1053584.

Mac OS X client communication methods

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

You can use the following communication methods with the Tivoli Storage Manager Version 7.1.8 Mac OS X client:

Table 8. Mac OS X client communication methods

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

Oracle Solaris client environment

This section contains client environment information, backup-archive client components, and hardware and software requirements for the Oracle Solaris platform.

Oracle Solaris client installable components

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

You can install the following client components:

- Backup-archive client
- Administrative client
- Backup-archive Java graphical user interface (GUI)
- Web backup-archive client
- Tivoli Storage Manager API

System requirements for Oracle Solaris clients

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

For software and hardware requirements for all supported versions of Tivoli Storage Manager Oracle Solaris clients, including the most recent fix packs, see the following IBM support pages:

- For Oracle Solaris SPARC client requirements, see technote 1052211.

- For Oracle Solaris x86_64 client requirements, see technote 1232956.

Oracle Solaris client communication methods

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

You can use the following communication methods with the Oracle Solaris client:

Table 9. Oracle Solaris client communication methods

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

NDMP support requirements (Extended Edition only)

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

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

NDMP support requires the following hardware and software:

- Tivoli Storage Manager Extended Edition
- Tape drive and tape library. For supported combinations, go to: product information

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

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

You must install the following software:

- Tivoli Storage Manager FastBack Version 6.1
- Tivoli Storage Manager client V6.1.3.x (where x is 1 or higher) or V6.2 or later
- Tivoli Storage Manager server V6.1.3 or higher
- Tivoli Storage Manager Administration Center V6.1.3
 - Required only if you want to use integrated Tivoli Storage Manager FastBack - administration.

Starting with Tivoli Storage Manager V7.1, the Administration Center component is no longer included in Tivoli Storage Manager distributions. FastBack users who have an Administration Center from a previous server release, can continue to use it to create and modify FastBack schedules. If you do not already have an Administration Center installed, you can download the previously-released version from <ftp://public.dhe.ibm.com/storage/tivoli-storage-management/maintenance/admincenter/v6r3/>. If you do not already have an Administration Center installed, you must create and modify

FastBack schedules on the Tivoli Storage Manager server. For information about creating schedules on the server, see the Tivoli Storage Manager server documentation.

The Tivoli Storage Manager FastBack environment must be running. For information about installing and setting up Tivoli Storage Manager FastBack, see the product information at IBM Tivoli Storage Manager Fastback.

For information about integrating Tivoli Storage Manager and Tivoli Storage Manager FastBack, see Integrating Tivoli Storage Manager FastBack and Tivoli Storage Manager.

You can install the Tivoli Storage Manager client in one of the following ways:

- Install the backup-archive client on a workstation where the FastBack Disaster Recovery Hub is installed. In this case, the prerequisites are: the FastBack Disaster Recovery Hub setup, and the FastBack shell.
- Install backup-archive client on a workstation where neither the FastBack server or the FastBack Disaster Recovery Hub is installed. In this case, the FastBack shell is still required.

Related concepts:

“Configuring the client to back up and archive Tivoli Storage Manager FastBack data” on page 94

Install the UNIX and Linux backup-archive clients

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

Note: You must log on as the root user to install the backup-archive client on a UNIX or Linux workstation.

The supported UNIX and Linux clients and the location of the installation instructions for each client are listed here.

- “Installing the AIX client”
- “Installing the HP-UX Itanium 2 client” on page 14
- “Installing the backup-archive client on Linux on Power Systems” on page 17
- “Installing the backup-archive client on Linux on Power Systems (Little Endian)” on page 22
- “Installing the backup-archive client on Ubuntu Linux on Power Systems (Little Endian)” on page 26
- “Installing the Linux x86_64 client” on page 28
- “Installing the Ubuntu Linux x86_64 client” on page 33
- “Installing the Linux on System z client” on page 36
- “Installing the Mac OS X client” on page 41
- “Installing the Oracle Solaris client” on page 43

Related concepts:

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

Installing the AIX client

You can install the AIX backup-archive client from the product installation media.

About this task

In Tivoli Storage Manager Version 7.1.8, a 64-bit version of the AIX client is provided in the distribution libraries. You cannot upgrade a previously installed 32-bit AIX client to new the 64-bit AIX client. If you have a 32-bit client that is installed from a previous version of Tivoli Storage Manager, use SMIT to perform the following steps:

1. Uninstall the 32-bit client (`tivoli.tsm.client.ba`).
2. Uninstall any national language files that were previously installed.
3. Uninstall the API (`tivoli.tsm.client.api.32bit`).

Next, use SMIT to install the following packages in the Tivoli Storage Manager V7.1.8 distribution libraries, in the following order:

1. Install the 64-bit API (`tivoli.tsm.client.api.64bit`).
2. Install the 64-bit client (`tivoli.tsm.client.ba.64bit`).

If you already have a 64-bit Tivoli Storage Manager V6.3 (or newer) client installed, you can upgrade the client instead of uninstalling it and reinstalling it.

If you have a 64-bit client from an earlier version of Tivoli Storage Manager installed (for example, V6.1, or V6.2) you must uninstall the client, language packs, and API. Then, install the new Tivoli Storage Manager API and client.

All of the packages that are needed to install the client are in the AIX client package, and they overwrite any older runtime applications on your system during installation. The LibC (C Set ++) runtime library is required.

When you use the **installp** command to install this client, do not change the default field values for the following two choices:

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

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

Install the following packages. They are all provided on the installation media. You need an Extended Edition license to use the NAS client.

The following files are listed in order of dependency. For example, the API is dependent on the Global Security Kit (GSKit). When you install all of them using SMIT, you can select them (F7) in any order.

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

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

tivoli.tsm.client.ba.64bit
Installs the following 64-bit client files:

- Backup-archive Java client (GUI)
- Backup-archive web client
- Hardware backup client
- Snapshot hardware backup client

- NAS backup client

tivoli.tsm.filepath_aix

Installs the file path kernel extension that is required for journal-based backup.

tivoli.tsm.client.jbb.64bit

Installs the journal-based backup component.

Each package is installed in the following default installation directory:

- The backup-archive, web client, and administrative client (**dsmadmc**) 64-bit files are installed in the `/usr/tivoli/tsm/client/ba/bin64` directory.
- The Tivoli Storage Manager 64-bit API files are installed in the `/usr/tivoli/tsm/client/api/bin64` directory.
- The sample system-options file, `dsm.sys.smp`, is placed in the installation directory.

This installation procedure can be used to install new distributions or updates from downloaded installation media. The downloaded files that you use to install the client might be compressed. Depending on the package file format, either copy or extract the files to disk and use these instructions to install the components.

You can download the appropriate package file from one of the following websites:

- Download the client package from Passport Advantage or Fix Central.
- For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.

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

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

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

Procedure

1. Log in as the root user.
2. Mount the volume that you are installing from.
3. From the AIX command line, type `smitty install` and press Enter.
4. Select **Install and Update Software** and press Enter.
5. Select **Install and Update From ALL Available Software** and press Enter.
6. At the INPUT device/directory for software prompt, press the F4 key and specify the directory that contains the installation images, and press Enter.
7. At the SOFTWARE to install prompt, press the F4 key. Select the Tivoli Storage Manager file sets you want to install by pressing the F7 key. Then, press the Enter key.
8. On the Install and Update From ALL Available Software panel, press the F4 key to change any entry fields, or use the default fields. Press Enter twice to begin the installation.
9. After the installation completes, press F10 to exit.

Results

When file sets are installed, they are automatically committed on the system. The previous version of backup-archive client software is replaced by the newly installed version.

The backup-archive client files are installed in the `/usr/tivoli/tsm/client/ba/bin64` directory. If you move the client files to another directory, you must perform the following steps:

1. Make sure that the permissions of the installed files have not changed.
2. Update the symbolic links for the installed files in the following directories:
 - The `/usr/bin` directory
 - The `/usr/lib` directory for Tivoli Storage Manager libraries
3. Ensure that every user of the backup-archive client sets the `DSM_DIR` environment variable to the newly installed directory.

What to do next

After the installation completes, see Chapter 2, “Configure the Tivoli Storage Manager client,” on page 47 for required and optional tasks to complete before you use the backup-archive client.

Note:

- AIX workload partitions (WPAR) are supported as follows:
 - Supported in global environments
 - Supported with non-shared system WPARs
 - Supported with shared system WPARs (backup-archive client logs and configuration files must be defined to non-default locations)
 - No support for application WPARs
 - No support for image backup
 - No support for backup set restore from tape
- On AIX Version 6.1, if you are using encrypted file systems (EFS) with the backup-archive client, and if the EFS user keystore password is different from the user login password, the EFS keystore is not automatically opened when you log on. If the EFS keystore is not open when you log on, the client might not restore a non-EFS file into an EFS file system. You can prevent the EFS file system restore problem one of the following ways:
 - Start the backup-archive client by using the **efskeymgr -o** command. For example: **efskeymgr -o ./dsmj**
 - Synchronize the keystore password with the user login password by using the **efskeymgr -n** command. For example: **efskeymgr -n**

Uninstalling the AIX client

You can use the following procedures to uninstall the Tivoli Storage Manager AIX backup-archive client.

Before you begin

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

Procedure

1. Enter the following AIX command: **smitty remove**.
2. Press the ENTER key.
3. In the **SOFTWARE** name field, press F4 to list the Tivoli Storage Manager file sets that you want to uninstall; press the ENTER key.
4. Select the Tivoli Storage Manager file sets that you want to uninstall; press the ENTER key.

Note: The journal-based backup feature is contained in two file sets. Select both `tivoli.tsm.client.jbb.64bit` and `tivoli.tsm.filepath_aix`. If you uninstall the file sets one at a time, uninstall the `tivoli.tsm.client.jbb.64bit` file set first.

5. In the **PREVIEW only?** field (the remove operation does not occur), select **No**; press the ENTER key.

Installing the HP-UX Itanium 2 client

You can install the HP-UX Itanium 2 backup-archive client from the product installation media.

About this task

The following source packages are available on the installation media:

tsmcli/hp11ia64/gskcrypt64-8.x.x.x.hpux.ia64.tar.Z and **tsmcli/hp11ia64/gskssl64-8.x.x.x.hpux.ia64.tar.Z**

Contains the GSKit. If you have a previous version of the GSKit, uninstall it before you install the new version.

tsmcli/hp11ia64/TIVsmC

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

TIVsm.CLIENT

Contains the backup-archive client (command-line and Java GUI), administrative client (command-line), and the web client.

TIVsm.CLIENT_API64

Contains the 64-bit API.

tsmcli/hp11ia64/TIVsmCapi64

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

Default installation directories

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

- The backup-archive, web client, and administrative client (dsmadm) files are installed in the `/opt/tivoli/tsm/client/ba/bin` directory.
- The Tivoli Storage Manager API 64 files are installed in the `/opt/tivoli/tsm/client/api/bin64` directory.
- The sample system-options file, `dsm.sys.smp`, is placed in the installation directory.

To remove previous backup-archive client versions, log in as the root user and enter the following command:

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

If you installed additional languages in a Version 7.1.2 or earlier client, run the following command to remove them:

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

Replace *lang* with the appropriate language code from Table 10.

Table 10. HP-UX Itanium 2 client: Language codes for installation packages

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

This installation procedure can be used to install new distributions or updates from downloaded installation media. The downloaded files that you use to install the client might be compressed. Depending on the package file format, either copy or extract the files to disk and use these instructions to install the components.

You can download the appropriate package file from one of the following websites:

- Download the client package from Passport Advantage or Fix Central.
- For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.

Procedure

1. Log in as the root user.
2. Mount the volume that you are installing from.
3. To install GSKit: If you have a previous version of GSKit installed, remove it before you install the new version. Extract the contents of gskcrypt64-8.x.x.x.hpux.ia64.tar.Z and gskssl64-8.x.x.x.hpux.ia64.tar.Z to a directory on your hard disk. Enter the following commands to install the packages:

```
/usr/sbin/swinstall -x mount_all_filesystems=false -v -s `pwd`
/gskcrypt64 gskcrypt64
/usr/sbin/swinstall -x mount_all_filesystems=false -v -s `pwd`
/gskssl64 gskssl64
```

4. If you downloaded from FTP, go to the directory where the installable image is located. Enter the following command:

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

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

Results

To install only the API, omit the last TIVsm64 from the command, and mark only the file set that you want to install. Raw logical volume backups require that the API is installed. Mark CLIENT to install the backup-archive client and the API. Mark CLIENT_API64 to install only the API.

Related concepts:

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

Increasing the default limit of the data segment size

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

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

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

Uninstalling the HP-UX Itanium 2 client

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

About this task

Follow these steps to uninstall all of the packages that are related to Tivoli Storage Manager, including the command-line client, the GUI, the web GUI, and the administrative client components.

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

Procedure

1. Enter the following command to uninstall the backup-archive client:

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

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

Note: If one or more Tivoli Storage Manager language messages packages were installed in a Version 7.1.2 or earlier client, you must remove them before you can remove the API package. To remove the language messages packages, enter the following command as root user:

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

Where xx_XX is the language you want to remove.

Replace xx_XX with the following language codes: cs_CZ, de_DE, es_ES, it_IT, fr_FR, hu_HU, ja_JP, ko_KR, pl_PL, pt_BR, ru_RU, zh_CN and zh_TW

2. To remove the CLIENT_API file set, enter the following command:

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

3. To remove the Global Security Kit (GSKit), enter the following commands:

```
/usr/sbin/swremove -x mount_all_filesystems=false gskssl64  
/usr/sbin/swremoveswremove -x mount_all_filesystems=false gskcrypt64
```

What to do next

After you uninstall the HP-UX client, several empty directories remain in the file system, such as the following directories:

- The license directory (/opt/tivoli/tsm/license)
- One or more language directories (/opt/tivoli/tsm/client/ba/bin/xx_XX), where xx_XX represents one of the following language codes: cs_CZ, de_DE, es_ES, it_IT, fr_FR, hu_HU, ja_JP, ko_KR, pl_PL, pt_BR, ru_RU, zh_CN and zh_TW
- /opt/tivoli/tsm/client/ba/bin/cit
- /opt/tivoli/tsm/client/ba/bin/images
- /opt/tivoli/tsm/client/ba/bin/plugin

If you want to remove these empty directories, you can manually remove them.

Installing the backup-archive client on Linux on Power Systems

You can install the backup-archive client from the product installation media.

Before you begin

You must be logged in as root to install the product.

About this task

If you have Tivoli Storage Manager Version 6.2 (or an earlier version) installed, remove it (**rpm -e**) and any other dependent software programs before you install a newer version.

If you have Tivoli Storage Manager V6.3 (or newer) installed, you can use the rpm upgrade option (**rpm -U**) or the rpm freshen option (**rpm -F**) to upgrade the

existing software to a newer version. The **rpm -U** command can be used to install new packages or upgrade existing packages; **rpm -F** can update only packages that are already installed.

Stop any running client processes before you uninstall or upgrade the Tivoli Storage Manager API or backup-archive client. If you are running a V7.1.2 or earlier client, you must uninstall any language packages before you proceed with the upgrade.

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

Table 11. Package names, contents, and default directory

Package Name	Contents	Default directory
gskcrypt64-8.x.x.x.linux.ppc.rpm gskssl64-8.x.x.x.linux.ppc.rpm	64-bit Global Security Kit (GSKit) packages	/usr/local/ibm/gsk8
TIVsm-API64.ppc64.rpm	Application programming interface (API), which contains the Tivoli Storage Manager API shared libraries and samples.	/opt/tivoli/tsm/client/api/bin64
TIVsm-BA.ppc64.rpm	Backup-archive client (command-line and GUI), administrative client (dsmadm), and the web client.	/opt/tivoli/tsm/client/ba/bin This directory is considered to be the default installation directory for many backup-archive client files. The sample system-options file (dsm.sys.smp) is written to this directory. If the DSM_DIR environment variable is not set, the dsmc executable file, the resource files, and the dsm.sys file are stored in this directory. If the DSM_CONFIG environment variable is not set, the client user-options file must be in this directory. If you do not set the DSM_LOG environment variable, the backup-archive client writes messages to the dsmerror.log and dsmsched.log files in the current working directory.

Table 11. Package names, contents, and default directory (continued)

Package Name	Contents	Default directory
TIVsm-APIcit.ppc64.rpm TIVsm-BACit.ppc64.rpm	Optional. These files provide the Common Inventory Technology components that you can use to obtain information about the number of client and server devices that are connected to the system, and the utilization of processor value units (PVUs) by server devices. For more information about PVUs, see Estimating processor value units in the Tivoli Storage Manager server documentation.	APIcit is installed in /opt/tivoli/tsm/client/api/bin64/cit/ BACit is installed in /opt/tivoli/tsm/client/ba/bin/cit/
TIVsm-filepath-version-linuxOS.ppc64.rpm Where: <i>version</i> indicates the Tivoli Storage Manager version, and <i>linuxOS</i> indicates the Linux operating system that you are installing the filepath component on. TIVsm-filepath-source.tar.gz TIVsm-JBB.ppc64.rpm	Files that are needed for journal-based backups.	Filepath is installed in /opt/filepath JBB is installed in /opt/tivoli/tsm/client/ba/bin

This installation procedure can be used to install new distributions or updates from downloaded installation media. The downloaded files that you use to install the client might be compressed. Depending on the package file format, either copy or extract the files to disk and use these instructions to install the components.

You can download the appropriate package file from one of the following websites:

- Download the client package from Passport Advantage or Fix Central.
- For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.

Procedure

1. Mount the volume that you are installing from.
2. Change to the directory where the installation packages are stored.
3. Install the 64-bit GSKit packages. In this example, the "8.x.x.x" characters represent the GSKit version:
rpm -U gskcrypt64-8.x.x.x.linux.ppc.rpm gskssl64-8.x.x.x.linux.ppc.rpm
4. Install the Tivoli Storage Manager API, and optionally install the Common Inventory Technology package that is needed to support processor value unit (PVU) calculations.
 - a. Required: Install the API:
rpm -i TIVsm-API64.ppc64.rpm

- b. Optional: Install the Common Inventory Technology package that is used by the API. This package is dependent on the API so it must be installed after the API package is installed.

```
rpm -i TIVsm-APIcit.ppc64.rpm
```

Tip: If you are upgrading the API and the Common Inventory Technology package was previously installed, you must upgrade both the API and Common Inventory Technology packages. For example, you can run the following command:

```
rpm -U TIVsm-API64.ppc64.rpm TIVsm-APIcit.ppc64.rpm
```

If you need only the API installed, you can stop here. The rest of the steps in this procedure describe how to install the backup-archive client components and an optional client package that is needed only if you want the client to send PVU metrics to the server. Also described in subsequent steps are the installation of the packages that are needed if you want to run journal-based backups.

5. Install the backup-archive client, and optionally install the Common Inventory Technology package that is needed to support processor value unit (PVU) calculations.

- a. Install the backup-archive client components.

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

- b. Optional: Install the Common Inventory Technology package the client uses to send PVU metrics to the server. This package is dependent on the client package so it must be installed after the client package is installed.

```
rpm -i TIVsm-BACit.ppc64.rpm
```

6. Optional: If you want to use journal-based backups, install the packages that are needed for the filepath component and journal-based backups. Install the filepath component first and ensure that you install the filepath package that is appropriate for the operating system on the client computer. Specify one of the following package names for the filepath package name:
 - TIVsm-filepath-version-rhel59.ppc64.rpm, for Red Hat Enterprise Linux 5.9.
 - TIVsm-filepath-version-rhel64.ppc64.rpm, for Red Hat Enterprise Linux 6.4.
 - TIVsm-filepath-version-sles11sp2.ppc64.rpm, for SUSE Linux Enterprise Server 11, SP 2.

Use the `ls` command to display the names of the filepath packages that are included in this release. Then, use `rpm -i filepath_package_name.rpm` to install the appropriate filepath package.

If the filepath packages in the distribution do not apply to your Linux version, you must compile filepath from source code. Extract TIVsm-filepath-source.tar.gz and see the README file for compile and install instructions. The Linux Filepath kernel module is licensed pursuant to the terms of the GNU General Public License ("GPL").

Related concepts:

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

Uninstalling the backup-archive client on Linux on Power Systems

You can uninstall the Tivoli Storage Manager Linux on Power Systems client.

Before you begin

You must be logged in as root to uninstall the product. Uninstall the packages in the order shown.

Procedure

To uninstall a previously installed Tivoli Storage Manager client package, enter the following commands to remove the packages for journal-based backup, the filepath component, the backup-archive client, the API, and the IBM Global Security Kit (GSKit).

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

1. To uninstall the journal-based backup components only, remove both packages (journal-based backup and filepath). The TIVsm-JBB package is dependent on the filepath package. If you use two separate **rpm -e** commands to uninstall the components one at a time, uninstall the TIVsm-JBB package first.

```
rpm -e TIVsm-JBB TIVsm-filepath
```

2. Uninstall the backup-archive client packages:

- a. If you installed the optional TIVsmBAcit package, uninstall it before you uninstall the client package.

```
rpm -e TIVsm-BAcit
```

- b. Uninstall the backup-archive client package.

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

Note: If language packages are installed in a Version 7.1.2 or earlier client, you must remove them before you remove the API package. Enter the following command, and replace *xx_xx* with the language code for each additional language that you installed. For a list of language code identifiers, see Table 12.

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

Table 12. Language pack identifiers

Language	Language identifier
Czech	CS_CZ
French	FR_FR
German	DE_DE
Hungarian	HU_HU
Italian	IT_IT
Japanese	JA_JP
Korean	KO_KR
Polish	PL_PL
Portuguese	PT_BR
Russian	RU_RU
Spanish	ES_ES
Traditional Chinese (EUC)	ZH_CN
Traditional Chinese Big5	ZH_TW

3. Uninstall any products that are dependent on the API, such as IBM Tivoli Storage Manager for Databases and IBM Tivoli Storage Manager for Mail. Any API-dependent products must be uninstalled before you uninstall the API

package. If you uninstall an API-dependent product, you must reinstall it after you install a newer version of the backup-archive client and API packages. Consult the documentation of the dependent product to determine what you need to do to prevent data loss when you uninstall and reinstall the products.

- a. If you installed the optional API common inventory package (TIVsm-APIcit), uninstall it before you uninstall the API package. Use the following command to uninstall the package:

```
rpm -e TIVsm-APIcit
```

- b. Uninstall the API package by using the following command:

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

4. Uninstall GSKit by entering this command:

```
rpm -e gskcrypt64 gskssl64
```

Related tasks:

“Installing the backup-archive client on Linux on Power Systems” on page 17

Installing the backup-archive client on Linux on Power Systems (Little Endian)

You can install the backup-archive client from the product installation media.

Before you begin

You must be logged in as root user to install the product.

This installation procedure can be used to install new distributions or updates from downloaded installation media. The downloaded files that you use to install the client might be compressed. Depending on the package file format, either copy or extract the files to disk and use these instructions to install the components.

You can download the appropriate package file from one of the following websites:

- Download the client package from Passport Advantage or Fix Central.
- For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.

About this task

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

Table 13. Package names, contents, and default directory

Package Name	Contents	Default directory
gskcrypt64-8.x.x.x.linux.ppcle.rpm gskssl64-8.x.x.x.linux.ppcle.rpm	64-bit Global Security Kit (GSKit) packages	/usr/local/ibm/gsk8
TIVsm-API64.ppc64le.rpm	Application programming interface (API), which contains the shared libraries and samples for the Tivoli Storage Manager API.	/opt/tivoli/tsm/client/api/bin64

Table 13. Package names, contents, and default directory (continued)

Package Name	Contents	Default directory
TIVsm-BA.ppc64le.rpm	Backup-archive client (command-line and GUI), administrative client (dsmadm), and the web client.	/opt/tivoli/tsm/client/ba/bin This directory is typically the default installation directory for many backup-archive client files. The sample system-options file (dsm.sys.smp) is written to this directory. If you do not set the DSM_DIR environment variable, the dsmc executable file, the resource files, and the dsm.sys file are stored in this directory. If you do not set the DSM_CONFIG environment variable, the client user-options file must be in this directory. If you do not set the DSM_LOG environment variable, the backup-archive client writes messages to the dsmmerror.log and dsmsched.log files in the current working directory.
TIVsm-APIcit.ppc64le.rpm TIVsm-BAcit.ppc64le.rpm	These files provide the Common Inventory Technology components that you can use to obtain information about the number of client and server devices that are connected to the system, and the utilization of processor value units (PVUs) by server devices. These files are optional. For more information about PVUs, see Estimating processor value units in the Tivoli Storage Manager server documentation.	APIcit is installed in /opt/tivoli/tsm/client/api/bin64/cit BAcit is installed in /opt/tivoli/tsm/client/ba/bin/cit
TIVsm-filepath-version-linuxOS.ppc64le.rpm Where: <i>version</i> indicates the Tivoli Storage Manager version, and <i>linuxOS</i> indicates the Linux operating system that you are installing the filepath component on. TIVsm-filepath-source.tar.gz TIVsm-JBB.ppc64le.rpm	Files that are needed for journal-based backups.	The filepath and TIVsm-JBB packages are only required if you plan to use journal-based backups. Filepath is installed in /opt/filepath The TIVsm-JBB.ppc64le.rpm package is installed in /opt/tivoli/tsm/client/ba/bin.

Procedure

1. Mount the volume that you are installing the packages from.
2. Change to the directory where the installation packages are stored.

3. Install the 64-bit GSKit packages. In the following command example, the "8.x.x.x" characters represent the GSKit version:

```
rpm -U gskcrypt64-8.x.x.x.linux.ppc64le.rpm gskssl64-8.x.x.x.linux.ppc64le.rpm
```
4. Install the Tivoli Storage Manager API:

```
rpm -ivh TIVsm-API64.ppc64le.rpm
```
5. Optional: Install the Common Inventory Technology package that is used by the API. This package is dependent on the API so it must be installed after the API package is installed.

```
rpm -ivh TIVsm-APIcit.ppc64le.rpm
```

Tip: If you are upgrading the API and the Common Inventory Technology package was previously installed, you must upgrade both the API and Common Inventory Technology packages. For example, you can run the following command:

```
rpm -U TIVsm-API64.ppc64le.rpm TIVsm-APIcit.ppc64le.rpm
```

If you need only the API installed, you can stop here. The rest of the steps in this procedure describe how to install the backup-archive client components, and an optional client package that is needed only if you want the client to send PVU metrics to the server. Also described in subsequent steps are the installation of the packages that are needed if you want to run journal-based backups.

6. Install the backup-archive client components:

```
rpm -ivh TIVsm-BA.ppc64le.rpm
```
7. Optional: Install the Common Inventory Technology package that the client uses to send PVU metrics to the server. This package is dependent on the client package so it must be installed after the client package is installed.

```
rpm -ivh TIVsm-BAcit.ppc64le.rpm
```
8. Optional: If you want to use journal-based backups, install the packages that are needed for the filepath component and journal-based backups.
 - a. Install the filepath component first and ensure that you install the filepath package that is appropriate for the operating system on the client computer. Specify one of the following package names for the filepath package name:
 - TIVsm-filepath-7.1.8-0-rhel7.ppc64le.rpm, for Red Hat Enterprise Linux 7.1.
 - TIVsm-filepath-7.1.8-0-sles12.ppc64le.rpm, for SUSE Linux Enterprise Server 12.

Use the **ls** command to display the names of the filepath packages that are included in this release. Then, use `rpm -ivh filepath_package_name.rpm` to install the appropriate filepath package.

If the filepath packages in the distribution do not apply to your Linux version, you must compile filepath from source code. Extract TIVsm-filepath-source.tar.gz and see the README file for compile and install instructions. The Linux Filepath kernel module is licensed pursuant to the terms of the GNU General Public License ("GPL").

- b. Install the journal-based backup package:

```
rpm -ivh TIVsm-JBB.ppc64le.rpm
```

Related concepts:

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

Uninstalling the backup-archive client on Linux on Power Systems (Little Endian)

You can uninstall the Tivoli Storage Manager client on Linux on Power Systems (Little Endian).

Before you begin

You must be logged in as root user to uninstall the product. You must uninstall the packages in the order that is shown, otherwise the uninstallation fails.

Procedure

To uninstall the backup-archive client, enter the following commands to remove the packages for journal-based backup, the filepath component, the backup-archive client, the API, and the IBM Global Security Kit (GSKit).

Tip: The version number of the packages is not required.

1. To uninstall the journal-based backup components only, remove both packages (journal-based backup and filepath). The TIVsm-JBB package is dependent on the filepath package. If you use two separate **rpm -e** commands to uninstall the components one at a time, uninstall the TIVsm-JBB package first.

```
rpm -e TIVsm-JBB TIVsm-filepath
```
2. Uninstall the backup-archive client package:

```
rpm -e TIVsm-BA
```
3. Uninstall the backup-archive client packages:
 - a. If you installed the client common inventory package (TIVsmBACit), uninstall it:

```
rpm -e TIVsm-BACit
```
 - b. Uninstall the backup-archive client package:

```
rpm -e TIVsm-BA
```
4. Uninstall products that are dependent on the API, such as IBM Tivoli Storage Manager for Databases and IBM Tivoli Storage Manager for Mail. Any API-dependent products must be uninstalled before you uninstall the API package. If you uninstall an API-dependent product, you must reinstall it after you install a newer version of the backup-archive client and API packages. Follow the instructions of the API-dependent products to determine what you need to do to prevent data loss when you uninstall and reinstall the products.
5. Uninstall the API package by using the following command:

```
rpm -e TIVsm-API64
```
6. Uninstall the API packages:
 - a. If you installed the API common inventory package (TIVsm-APIcit), uninstall it:

```
rpm -e TIVsm-APIcit
```
 - b. Uninstall the API package by using the following command:

```
rpm -e TIVsm-API64
```
7. Uninstall GSKit by entering the following command:

```
rpm -e gskcrypt64 gskssl64
```

Related tasks:

“Installing the backup-archive client on Linux on Power Systems (Little Endian)” on page 22

Installing the backup-archive client on Ubuntu Linux on Power Systems (Little Endian)

You can install the backup-archive client from the product installation media.

Before you begin

You must be logged in as the root user to install the product.

This installation procedure can be used to install new distributions or updates from downloaded installation media. The downloaded files that you use to install the client might be compressed. Depending on the package file format, either copy or extract the files to disk and use these instructions to install the components.

You can download the appropriate package file from one of the following websites:

- Download the client package from Passport Advantage or Fix Central.
- For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.

About this task

The following installation packages are available on the installation media.

Table 14. Package names, contents, and default directory

Package Name	Contents	Default directory
gskcrypt64_8.x.x.x.ppc64el.deb gskssl64_8.x.x.x.ppc64el.deb	64-bit Global Security Kit (GSKit) packages	/usr/local/ibm/gsk8
tivsm-api64.ppc64el.deb	Application programming interface (API), which contains the shared library and samples for the Tivoli Storage Manager API.	/opt/tivoli/tsm/client/api/bin64
tivsm-ba.ppc64el.deb	Backup-archive client (command-line and GUI), administrative client (dsmadm), and the web client.	/opt/tivoli/tsm/client/ba/bin This directory is typically the default installation directory for many backup-archive client files. The sample system-options file (dsm.sys.smp) is written to this directory. If you do not set the DSM_DIR environment variable, the dsmc executable file, the resource files, and the dsm.sys file are stored in this directory. If you do not set the DSM_CONFIG environment variable, the client user-options file must be in this directory. If you do not set the DSM_LOG environment variable, the backup-archive writes messages to the dsmerror.log and dsmsched.log files in the current working directory.

Table 14. Package names, contents, and default directory (continued)

Package Name	Contents	Default directory
TIVsm-filepath-source.tar.gz	Files that are required for journal-based backups.	The TIVsm-filepath-source.tar.gz package is installed in the /opt/filepath directory.
tivsm-jbb.ppc64el.deb		The tivsm-jbb.ppc64el.rpm package is installed in the /opt/tivoli/tsm/client/ba/bin directory.

Procedure

1. Mount the volume that you are installing the packages from.
2. Change to the directory where the installation packages are stored.
3. Install the 64-bit GSKit packages. In the following command example, the "8.x.x.x" characters represent the GSKit version:

```
dpkg -i gskcrypt64_8.x.x.x.ppc64el.deb gskssl64_8.x.x.x.ppc64el.deb
```
4. Install the Tivoli Storage Manager API:

```
dpkg -i tivsm-api64.ppc64el.deb
```

If you need to install only the API, you can stop here. Use the following steps to install the backup-archive client and the packages that are required to run journal-based backups.
5. Install the backup-archive client:

```
dpkg -i tivsm-ba.ppc64el.deb
```
6. Optional: If you want to use journal-based backups, install the following packages:
 - a. Extract TIVsm-filepath-source.tar.gz and review the README file for instructions about how to compile and install the software. The Linux Filepath kernel module is licensed pursuant to the terms of the GNU General Public License ("GPL").
 - b. Install the journal-based backup package:

```
dpkg -i tivsm-jbb.ppc64el.deb
```

Related concepts:

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

Uninstalling the client on Ubuntu Linux on Power Systems (Little Endian)

You can uninstall the Tivoli Storage Manager backup-archive client on Ubuntu Linux on Power Systems (Little Endian).

Before you begin

You must be logged in as the root user to uninstall the product.

Requirement: You must uninstall the packages in the order that is shown, otherwise the uninstallation fails.

Procedure

To uninstall the backup-archive client, enter the following commands to remove the packages for journal-based backup, the backup-archive client, the API, and the

IBM Global Security Kit (GSKit). Instructions for uninstalling the filepath component are provided with the source code for filepath, when you obtain the software from IBM.

Tip: The version number of the packages is not required.

1. To uninstall only the journal-based backup components, remove both the `tivsm-jbb` and `filepath` packages. The `tivsm-jbb` package depends on the `filepath` package. Uninstall the `tivsm-jbb` package first.
 - a. `dpkg -r tivsm-jbb`
 - b. `dpkg -r TIVsm-filepath`
2. Uninstall the backup-archive client package:
`dpkg -r tivsm-ba`
3. Uninstall any products that depend on the API, such as Tivoli Storage Manager for Databases and Tivoli Storage Manager for Mail.
If you uninstall an API-dependent product, you must reinstall it after you install a newer version of the backup-archive client and API packages. Follow the instructions of the API-dependent products to determine what you need to do to prevent data loss when you uninstall and reinstall the products.
4. Uninstall the API package by issuing the following command:
`dpkg -r tivsm-api64`
5. Remove the GSKit packages:
`dpkg -r gskcrypt64 gskssl64`

Related tasks:

“Installing the backup-archive client on Ubuntu Linux on Power Systems (Little Endian)” on page 26

Installing the Linux x86_64 client

You can install the Linux x86_64 backup-archive client from the product installation media.

Before you begin

- You must be logged in as root to install the product.
- If you have Tivoli Storage Manager Version 6.2 (or an earlier version) installed, remove it (**rpm -e**) and any other dependent software programs before you install a newer version.
- If you have Tivoli Storage Manager V6.3 (or later) installed, you can use the rpm upgrade option (**rpm -U**) or the rpm freshen option (**rpm -F**) to upgrade the existing software to a newer version. The **rpm -U** command can be used to install new packages or upgrade existing packages only if you did not previously install any language packages. The **rpm -F** command can update only packages that are already installed.
- Stop any running client processes before you uninstall or upgrade the Tivoli Storage Manager API or backup-archive client.
- If any language packages are installed, you must uninstall them before you install or upgrade the Tivoli Storage Manager API or backup-archive client.

About this task

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

Table 15. Package names, contents, and default directory

Package Name	Contents	Default directory
gskcrypt64-8.x.x.x.linux.x86_64.rpm gskssl64-8.x.x.x.linux.x86_64.rpm	64-bit Global Security Kit (GSKit) packages	/usr/local/ibm/gsk8
TIVsm-API64.x86_64.rpm	Application programming interface (API), which contains the Tivoli Storage Manager API shared libraries and samples.	/opt/tivoli/tsm/client/api/bin64
TIVsm-BA.x86_64.rpm	Backup-archive client (command-line and GUI), administrative client (dsmadmc), and the web client.	/opt/tivoli/tsm/client/ba/bin This directory is considered to be the default installation directory for many backup-archive client files. The sample system-options file (<code>dsm.sys.smp</code>) is written to this directory. If the <code>DSM_DIR</code> environment variable is not set, the <code>dsmc</code> executable file, the resource files, and the <code>dsm.sys</code> file are stored in this directory. If <code>DSM_CONFIG</code> is not set, the client user-options file must be in this directory. If you do not define <code>DSM_LOG</code> , writes messages to the <code>dsmerror.log</code> and <code>dsmsched.log</code> files in the current working directory.
TIVsm-APIcit.x86_64.rpm TIVsm-BAcit.x86_64.rpm	Optional. These files provide the Common Inventory Technology components that you can use to obtain information about the number of client and server devices that are connected to the system, and the utilization of processor value units (PVUs) by server devices. For more information about PVUs, see Estimating processor value units in the Tivoli Storage Manager server documentation.	APIcit is installed in <code>/opt/tivoli/tsm/client/api/bin64/cit/</code> BAcit is installed in <code>/opt/tivoli/tsm/client/ba/bin/cit/</code>
TIVsm-filepath-version-linuxOS.x86_64.rpm where: <i>version</i> indicates the Tivoli Storage Manager version, and <i>linuxOS</i> indicates the Linux operating system that you are installing the filepath component on. TIVsm-filepath-source.tar.gz TIVsm-JBB.x86_64.rpm	Files needed to support journal-based backups.	Filepath is installed in <code>/opt/filepath</code> JBB is installed in <code>/opt/tivoli/tsm/client/ba/bin</code>

Table 15. Package names, contents, and default directory (continued)

Package Name	Contents	Default directory
TIVsm_BAhdw.x86_64.rpm	Provides support for snapshot incremental backup for NetAPP and N-Series file servers.	/opt/tivoli/tsm/client/ba/bin/plugins

This installation procedure can be used to install new distributions or updates from downloaded installation media. The downloaded files that you use to install the client might be compressed. Depending on the package file format, either copy or extract the files to disk and use these instructions to install the components.

You can download the appropriate package file from one of the following websites:

- Download the client package from Passport Advantage or Fix Central.
- For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.

Procedure

To install the Linux x86_64 backup-archive client, complete the following steps:

1. Mount the volume that you are installing from.
2. Change to the directory where the installation packages are stored.
3. Install the 64-bit GSKIT packages. In this example, the "8.x.x.x" characters represent the GSKIT version on the installation media:

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

4. Install the Tivoli Storage Manager API, and optionally install the Common Inventory Technology package that is necessary to support processor value unit (PVU) calculations.

- a. Required: Install the API:

```
rpm -i TIVsm-API64.x86_64.rpm
```

- b. Optional: Install the Common Inventory Technology package that is used by the API. This package depends on the API so it must be installed after the API package is installed.

```
rpm -i TIVsm-APIcit.x86_64.rpm
```

Tip: If you are upgrading the API and the Common Inventory Technology package was previously installed, you must upgrade both the API and Common Inventory Technology packages. For example, you can run the following command:

```
rpm -U TIVsm-API64.x86_64.rpm TIVsm-APIcit.x86_64.rpm
```

If you need only the API installed, you can stop here. The rest of the steps in this procedure describe how to install the backup-archive client components and an optional client package that is needed only if you want the client to send PVU metrics to the server. Also described in subsequent steps are the installation of the packages that are needed if you want to perform journal-based backups.

5. Install the backup-archive client, and optionally install the Common Inventory Technology package that is necessary to support processor value unit (PVU) calculations.
 - a. Install the backup-archive client components.

```
rpm -i TIVsm-BA.x86_64.rpm
```

- b. Optional: Install the Common Inventory Technology package the client uses to send PVU metrics to the server. This package depends on the client package so it must be installed after the client package is installed.

```
rpm -i TIVsm-BAcit.x86_64.rpm
```

6. Optional: If you want to use journal-based backups, install the packages that are necessary to support the filepath component and journal-based backups. Install the filepath component first and ensure that you install the filepath package that is appropriate for the operating system on the client computer. Specify one of the following package names for the filepath package name:

- `TIVsm-filepath-version-rhel59.x86_64.rpm`, for Red Hat Enterprise Linux 5.9.
- `TIVsm-filepath-version-rhel64.x86_64.rpm`, for Red Hat Enterprise Linux 6.4
- `TIVsm-filepath-version-sles11sp2.x86_64.rpm`, for SUSE Linux Enterprise Server 11, SP 2.

Use the `ls` command to display the names of the filepath packages that are included in this release. Then, use `rpm -i filepath_package_name.rpm` to install the appropriate filepath package.

If the filepath packages in the distribution do not apply to your Linux version, you must compile filepath from source code. Extract `TIVsm-filepath-source.tar.gz` and see the README file for compile and install instructions. The Linux Filepath kernel module is licensed pursuant to the terms of the GNU General Public License ("GPL").

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

```
rpm -i TIVsm-BAhdw.x86_64.rpm
```

Related concepts:

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

Uninstalling the Linux x86_64 client

You can use the following procedure to uninstall the Tivoli Storage Manager Linux x86_64 client.

Before you begin

You must be logged in as root to uninstall the product. Uninstall the packages in the order shown.

Procedure

To uninstall a previously installed Tivoli Storage Manager client package, enter the following commands to remove the packages for journal-based backup, the filepath component, the backup-archive client, the API, and the IBM Global Security Kit (GSKit).

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

1. To uninstall the journal-based backup components only, remove both packages (journal-based backup and filepath). The `TIVsm-JBB` package depends on the filepath package. If you use two separate `rpm -e` commands to uninstall the components one at a time, uninstall the `TIVsm-JBB` package first.

```
rpm -e TIVsm-JBB TIVsm-filepath
```

2. Uninstall the backup-archive client packages:
 - a. If you installed the optional TIVsm-BACit package, uninstall it before you uninstall the client:


```
rpm -e TIVsm-BACit
```
 - b. Uninstall the backup-archive client.


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

Note: If language packages are installed in a Version 7.1.2 or earlier client, you must remove them before you remove the API package. Enter the following command, and replace *xx_xx* with the language code for each additional language that you installed. For a list of language code identifiers, see Table 16.

```
rpm -e TIVsm-msg.xx_xx
```

Table 16. Language pack identifiers

Language	Language identifier
Czech	CS_CZ
French	FR_FR
German	DE_DE
Hungarian	HU_HU
Italian	IT_IT
Japanese	JA_JP
Korean	KO_KR
Polish	PL_PL
Portuguese	PT_BR
Russian	RU_RU
Spanish	ES_ES
Traditional Chinese (EUC)	ZH_CN
Traditional Chinese Big5	ZH_TW

3. Uninstall any products that depend on the API, such as IBM Tivoli Storage Manager for Databases and IBM Tivoli Storage Manager for Mail. Any API-dependent products must be uninstalled before you uninstall the API package. If you uninstall an API-dependent product, you must reinstall it after you install a newer version of the backup-archive client and API packages. Consult the documentation of the dependent product to determine what you need to do to prevent data loss when you uninstall and reinstall the products.
 - a. If you installed the optional API common inventory package (TIVsm-APIcit), uninstall it before you uninstall the API package. Use the following command to uninstall the package:


```
rpm -e TIVsm-APIcit
```
 - b. Uninstall the API package by using the following command:


```
rpm -e TIVsm-API64
```
4. To remove the GSKit 64-bit package, enter the following command:


```
rpm -e gskcrypt64 gskss164
```

Related tasks:

“Installing the Linux x86_64 client” on page 28

Installing the Ubuntu Linux x86_64 client

You can install the Ubuntu Linux 64-bit backup-archive client from the product installation media.

About this task

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

Table 17. Package names, contents, and default directory

Package Name	Contents	Default directory
gskcrypt64_8.0-50.40.linux.x86_64.deb gskssl64_8.0-50.40.linux.x86_64.deb	64-bit Global Security Kit (GSKit) packages	/usr/local/ibm/gsk8
tivsm-api64.amd64.deb	Application programming interface (API), which contains the Tivoli Storage Manager API shared libraries and samples.	/opt/tivoli/tsm/client/api/bin64
tivsm-ba.amd64.deb	Backup-archive client (command-line and GUI), administrative client (dsmadm), and the web client.	/opt/tivoli/tsm/client/ba/bin This directory is considered to be the default installation directory for many backup-archive client files. The sample system-options file (dsm.sys.smp) is written to this directory. If the DSM_DIR environment variable is not set, the dsmc executable file, the resource files, and the dsm.sys file are stored in this directory. If DSM_CONFIG is not set, the client user-options file must be in this directory. If you do not define DSM_LOG, writes messages to the dsmerror.log and dsmsched.log files in the current working directory.

Table 17. Package names, contents, and default directory (continued)

Package Name	Contents	Default directory
tivsm-apicit.amd64.deb tivsm-bacit.amd64.deb	Optional. These files provide the Common Inventory Technology components that you can use to obtain information about the number of client and server devices that are connected to the system, and the utilization of processor value units (PVUs) by server devices. For more information about PVUs, see Estimating processor value units in the Tivoli Storage Manager server documentation.	APicit is installed in /opt/tivoli/tsm/client/api/bin64/cit BAcit is installed in /opt/tivoli/tsm/client/ba/bin/cit
tivsm-filepath-source.tar.gz tivsm-jbb.amd64.deb	Files needed to support journal-based backups.	The filepath and tivsm-jbb packages are only required if you plan to use journal-based backups. The tivsm-jbb.x86_64.deb package is installed in /opt/tivoli/tsm/client/ba/bin.
tivsm-bahdw.amd64.deb	Provides support for snapshot incremental backup for NetAPP and N-Series file servers.	/opt/tivoli/tsm/client/ba/bin/plugins

This installation procedure can be used to install new distributions or updates from downloaded installation media. The downloaded files that you use to install the client might be compressed. Depending on the package file format, either copy or extract the files to disk and use these instructions to install the components.

You can download the appropriate package file from one of the following websites:

- Download the client package from Passport Advantage or Fix Central.
- For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.

Procedure

To install the Ubuntu Linux x86_64 backup-archive client, complete the following steps.

1. Mount the volume that you are installing from.
2. Change to the directory where the installation packages are stored.
3. Install the 64-bit GSKit packages.

```
sudo dpkg -i gskcrypt64_8.0-50.40.linux.x86_64.deb gskssl64_8.0-50.40.linux.x86_64.deb
```
4. Install the Tivoli Storage Manager API, and optionally install the Common Inventory Technology package that is necessary to support processor value unit (PVU) calculations.

- a. Required: Install the API:

```
sudo dpkg -i tivsm-api64.amd64.deb
```

- b. Optional: Install the Common Inventory Technology package that is used by the API. This package depends on the API so it must be installed after the API package is installed.

```
sudo dpkg -i tivsm-apicit.amd64.deb
```

If you need only the API installed, you can stop here. The rest of the steps in this procedure describe how to install the backup-archive client components and an optional client package that is needed only if you want the client to send PVU metrics to the server. Also described in subsequent steps are the installation of the packages that are needed if you want to perform journal-based backups.

5. Install the backup-archive client, and optionally install the Common Inventory Technology package that is necessary to support processor value unit (PVU) calculations.

- a. Install the backup-archive client components.

```
sudo dpkg -i tivsm-ba.amd64.deb
```

- b. Optional: Install the Common Inventory Technology package that the client uses to send PVU metrics to the server. This package depends on the client package so it must be installed after the client package is installed.

```
sudo dpkg -i tivsm-bacit.amd64.deb
```

6. Optional: Complete this step only if you plan to use journal-based backups.

- a. Extract `tivsm-filepath-source.tar.gz` and see the README file for compile and install instructions. The filepath kernel module is licensed pursuant to the terms of the GNU General Public License ("GPL").

- b. Install the journal-based backup package: `dpkg -i tivsm-jbb.amd64.deb`.

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

```
sudo dpkg -i tivsm-bahdw.amd64.deb
```

Related concepts:

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

Uninstalling the Ubuntu Linux x86_64 client

Use the following procedure to uninstall the Tivoli Storage Manager Ubuntu Linux 64-bit client.

Procedure

To uninstall a previously installed Tivoli Storage Manager client package, enter the following commands to remove the packages for journal-based backup, the backup-archive client, the API, and the IBM Global Security Kit (GSKit). Instructions to uninstall the filepath component are provided with the source code for filepath, when you obtain the software from IBM.

1. To uninstall only the journal-based backup components, remove both the `tivsm-jbb` and the filepath component. The `tivsm-jbb` package depends on the filepath package. Uninstall the `tivsm-jbb` package first.

- a. `sudo dpkg -r tivsm-jbb`

- b. `sudo dpkg -r tivsm-filepath`

2. Uninstall the backup-archive client packages:

- a. If you installed the optional `tivsm-bacit` package, uninstall it before you uninstall the client:

```
sudo dpkg -r tivsm-bacit
```

- b. Uninstall the backup-archive client.

```
sudo dpkg -r tivsm-ba
```

Note: If language packages are installed in a Version 7.1.2 or earlier client, you must remove them before you remove the API package. Enter the following command, and replace *xx-xx* with the language code for each additional language that you installed. For a list of language code identifiers, see Table 18.

```
dpkg -r tivsm-msg.xx-xx
```

Table 18. Language pack identifiers

Language	Language identifier
Czech	cs-cz
French	fr-fr
German	de-de
Hungarian	hu-hu
Italian	it-it
Japanese	ja-jp
Korean	ko-kr
Polish	pl-pl
Portuguese	pt-br
Russian	ru-ru
Spanish	es-es
Traditional Chinese (EUC)	zh-cn
Traditional Chinese Big5	zh-tw

3. Uninstall any products that depend on the API, such as IBM Tivoli Storage Manager Data Protection products. Any API-dependent products must be uninstalled before you uninstall the API package. If you uninstall an API-dependent product, you must reinstall it after you install a newer version of the backup-archive client and API packages. Consult the documentation of the dependent product to determine what you need to do to prevent data loss when you uninstall and reinstall the products.

- a. If you installed the optional API common inventory package (*tivsm-apicit*), uninstall it before you uninstall the API package. Use the following command to uninstall the package:

```
sudo dpkg -r tivsm-apicit
```

- b. Uninstall the API package by using the following command:

```
sudo dpkg -r tivsm-api64
```

4. To remove the GSKit 64-bit packages, enter the following command:

```
sudo dpkg -r gskcrypt64 gskssl64
```

Related tasks:

“Installing the Ubuntu Linux x86_64 client” on page 33

Installing the Linux on System z client

You can install the Linux on System z backup-archive client from the product installation media.

Before you begin

You must be logged in as root to install the product.

About this task

If you have Tivoli Storage Manager Version 6.2 (or an earlier version) installed, remove it (**rpm -e**) and any other dependent software programs before you install a newer version.

If you have Tivoli Storage Manager V6.3 (or newer) installed, you can use the rpm upgrade option (**rpm -U**) or the rpm freshen option (**rpm -F**) to upgrade the existing software to a newer version. The **rpm -U** command can be used to install new packages or upgrade existing packages; **rpm -F** can update only packages that are already installed.

Stop any running client processes before you uninstall or upgrade the Tivoli Storage Manager API or backup-archive client. If you are running a V7.1.2 or earlier client, you must uninstall any language packages before you proceed with the upgrade.

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

Table 19. Package names, contents, and default directory

Package Name	Contents	Default directory
gskcrypt64-8.x.x.x.linux.s390x.rpm gskssl64-8.x.x.x.linux.s390x.rpm	64-bit Global Security Kit (GSKit) packages	/usr/local/ibm/gsk8
TIVsm-API64.s390x.rpm	Application programming interface (API), which contains the Tivoli Storage Manager API shared libraries and samples.	/opt/tivoli/tsm/client/api/bin64
TIVsm-BA.s390x.rpm	Backup-archive client (command-line and GUI), administrative client (dsmadmc), and the web client.	/opt/tivoli/tsm/client/ba This directory is considered to be the default installation directory for many backup-archive client files. The sample system-options file (dsm.sys.smp) is written to this directory. If the DSM_DIR environment variable is not set, the dsmc executable file, the resource files, and the dsm.sys file are stored in this directory. If DSM_CONFIG is not set, the client user-options file must be in this directory. If you do not define DSM_LOG, the backup-archive client writes messages to the dsmerror.log and dsmsched.log files in the current working directory.

Table 19. Package names, contents, and default directory (continued)

Package Name	Contents	Default directory
TIVsm-APIcit.s390x.rpm TIVsm-BACit.s390x.rpm	Optional. These files provide the Common Inventory Technology components that you can use to obtain information about the number of client and server devices that are connected to the system, and the utilization of processor value units (PVUs) by server devices. For more information about PVUs, see Estimating processor value units in the Tivoli Storage Manager server documentation.	APIcit is installed in /opt/tivoli/tsm/client/api/bin64/cit BACit is installed in /opt/tivoli/tsm/client/ba/bin/cit
TIVsm-filepath-version-linuxOS.s390x.rpm Where: <i>version</i> indicates the Tivoli Storage Manager version, and <i>linuxOS</i> indicates the Linux operating system that you are installing the filepath component on. TIVsm-filepath-source.tar.gz TIVsm-JBB.s390x.rpm	Files needed to support journal-based backups.	Filepath is installed in /opt/filepath JBB is installed in /opt/tivoli/tsm/client/ba/bin

This installation procedure can be used to install new distributions or updates from downloaded installation media. The downloaded files that you use to install the client might be compressed. Depending on the package file format, either copy or extract the files to disk and use these instructions to install the components.

You can download the appropriate package file from one of the following websites:

- Download the client package from Passport Advantage or Fix Central.
- For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.

Procedure

1. Mount the volume that you are installing from.
2. Change to the directory where the packages are stored.
3. Install the 64-bit GSKit packages. In this example, the "8.x.x.x" characters represent the GSKit version on the installation media:
rpm -U gskcrypt64-8.x.x.x.linux.s390x.rpm gskssl64-8.x.x.x.linux.s390x.rpm
4. Install the Tivoli Storage Manager API, and optionally install the Common Inventory Technology package that is needed to support processor value unit (PVU) calculations.
 - a. Required: Install the API:
rpm -i TIVsm-API64.s390x.rpm

- b. Optional: Install the Common Inventory Technology package that is used by the API. This package is dependent on the API so it must be installed after the API package is installed.

```
rpm -i TIVsm-APIcit.s390x.rpm
```

Tip: If you are upgrading the API and the Common Inventory Technology package was previously installed, you must upgrade both the API and Common Inventory Technology packages. For example, you can run the following command:

```
rpm -U TIVsm-API64.s390x.rpm TIVsm-APIcit.s390x.rpm
```

If you need only the API installed, you can stop here. The rest of the steps in this procedure describe how to install the backup-archive client components and an optional client package that is needed only if you want the client to send PVU metrics to the server. Also described in subsequent steps are the installation of the packages that are needed if you want to perform journal-based backups.

5. Install the backup-archive client, and optionally install the Common Inventory Technology package that is needed to support processor value unit (PVU) calculations.

- a. Install the backup-archive client components.

```
rpm -i TIVsm-BA.s390x.rpm
```

- b. Optional: Install the Common Inventory Technology package the client uses to send PVU metrics to the server. This package is dependent on the client package so it must be installed after the client package is installed.

```
rpm -i TIVsm-BAcit.s390x.rpm
```

6. Optional: If you want to use journal-based backups, install the packages that are needed to support the filepath component and journal-based backups. Install the filepath component first and ensure that you install the filepath package that is appropriate for the operating system on the client computer. Specify one of the following package names for the filepath package name.
 - `TIVsm-filepath-version-rhel64.s390x.rpm`, for Red Hat Enterprise Linux 6.4.
 - `TIVsm-filepath-version-rhel70.s390x.rpm`, for Red Hat Enterprise Linux 7.0.
 - `TIVsm-filepath-version-sles11sp2.s390x.rpm`, for SUSE Linux Enterprise Server 11, SP 2.

Use the `ls` command to display the names of the filepath packages that are included in this release. Then, use `rpm -i filepath_package_name.rpm` to install the appropriate filepath package.

If you are running GPFS™ for Linux on z Systems, you must specify the filepath package name for either Red Hat Enterprise Linux 7.0 or SUSE Linux Enterprise Server 11, SP 2.

If the filepath packages in the distribution do not apply to your Linux version, you must compile filepath from source code. Extract `TIVsm-filepath-source.tar.gz` and see the README file for compile and install instructions. The Linux Filepath kernel module is licensed pursuant to the terms of the GNU General Public License ("GPL").

Related concepts:

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

Uninstalling the Linux on System z client

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

Before you begin

You must be logged in as root to install the product. Uninstall the packages in the order shown.

About this task

To uninstall a previously installed Tivoli Storage Manager client package, enter the following commands to remove the packages for journal-based backup, the filepath component, the backup-archive client, the API, and the IBM Global Security Kit (GSKit).

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

Procedure

1. To uninstall the journal-based backup components only, remove both packages (journal-based backup and filepath). The TIVsm-JBB package is dependent on the filepath package. If you use two separate **rpm -e** commands to uninstall the components one at a time, uninstall the TIVsm-JBB package first.

```
rpm -e TIVsm-JBB TIVsm-filepath
```

2. Uninstall the backup-archive client packages:

- a. If you installed the optional TIVsm-BACit package, uninstall it before you uninstall the client:

```
rpm -e TIVsm-BACit
```

- b. Uninstall the backup-archive client.

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

Note: If language packages are installed in a Version 7.1.2 or earlier client, you must remove them before you remove the API package. Enter the following command, and replace *xx_xx* with the language code for each additional language that you installed. For a list of language code identifiers, see Table 20.

```
rpm -e TIVsm-msg.xx_xx
```

Table 20. Language pack identifiers

Language	Language identifier
Czech	CS_CZ
French	FR_FR
German	DE_DE
Hungarian	HU_HU
Italian	IT_IT
Japanese	JA_JP
Korean	KO_KR
Polish	PL_PL
Portuguese	PT_BR
Russian	RU_RU
Spanish	ES_ES
Traditional Chinese (EUC)	ZH_CN
Traditional Chinese Big5	ZH_TW

3. Uninstall any products that are dependent on the API, such as IBM Tivoli Storage Manager for Databases and IBM Tivoli Storage Manager for Mail. Any API-dependent products must be uninstalled before you uninstall the API package. If you uninstall an API-dependent product, you must reinstall it after you install a newer version of the backup-archive client and API packages. Consult the documentation of the dependent product to determine what you need to do to prevent data loss when you uninstall and reinstall the products.
 - a. If you installed the optional API common inventory package (TIVsm-APIcit), uninstall it before you uninstall the API package. Use the following command to uninstall the package:

```
rpm -e TIVsm-APIcit
```
 - b. Uninstall the API package by using the following command:

```
rpm -e TIVsm-API64
```
4. To remove the GSKit 64-bit package, enter the following command:

```
rpm -e gskcrypt64 gskssl64
```

Related tasks:

“Installing the Linux on System z client” on page 36

Installing the Mac OS X client

You can install the IBM Tivoli Storage Manager Mac OS X backup-archive client from the product installation media.

Before you begin

You must be a system administrator to install the backup-archive client.

About this task

This installation procedure can be used to install new distributions or updates from downloaded installation media. The downloaded files that you use to install the client might be compressed. Depending on the package file format, either copy or extract the files to disk and use these instructions to install the components.

You can download the appropriate package file from one of the following websites:

- Download the client package from Passport Advantage or Fix Central.
- For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.

For MAC OS X clients, you can use an installation wizard that prompts you for information as the product is installed, or you can also install the client from the command line. When you install the client by using the command-line installation procedure, the installation runs without user interaction. The command-line procedure is useful if you want to script the installation and run it on many nodes, or if you must install the software on a system that does not have a monitor.

Procedure

Select an installation method and install the client. Use either the installation wizard method or install the client from the command line.

Installation method	Procedure
Installation wizard	1. Double-click the 7.1.8.0-TIV-TSMBAC-Mac.dmg file to mount the disk image. 2. Double-click the Tivoli Storage Manager installation package icon and follow the prompts to complete the installation.
Command line	1. Change directories to where the Tivoli Storage Manager installer is located. 2. Install the custom installation package with the following command: <pre>/usr/sbin/installer -pkg "/Volumes/Tivoli Storage Manager/ Tivoli Storage Manager.pkg" -target /</pre>

What to do next

A sample client system options file, called `dsm.sys.smp`, is created in the installation directory. You can copy this file and modify it to create the client systems options file for your node. The default name for the client systems option file is `dsm.sys`.

After you install the client, you might need to set environment variables before you use it. For more information about setting environment variables, see “Set processing environment variables” on page 57.

Uninstalling the Mac OS X client

You can uninstall the Tivoli Storage Manager Mac OS X client if you no longer need it.

Before you begin

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

About this task

You can use a shell script to uninstall the backup-archive client. The shell script name is `uninstall.sh` and it is in the default installation directory, which is `/Library/Application Support/tivoli/tsm/client/ba/bin`. Use the **sudo** command to run the script.

Alternately, you can complete the following steps instead of using the script:

Procedure

1. Move the following folders to the trash:
 - `/Applications/Tivoli Storage Manager`
 - `/Library/Application Support/tivoli`
2. Remove the following symbolic links:
 - `/usr/bin/dsmc`

- /usr/bin/dsmcad
- /usr/bin/dsmadm
- /usr/bin/dsmtrace
- /usr/bin/dsmagent
- /usr/lib/libxmlutil-6.2.0.dylib
- /usr/lib/libtsm620xerces-c1_6_0.dylib

3. Optional: Remove the log files and options files if you do not want to preserve them. The uninstall process leaves them on disk so your settings are retained in case you reinstall the product later.

The backup-archive client might have created log files in these locations:

- a. /Library/Logs/tivoli
- b. ~/Library/Logs/tivoli

The client option files (dsm.opt and dsm.sys) are typically saved in the following locations:

- a. /Library/Preferences/Tivoli Storage Manager
- b. ~/Library/Preferences/Tivoli Storage Manager

Installing the Oracle Solaris client

You can install the Tivoli Storage Manager Oracle Solaris x86_64 backup-archive client from the product installation media.

About this task

If a previous version of the backup-archive client is installed, remove it before you install a new version. For information about removing previous Solaris client packages, see “Uninstalling the Oracle Solaris client” on page 45.

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

Table 21. Installation package names and descriptions

Package	Package Name	Package Description
IBM Global Security Kit (GSKit) 64 bit	gsk8cry64.pkg and gsk8ssl64.pkg	Contains the IBM GSKit that provides Secure Sockets Layer (SSL) 64-bit data encryption between the Tivoli Storage Manager client and server.
Tivoli Storage Manager application programming interface (API)	TIVsmCapi.pkg	Contains the Tivoli Storage Manager 64-bit API shared library and samples.

Table 21. Installation package names and descriptions (continued)

Package	Package Name	Package Description
Backup-archive client	TIVsmCba.pkg	Contains the following 64-bit components: • Backup-archive client (command-line and GUI) • Administrative client (command-line) • Web backup-archive client Note: 1. TCP/IP and Shared memory are supported as communication methods. 2. The web client is a part of the backup-archive client package and cannot be installed without it.

This installation procedure can be used to install new distributions or updates from downloaded installation media. The downloaded files that you use to install the client might be compressed. Depending on the package file format, either copy or extract the files to disk and use these instructions to install the components.

You can download the appropriate package file from one of the following websites:

- Download the client package from Passport Advantage or Fix Central.
- For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.

Install these packages in the order shown; some packages depend on the presence of others. For example, GSKit is a prerequisite of the API, and the API is a prerequisite of the backup-archive client package.

Procedure

1. Log in as the root user.
2. Mount the volume that you are installing from.
3. Change to the directory where the packages are stored.
4. The IBM GSKit; it is a prerequisite of the Tivoli Storage Manager API package. Install GSKit by using the following commands:

```
pkgadd -n -a ./tsmadmin -d ./gsk8cry64.pkg gsk8cry64
pkgadd -n -a ./tsmadmin -d ./gsk8ssl64.pkg gsk8ssl64
```

Note: On Solaris 10, these commands install the 64-bit GSKit in the global zone and in all running non-global zones. To install the client in a sparse-root, non-global zone only, GSKit must first be installed in the global zone. On Solaris 11, the packages are only installed in the zone where these commands are run.

5. Use the following command to install the Tivoli Storage Manager API:

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

Note: On Solaris 10, this command installs the Tivoli Storage Manager 64-bit API in the global zone and in all running non-global zones. If you want to

install it in the global zone only, use the **-G** parameter of the **pkgadd** command. On Solaris 11, the API is only installed in the zone where this command is run.

6. Use the following command to install the backup-archive client:

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

Note: On Solaris 10, this command installs the backup-archive client components in the global zone and in all running non-global zones. If you want to install them in the global zone only, use the **-G** parameter of the **pkgadd** command. On Solaris 11, the client components must be only installed in the zone where this command is run.

Results

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

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

Related concepts:

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

Uninstalling the Oracle Solaris client

You can uninstall all the packages that are related to Tivoli Storage Manager Oracle Solaris client, including the command-line, GUI, web GUI, and administrative client components.

About this task

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

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

Procedure

1. Enter the following command to uninstall the backup-archive client:

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

This command uninstalls all of the components of the backup-archive client (command-line, GUI, web client, and the administrative client). You cannot uninstall individual components of this package (for example, the command-line client).

Note: If one or more language messages packages are installed in Version 7.1.2 or earlier clients, remove them before you remove the API package. Enter the following command as the root user:

```
pkgrm -n -a ./tsmadmin TIVsmC1Cs TIVsmC1De TIVsmC1Es TIVsmC1Fr \
TIVsmC1Hu TIVsmC1It TIVsmC1Ja TIVsmC1Ko \
TIVsmC1Pl TIVsmC1Pt TIVsmC1Ru TIVsmC1Sc TIVsmC1Tc
```

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

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

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

3. Enter the following commands to uninstall the GSKit:

```
pkgrm -n -a ./tsmadmin gsk8ssl64
```

```
pkgrm -n -a ./tsmadmin gsk8cry64
```

Software updates

Software updates might periodically be made available by IBM for download.

For the latest information, updates, and maintenance fixes, go to the support website: http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager.

Installing the client management service to collect diagnostic information

You can install Tivoli Storage Manager client management services to collect diagnostic information about the backup-archive client. The client management service makes the information available to the Tivoli Storage Manager Operations Center for basic monitoring capability.

About this task

After you install the backup-archive client, install the client management service on the same computer so that the Tivoli Storage Manager server administrator can view diagnostic information from the Operations Center.

The client management service is available for installation on Linux backup-archive client systems.

For installation instructions and more information about the client management service, see Collecting diagnostic information with Tivoli Storage Manager client management services.

Chapter 2. Configure the Tivoli Storage Manager client

After installing the backup-archive client, you must configure it before performing any operations.

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

Some configuration tasks are required, while other tasks are optional. The following configuration tasks are required:

- “Creating and modifying the client system-options file” on page 51
- “Register your workstation with a server” on page 113

The following configuration tasks are optional:

- “Creating a default client-user options file” on page 53
- “Creating a customized client user-options file” on page 55
- “Environment variables” on page 55
- “Configuring the web client on AIX, HP-UX, Linux, Mac, and Solaris systems” on page 60
- “Configuring the scheduler” on page 61
- “Creating an include-exclude list” on page 115
- Configuring parallel backups of VMware virtual machines. See “Parallel backups of virtual machines” on page 205

UNIX and Linux client root and authorized user tasks

An authorized user is any non-root user who has read and write access to the stored password (`TSM.sth` file), or anyone who knows the password and enters it interactively. Authorized users use the `passworddir` option to define the directory where their copy of the password file is saved.

Tasks for root users and authorized users shows the tasks that can and cannot be performed by the root user, authorized users, and other users.

Table 22. Tasks for root users and authorized users

Task	Root user	Authorized user
Log on to a Tivoli Storage Manager server, using an LDAP server to authenticate credentials.	Yes	Yes
Register new nodes with the Tivoli Storage Manager server (if registration is set to open on the server).	Yes	Yes
Set or re-create the Tivoli Storage Manager password for client workstations	Yes	Yes

Table 22. Tasks for root users and authorized users (continued)

Task	Root user	Authorized user
Backup	Yes Note: The Tivoli Storage Manager administrator can specify an option on either the Register Node or Update Node commands to specify who is allowed to back up data for a node. Setting BACKUPINITiation to root restricts backups so that only root or authorized users can back up files on a node. Setting BACKUPINITiation to all allows any user to back up data on a node. For information about these commands and options, see the Tivoli Storage Manager server documentation.	Yes, if you have read permission, regardless of ownership
Restore	Yes; when restoring to a new location or the same location, file permission and ownership are preserved	Yes; however, the operating system prevents writing to the same location if the file has read only permission. When restoring to the same location, file permissions and ownership are preserved. When restoring to a different location, the permissions of the restored file are preserved but the ownership changed to the current user.
Archive	Yes	Yes, if you have read permission, regardless of ownership
Retrieve	Yes. When retrieving to a new location or to the same location, file permissions and ownership are preserved.	Yes. However, the operating system prevents writing to the same location if the file has read only permission. Ownership of all retrieved objects is changed to the current user.
Client scheduler	Yes	Yes, if not using the Client Acceptor Daemon. You must be root to manage the Client Acceptor Daemon. A non-root authorized user can use the scheduler (dsmc sched).
Grant user access to files on the Tivoli Storage Manager server	Yes	Yes
Delete Tivoli Storage Manager server file spaces	Yes, if the node is granted backup or archive delete authority by a Tivoli Storage Manager server administrator	Yes, if the node is granted backup or archive delete authority by a Tivoli Storage Manager server administrator

On Mac OS X systems, a system administrator is any user that is allowed to administer the system. You can check your account type using the **System Preferences > Accounts** tool. System Administrators have an account type of **Admin**.

The system administrator is responsible for configuring the backup-archive client so non-administrators can manage their own data. Non-administrators (or non-authorized users) meet the following criteria:

- They do not have a user ID of 0. They are not the root user.
- They have a user account that has not been configured as a system administrator.

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

Table 23. Mac OS X authorization tools and associated Tivoli Storage Manager applications

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

Enable non-root users to manage their own data

To enable non-root users to use the backup-archive client to manage their own data, the system administrator must complete steps in addition to the normal configuration steps to setup first-time Authorized users for non-root users.

In addition to the normal configuration steps, the system administrator must complete the following steps to setup Authorized users for non-root users:

1. Add a stanza in the client system-options file, `dsm.sys`, for the non-root user.
2. In this stanza, use the `passworddir` option to point to a directory that is owned by the non-root user. The non-root user can then create a file in this `passworddir` directory.
3. Assign the non-root user with a unique TSM node name.
4. Ensure that an earlier `TSM.PWD` file that is not owned by the non-root user, does not exist in the `passworddir` directory. If such a file exists, change ownership of this file to the non-root user or remove the file.
5. Ensure that `TSM.KDB`, `TSM.IDX` or `TSM.sth` files that are not owned by the non-root user, do not exist in the `passworddir` directory. If such files exist, remove them.

On completion of the steps by the system administrator, the non-root user must complete the following steps:

1. Create a client system-options file, `dsm.opt`, and use the `servername` option to specify the stanza name.
2. Ensure that the `dsm.opt` file can be read by default by the `DSM_CONFIG` environment variable. Issue the **export DSM_CONFIG=<dsm.opt>** command from a shell command window to check.
3. Run the **dsmc q f** command to use password files that are pointed to by the `passworddir` option. If no password files exist, the user is prompted.

Enabling encryption for backup-archive client users

If you configure the backup-archive client to encrypt data during backup and archive operations, and if you specify the option to store the encryption key password (**encryptkey save**), by default, only root and Tivoli Storage Manager authorized users can use the stored password to encrypt or decrypt files.

Authorized users include any non-root users who have read and write access to the stored password (TSM.sth file), or users who know the password and enter it interactively.

Client options file overview

You set (specify) client options and values in a client options file. Client options can also be set on the server in a *client option set*. Client options that are set on the server in a client option set override client options that are set in the client options file.

On AIX, HP-UX, Linux, Mac, and Solaris systems, the default client options file is named `dsm.opt`. For these operating systems, two files contain backup-archive client options:

- The *client-user options* file. The default name for this file is `dsm.opt`. For brevity, this file is often called the *client options file*.
- The *client-system options* file. The default name for this file is `dsm.sys`. The client-system options file is an editable file that identifies the server and communication method, and provides the configuration for backup, archiving, hierarchical storage management, and scheduling. For brevity, this file is often called the *system options file*.

You can create multiple client options files. If your client options file is not named `dsm.opt`, or if `dsm.opt` is not in the default directory, use the `OPTFILE` client option to tell the backup-archive client which file to read the options and parameters from when the backup-archive client is started.

You cannot change the name of the client-system option file. It must be named `dsm.sys`.

You can use a text editor application to directly edit the client options file. You can also set options by using the backup-archive client GUI. In the GUI, select **Edit > Preferences** and use the Preferences Editor to set client options. Options that you set in the Preferences Editor are stored in the client options file. Not all client options can be set by using the Preferences Editor.

Restriction: For Mac OS X, the client-user options file and client-system options file must be plain text files, encoded as Unicode (UTF-8). By default, TextEdit does not save files as plain text. Select **Format > Make Plain Text** to save the files as plain text files. Select **Unicode (UTF-8)** in the **Plain Text Encoding** drop down list. Do not add the `.txt` extension when you save the file.

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

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 in either of the following ways:

```
verbose  
ve
```

Follow these rules when you add options to your options files:

- You can annotate option settings by adding comments to the options file. Begin each comment with an asterisk (*) as the first character on the line.

- Do not specify options on a line that contains a comment.
- You can optionally indent options with spaces or tabs, to make it easier to view the options and values that you specify in the file.
- Enter each option on a separate line and enter all parameters for an option on the same line, as shown in the following examples:

```
domain /home /mfg /planning /mrkting /mgmt
domain / /Volumes/fs2 /Volumes/fs2 /Volumes/fs3 /Volumes/fs4
```
- To set an option in this file, enter the option name and one or more blank spaces, followed by the option value.
- Enter one or more blank spaces between parameters.
- The lengths of file and path names in the client options files cannot exceed the following limits:
 - On AIX, Mac OS, HP-UX, and Solaris, the maximum length for a file name is 255 bytes. The maximum combined length of the file name and path name is 1024 characters. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.
 - On Linux, the maximum length for a file name is 255 bytes. The maximum combined length of the file name and path name is 4096 bytes. This matches the **PATH_MAX** that is supported by the operating system. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that comprises a path and file name can vary. The limitation is the number of bytes in the path and file components, which might or might not correspond to an equal number of characters.
 - For archive or retrieve operations, the maximum length that you can specify for a path and file name, combined, is 1024 bytes.

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

Related reference:

“Optfile” on page 467

“Query Options” on page 682

Creating and modifying the client system-options file

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

About this task

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

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

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

For Mac OS X, copy the dsm.sys.smp file to dsm.sys in one of the following locations. The default locations are listed in the order that they are searched.

1. A location identified by the DSM_DIR environment variable
2. /Library/Application Support/tivoli/tsm/client/ba/bin/
3. /Library/Preferences/Tivoli Storage Manager/

The client uses the first options file that is found. You must use the name `dsm.sys` for this file. The `dsm.sys` file is controlled by the system administrator.

For Oracle Solaris systems, copying `dsm.sys.smp` to `dsm.sys` is not required. The client options files (`dsm.opt` and `dsm.sys`) are automatically created in `/usr/bin`, if they do not already exist, and they are linked to the client installation directory when you install the client. Note that the files are not removed if you uninstall the client, so you can reuse your settings if you upgrade or reinstall the client.

For the other platforms, as the root user, copy the `dsm.sys.smp` file to `dsm.sys` and then edit that file to configure your settings. The client looks for `dsm.sys` in the directory specified by the **DSM_DIR** environment variable (if it is set and exported), and then in the installation directory.

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

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

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

Servname	server_a
COMMethod	TCPip
TCPPort	1500
TCPServeraddress	node.domain.company.com

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

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

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

The `dsm.sys` file can also contain the following option categories:

- Communication options
- Backup and archive processing options
- Restore and retrieve processing options
- Scheduling options
- Authorization options
- Error processing options
- Transaction processing option
- Web client options

You can modify your `dsm.sys` file using one of the following methods:

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

Important: For Mac OS X, the system-options file must be a plain text file, encoded as Unicode (UTF-8). By default, TextEdit does not save files as plain text. Select **Format > Make PlainText** to save the user-options file as a plain text file. Set the **Plain Text Encoding:** to Unicode (UTF-8). Do not add the `.txt` extension.

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

Related concepts:

“Client options file overview” on page 50

Chapter 10, “Processing options,” on page 287

Related reference:

“Defaultserver” on page 350

“Passwordaccess” on page 470

Creating a default client-user options file

A client-user options file stores the backup-archive client processing options. The backup-archive installation program places a sample client-user options file on disk when you install the backup-archive client. A system administrator or root can edit this file to create a default client options file, and makes the file accessible to workstation users who use the backup-archive client. Individual users can create and use their own client options file.

Before you begin

You must be root or a system administrator to complete this procedure.

About this task

Creating a default client-user options file is an optional task.

By default, the client-user options file is named `dsm.opt`, and the file contains the following types of client options:

- Backup and archive processing options
- Restore and retrieve processing options
- Scheduling options
- Format options
- Command processing options
- Authorization options
- Error processing options
- Transaction processing option
- Web client options

For Mac clients, the client installation program places a sample client-user options file named `dsm.opt.smp` in `/Libraries/Preferences/Tivoli Storage Manager/`. This directory is the same directory that the installation program places a sample client-system option file (`dsm.sys.smp`) in.

For AIX, HP-UX, and Linux clients, the client installation program places a sample client-user options file named `dsm.opt.smp` in the default client installation

directory. This directory is the same directory that the installation program places a sample client-system option file (dsm.sys.smp) in.

For Oracle Solaris clients, the installation program places an initial client-user options file named dsm.opt in the /usr/bin directory. This directory is the same directory that the installation program places a sample client-system option (dsm.sys) file in.

For all client operating systems, the following procedure instructs you to edit the sample client-user options file and save it with the default name, dsm.opt. You can save the file with a different name or path, if you want to, but if you change the file name or if you move the file from the default installation directory, you must use either of the following methods to specify the path and name of the client-user options file:

- Set the DSM_CONFIG environment variable to indicate the path and file name of the client-user option file (dsm.opt). Set the DSM_DIR environment variable to indicate the path and file name of the client-system option file (dsm.sys). For more information about the environment variables, see “Set processing environment variables” on page 57.
- Specify the backup-archive client optfile option to specify the path and file name of the client-user options file.

Note: All node users must have read access to the disk location where you store the client-user options file.

Procedure

1. Change to the directory that contains the sample client-user options file.
2. Copy the file to dsm.opt.
3. Add options for your node to the dsm.opt file. Use either of the following methods to set the client-user options:
 - Edit dsm.opt with a text editor to add the options that are needed in the node.

Note: On Mac OS X, the dsm.opt file must be saved as a plain text file and use Unicode (UTF-8) as the encoding scheme. By default, TextEdit does not save files as plain text. To save dsm.opt, in TextEdit, select **Format > Make Plain Text**. In the **Plain Text Encoding** drop-down list, select **Unicode (UTF-8)**. Do not add the .txt extension to the file name.

- Set client options by using the preferences editor. In the backup-archive client GUI, select **Edit > Client Preferences** and select the options that you want to configure. The preferences editor updates the client configuration files, dsm.opt, and dsm.sys if you add, change, or remove options. If you update the dsm.opt file while the backup-archive client is running, you must restart the backup-archive client so the updates are recognized.

The preferences editor uses the DSM_DIR environment variable to locate the client-system options file (dsm.sys) and the DSM_CONFIG environment variable to locate the client user-options file (dsm.opt). If you want dsm.opt to be in a non-default location, set DSM_CONFIG before you start backup-archive client and then use the preferences editor to set the options. The preferences editor queries the server for options on the server, but cannot change the server options file.

Related concepts:

Chapter 10, “Processing options,” on page 287

“Set processing environment variables” on page 57

Related tasks:

“Creating and modifying the client system-options file” on page 51

Creating a customized client user-options file

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

About this task

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

Procedure

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

Results

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

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

Important: For Mac OS X, the system-options file must be a plain text file, encoded as Unicode (UTF-8). By default, TextEdit does not save files as plain text. Select **Format > Make PlainText** to save the user-options file as a plain text file. Set the **Plain Text Encoding** drop-down list selection to Unicode (UTF-8). Do not add the `.txt` extension.

Related concepts:

“Environment variables”

“Client options file overview” on page 50

Environment variables

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

Set language environment variables

The backup-archive client automatically detects the language of the system locale and displays in that language.

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

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

Note: The operating system locale, the terminal character set, and the file name character set encoding must match in order for files names to be displayed or entered correctly.

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

```
export LANG=fr_FR
```

Note:

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

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

If the locale of the backup-archive client is the same as the character encoding of the file names, all of those files are backed up or restored correctly. If you are running in any single-byte character set (SBCS), then all file names are valid and are backed up or restored by the backup-archive client.

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

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

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

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

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

Set processing environment variables

Some circumstances require you to set environment variables to ensure that Tivoli Storage Manager applications can locate the files that are needed to perform client operations, and that applications can create log files that record events and errors that occur during client operations.

You must set the environment variables in any of the following circumstances:

- You want to invoke the backup-archive client from a directory other than the directory where the backup-archive client is installed
- You want to specify a different options file for the backup-archive client, the administrative client, or both.
- You do not want log files to be written to the default installation directory.

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

There are four environment variables you can set which affect backup-archive client processing:

PATH Includes the directory where the executable file for the client executables (`dsmc`, `dsmadm`, `dsmj`) resides.

DSM_DIR

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

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

When you request an image backup, image restore, snapshot-based file backup, NAS backup, or NAS restore, the client uses the `DSM_DIR` environment variable to locate the corresponding plug-in library. If `DSM_DIR` is not set, the client looks for the plug-in library in the following directories:

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

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

DSM_CONFIG

Specifies the fully-qualified path and file name of the client user options file for users who create their own personalized options file. If `DSM_CONFIG` is not set, or the client `optfile` option is not used, the client user options file is expected to satisfy these requirements:

1. The options file must be named `dsm.opt`.
2. For UNIX clients other than Mac OS X, if `DSM_DIR` is *not* set, then the file must reside in the default installation directory. If `DSM_DIR` *is* set, then the file must reside in the directory specified by `DSM_DIR`.
3. For Mac OS X, the file can reside in any of the following locations. These directories are searched in order, and the first option file found is used. `~/Library Preferences/Tivoli Storage Manager`, `/Library Preferences/Tivoli Storage Manager`, or `/Library/Application Support/tivoli/tsm/client/ba/bin`.

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

DSM_LOG

Points to the directory where you want the Tivoli Storage Manager log files to reside. You cannot specify the root (/) directory for DSM_LOG. The log files contain information about errors and events that occur during processing. The client creates the logs to help the technical support team diagnose severe errors.

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

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

Note:

1. The **errorlogname** and **schedlogname** options override DSM_LOG. If you specify the **errorlogname** client option, the file is stored in the directory specified by the **errorlogname** option and not in the location specified by DSM_LOG. If you specify the **schedlogname** client option, it is written to the directory specified by the **schedlogname** option and not in the location specified by DSM_LOG.
2. The log files cannot be symbolic links. Tivoli Storage Manager detects any such links, delete the links, then exit the operation. This prevents Tivoli Storage Manager from overwriting protected data. The affected logs are created as files in a subsequent operation.

To use the backup-archive client Java GUI program, you must export the directory where you installed the java binary file. For example, enter the following command:

```
export PATH=$PATH:java_bin_dir
```

where: *java_bin_dir* is the path to the runnable Java binary file in your file system.

Related reference:

"Optfile" on page 467

Set Bourne and Korn shell variables

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

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


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

Set C shell variables

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

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

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

Set API environment variables

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

DSMI_DIR

Points to your installation directory. The file dsm.sys must reside in the directory pointed to by DSMI_DIR. This environment variable must be present.

DSMI_CONFIG

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

DSMI_LOG

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

Note: End users of applications that are developed with the API can consult the installation directions for that application for special path names or guidelines for options.

For more information about the Tivoli Storage Manager API, see Developing solutions with the application programming interface.

Web client configuration overview

The Tivoli Storage Manager web client provides remote management of a client node from a web browser. The procedures to configure the web client vary depending on which operating system is on the client node.

Beginning with Tivoli Storage Manager Version 7.1.8, you can no longer use the web client to connect to the Tivoli Storage Manager V7.1.8 or later or V8.1.2 or later server. However, you can still use the web client to connect to Tivoli Storage Manager V7.1.7 and earlier servers, or V8.1.1 and earlier V8 servers.. For more information, see “Using the web client in the new security environment” on page 139.

Backup-archive client options are used to configure web client settings. These options include httpport, managedservices, webports, and revokeremoteaccess.

On AIX, HP-UX, Linux, Mac, and Solaris client nodes, you add the web client options to the client-systems option file (dsm.sys).

To use the web client from the Tivoli Storage Manager Operations Center interface, specify the web client address in the URL parameter of the **REGISTER NODE** or **UPDATE NODE** command. The web address must include the DNS name or IP address of the node, and the port number that the web client uses. For example, `http://node.example.com:1581`. Replace this example host name with the IP address or host name of your client node. When you access the web client by using a web browser, enter the same URL syntax in the browser address bar.

All web client messages are written to the web client log file, which is named `dsmwebcl.log`. By default, the `dsmwebcl.log` file and the backup-archive client error log file (`dsmerror.log`) are created in the client installation directory. You can use the `DSM_LOG` environment variable to override the default locations for the error logs. If you do set the `DSM_LOG` environment variable, do not specify the root directory as location for the error logs. You can also use the backup-archive client `errorlogname` option, to change the location of the error log files. If you specify this option, it overrides the `DSM_LOG` environment variable setting.

Related concepts:

“Web client options” on page 306

Related tasks:

“Configuring the web client on AIX, HP-UX, Linux, Mac, and Solaris systems”

Configuring the web client on AIX, HP-UX, Linux, Mac, and Solaris systems

To configure the web client, edit the client-system options file (`dsm.sys`) to specify the required options, and then start the client acceptor daemon.

Procedure

1. Set the following options in the `dsm.sys` file: `managedservices webclient schedule` and `passwordaccess generate`.
2. Generate the Tivoli Storage Manager password. Enter `dsmc query session`. When you are prompted for credentials, enter the Tivoli Storage Manager user name and password.

On Mac OS X systems, you can also generate the password by using the TSM Tools for Administrators application. In the application, select **Tivoli Storage Manager** to start the client.

3. Start the client acceptor daemon. Enter `dsmcad`.

On Mac OS X, you can also start the client acceptor daemon with the TSM Tools for Administrators application. In the application, select **Start the Client Acceptor Daemon**.

4. To access the web client from a browser, specify the host name or IP address of the client node in the browser address bar, followed by the web client port number. The default port number is 1581. For example, to access the web client on the node that is named `myserver.example.com`, specify:
`http://myserver.example.com:1581`

If you must change the default web client port number, use the backup-archive client `httpport` option to assign a different port number.

What to do next

After you configure the web client, you can use the Tivoli Storage Manager Operations Center or a browser to backup or restore, or archive or retrieve, data on a node.

Related concepts:

“Scheduling options” on page 302

“Web client options” on page 306

Related tasks:

“Starting a web client session” on page 140

Related reference:

“Httpport” on page 414

“Passwordaccess” on page 470

Configuring the scheduler

Your Tivoli Storage Manager administrator can schedule the client to perform tasks automatically. For scheduled events to occur on the client, you must configure the client scheduler to communicate with the Tivoli Storage Manager server.

About this task

For example, you can automatically back up files at the end of each day or archive some of your files every Friday. This procedure, which is 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 that is 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. After you start the client scheduler, no further intervention is required.

With client scheduling, you can perform the following tasks:

- Display information about available schedules.
- Display information about work that the schedule completed.
- Modify scheduling options in the `dsm.sys` file.

The most effective way to manage the client scheduler is to use the client acceptor service. You can read about a comparison between using the CAD and traditional scheduler services to manage the scheduler. You can also learn how to configure the client to use the CAD to manage the scheduler.

Comparison between client acceptor-managed services and traditional scheduler services

You can use either the client acceptor service or the traditional scheduler service to manage the Tivoli Storage Manager scheduler. A comparison of these methods is provided.

The following table shows the differences between the client acceptor-managed services and the default traditional scheduler services methods.

Table 24. Client acceptor-managed services versus traditional scheduler services

Client acceptor-managed services	Tivoli Storage Manager traditional scheduler services
Defined by using the <code>manageservices schedule</code> option and started with client acceptor services. The client acceptor daemon is started with the dsmscad command	Started with command dsmsc sched command.
The client acceptor service starts and stops the scheduler process as needed for each scheduled action.	Remains active, even after scheduled backup is complete.
Requires fewer system resources when idle.	Requires higher use of system resources when idle.
Client options and Tivoli Storage Manager server override options are refreshed each time the client acceptor services start a scheduled backup.	Client options and Tivoli Storage Manager server override options are only processed after dsmsc sched is started.
Cannot be used with <code>SESSIONINITiation=SERVEROnly</code> backups.	You must restart the scheduler process for updated client options to take effect. Important: If you run the client scheduler on the command line, the scheduler does not run as a background service. Tip: Restart the traditional scheduler periodically to free system resources previously used by system calls.

Configuring the client to use the client acceptor service to manage the scheduler

One of the most effective ways of managing the client scheduler is to use the client acceptor. You must configure the client to use the client acceptor to manage the scheduler.

Before you begin

- If you include files for encryption, ensure that the **encryptkey** option is set to save in the options file. This option is set by selecting **Save Encryption Key Password Locally** on the Authorization tab in the preference editor. Setting this option enables unattended scheduled services. If the encryption key was not previously saved, you must run an attended backup of at least one file so that you get the encryption prompt to save the key.
- You cannot use the client acceptor for scheduling when the **sessioninitiation** option is set to `serveronly`.

About this task

The client acceptor 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 client acceptor restarts the scheduler when it is time to run the scheduled event. This action reduces the number of background processes on your workstation and resolves memory retention problems that can occur when the scheduler is run without client acceptor management.

The client acceptor service is also known as the client acceptor daemon.

Procedure

Complete the following steps to use the client acceptor to manage the client scheduler:

1. From the backup-archive client GUI, select **Edit > Preferences**.
2. Click the **Web Client** tab.
3. In the **Managed Services Options** field, click **Schedule**. If you also want the client acceptor to manage the web client, click **Both** option.
4. Start the client acceptor daemon by running the following command on the command line:

```
dsmcad
```

Tip:

- You can also use the **managedservices** option in the client system-options file (dsm.sys) to specify whether the client acceptor manages the scheduler.
- If you need the client acceptor to manage the scheduler in polling mode without opening a listening port, use the **cadlistenonport** option in the dsm.sys file.

Related concepts:

"Web client configuration overview" on page 59

"Enable or disable scheduled commands" on page 269

"Scheduling options" on page 302

Related tasks:

"Setting the client scheduler process to run as a background task and start automatically at startup" on page 264

Related reference:

"Cadlistenonport" on page 332

"Managedservices" on page 448

"Sessioninitiation" on page 514

Start the client scheduler

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

Scheduling events using the command-line client

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

About this task

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

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

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

Procedure

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

Results

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

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

1. Change to the backup-archive client installation directory and enter the following command:

```
dsmc schedule
```

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

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

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

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

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

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

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

Here are some examples:

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

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

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

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

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

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

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

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

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

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

For Mac OS X:

The system administrator must generate a password so that the client can store the password in the password (TSM.sth). This can be done either with TSM Tools for Administrators or with the command line.

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

Method 1 (preferred)

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

Method 2

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

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

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

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

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

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

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

The client scheduler can fail to properly initialize at system startup because TCP/IP is not fully initialized. You might need to delay the scheduler service start up to allow time for TCP/IP to initialize.

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

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

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

Related tasks:

“Configuring the scheduler” on page 61

Related reference:

“Managedservices” on page 448

“Passwordaccess” on page 470

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.

About this task

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

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

Method 1:

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

TCP/IP port

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

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

HTTP port

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

TCP/IP ports for the remote workstation

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

TCP/IP port for administrative sessions

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

Method 2:

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

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

Note:

1. If *sessioninitiation* is set to *serveronly*, the value for the *tcpclientaddress* client option must be the same as the value for the *HLAddress* option of the **update node** or **register node** server command. The value for the *tcpclientport* client option must be the same as the value for the *LLAddress* option of the **update node** or **register node** server command.
2. If you set the *sessioninitiation* option to *serveronly*, with the exception of CAD-managed schedulers, the command-line client, backup-archive client Java GUI, and web client GUI still attempts to initiate sessions, but are blocked by the Tivoli Storage Manager server for nodes that have the *sessioninitiation* option set to *serveronly*.
3. When configuring the scheduler on a client workstation for the first time, the scheduler service might be unable to authenticate to the server when the server contacts the client scheduler to run a schedule. This can happen when the *passwordaccess* is set to generate and the Tivoli Storage Manager server is behind a firewall and the encrypted password cannot be locally stored before the scheduler is started. To correct this problem, you need to run the scheduler from the command line (`dsmc schedule`), wait until a scheduled operation starts, and enter the password for your node when prompted.
4. The client cannot prompt for the encryption key password in scheduler mode. If you are using Tivoli Storage Manager data encryption, you must run an initial interactive backup once to set up the encryption key by opening the TCP/IP connection from the client workstation to the server workstation. See **Method 1** for more information about setting up this communication. After the encryption key is set, you can use server-initiated sessions to back up the files using encryption.

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

When using Tivoli Storage Manager across a firewall in *prompted* mode, the server needs to contact the client. In order to complete this action, 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 method is typically called *socksifying* a system. Proxies are not supported, because they only route a few types of communication protocols (HTTP, FTP, GOPHER). Tivoli Storage Manager communications are not routed by proxies. It is important to note that the client creates a new connection to the Tivoli Storage Manager server when prompted. This means that the firewall configuration discussed above must be in place.

Related tasks:

“Configuring the scheduler” on page 61

Related reference:

“Sessioninitiation” on page 514

“Tcpadminport” on page 542

“Tcpport” on page 547

“Webports” on page 599

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

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

About this task

The following client components support SSL:

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

Only outgoing client/server connections support SSL. A V7.1.8 client communicating with a down-level servers supports SSL. A V7.1.8 client communicating with a V7.1.8 server must use SSL. Incoming connections (for example, client acceptor, server-initiated schedule connections) do not support SSL. Client-to-client communications do support SSL. Web GUI does not support SSL. The Web GUI is no longer supported when communicating with a V7.1.8 server.

Each Tivoli Storage Manager server that is enabled for SSL must have a unique certificate. The certificate can be one of the following types:

- A certificate that is self-signed by Tivoli Storage Manager.
- A certificate that is issued by a certificate authority (CA). The CA can be from a company such as VeriSign or Thawte, or an internal CA, maintained within your company.

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

1. Obtain the Tivoli Storage Manager server self-signed certificate (cert256.arm)
Use the cert.arm certificate file when the server is not setup to use Transport Layer Security (TLS) 1.2; otherwise, use the cert256.arm file. The client certificate file must be the same as the certificate file that the server uses.
2. Configure the clients. To use SSL, each client must import the self-signed server certificate.
Use the dsmscert utility to import the certificate.

3. For a disaster recovery of the Tivoli Storage Manager server, if the certificate has been lost, a new one is automatically generated by the server. Each client must obtain and import the new certificate.

For fast path details for communication between a V7.1.8 client and a V7.1.8 server, you can use the `SSLACCEPTCERTFROMSERV` option to automatically accept a self-signed certificate. See “Configuring by using the default security settings (fast path)” on page 129 for details.

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

1. Obtain the CA root certificate.
2. Configure the clients. To use SSL, each client must import the self-signed server certificate.

Use the `dsmcert` utility to import the certificate.

Tip: After you complete this step, if the server gets a new certificate that is signed by the same CA, the client does not need to import the root certificate again.

3. If you are recovering the backup-archive client as part of disaster recovery, you must install the SSL certificate on the server again. If the certificate was lost, you must get a new one. You do not need to reconfigure the client if the new certificate has been signed by a CA.

Next, you must import the server certificate, or the CA root certificate.

If you use a self-signed certificate

Each Tivoli Storage Manager server generates its own certificate. The certificate has a fixed file name of either `cert.arm` or `cert256.arm`. The certificate file is stored on the server workstation in the server instance directory, for example, `/opt/tivoli/tsm/server/bin/cert256.arm`. If the certificate file does not exist and you specify the **SSLTCP** or **SSLTCPADMIN** server option, the certificate file is created when you restart the server with these options set. Tivoli Storage Manager V6.3 servers (and newer versions) generate files named `cert256.arm` and `cert.arm`. Tivoli Storage Manager servers older than V6.3 generate only certificate files named `cert.arm`. You must choose the certificate that is set as the default on the server.

Follow these steps to set up the SSL connection to a server:

1. Obtain the certificate from the server administrator.
2. Import the certificate into the client key database by using the following command:

```
dsmcert -add -server <servername> -file <path_to_cert256.arm>
```

If you use a certificate from a certificate authority

If the certificate was issued by a certificate authority (CA) such as VeriSign or Thawte, the client is ready for SSL and you can skip the following steps.

For the list of preinstalled root certificates from external certificate authorities, see “Certificate Authorities root certificates” on page 72. If the certificate was not issued by one of the well-known certificate authorities, follow these steps:

1. Obtain the root certificate of the signing CA.
2. Import the certificate into the client key database by using the following command:

```
dsmcert -add -server <servername> -file <path_to_cert256.arm>
```

Important:

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

```
# chmod go+r dsmcert.*
```
5. For performance reasons, use SSL only for sessions where it is needed. A V7.1.8 client communicating with a V7.1.8 server must use SSL. SSL No (the default value) indicates that encryption is not used when data is transferred between the client and a server earlier than V7.1.8. When the client connects to a V7.1.8 or later server, the default value No indicates that object data is not encrypted. All other information is encrypted, when the client communicates with the server. When the client connects to a V7.1.8 or later server, the value Yes indicates that SSL is used to encrypt all information, including object data, when the client communicates with the server. Consider adding more processor resources on the Tivoli Storage Manager server system to manage the increased requirements.
6. In order for a client to connect to a server that is using Transport Layer Security (TLS) Version 1.2, the certificate's signature algorithm must be SHA-1 or stronger. If you are using a self-signed certificate, you must use the cert256.arm certificate. Your Tivoli Storage Manager administrator might need to change the default certificate on the Tivoli Storage Manager server. See the SSLTLS12 server option topic for details.

Additional details for a V7.1.8 client communicating with a server V7.1.7 and earlier levels.

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

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

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

If the value of **tcpport** is incorrect, the client cannot connect to the server. Specify the correct port number on the **tcpport** option.

To disable security protocols that are less secure than TLS 1.2, add the SSLDISABLELEGACYtls yes option to the client options file, or within the Java GUI select the **Require TLS 1.2 or above** checkbox on the **Communication** tab of the **Preferences editor**. Requiring TLS 1.2 or above helps prevent attacks by malicious programs.

Related reference:

“Ssl” on page 533

“Sslfipsmode” on page 536

Creating a symbolic link to access the latest GSKit library

You can create a symbolic link to point the directory where the older version of GSKit is installed to the location of the latest GSKit libraries on the system.

About this task

When you install DB2 for Linux, UNIX, and Windows, on UNIX and Linux, local GSKit libraries are also installed. Those libraries are stored in <db2_install_path>/lib64/gskit_db2 or <db2_install_path>/lib32/gskit_db2. On Windows, the default location is C:\Program Files\ibm\gsk8.

During the installation of other IBM products, such as Tivoli Storage Manager, another copy of the GSKit libraries might be installed. Depending on the product, these libraries might be either local GSKit libraries or global GSKit libraries. When DB2 for Linux, UNIX, and Windows and another IBM product that includes GSKit libraries are both installed on the same system, some interoperability issues might arise. These interoperability issues might occur because GSKit allows only libraries from a single GSKit source to exist in any single process. The interoperability issues might lead to unpredictable behavior and runtime errors.

To ensure that a single source of GSKit libraries is used, the symbolic link approach can be used. During an initial DB2 for Linux, UNIX, and Windows installation, the installer creates a symbolic link <db2_install_path>/lib64/gskit or <db2_install_path>/lib32/gskit to <db2_install_path>/lib64/gskit_db2 or <db2_install_path>/lib32/gskit_db2. These symbolic links are the default locations from where GSKit libraries are loaded. Products that bundle DB2 for Linux, UNIX, and Windows, and change the symbolic link from the default directory to the library directory of another copy of GSKit must ensure that the newly installed GSKit is at the same or a newer level. This restriction applies whether the libraries are global or local. During an upgrade or update of DB2 for Linux, UNIX, and Windows, the symbolic link is preserved. If the newly installed copy has a symbolic link to the default location, the symbolic link that is associated with the older installation copy is preserved. If the newly installed copy does not have a symbolic link to the default location, the symbolic link that is associated with the newer installation copy is preserved.

Some limitations exist since the symbolic link <db2_install_path>/lib64/gskit or <db2_install_path>/lib32/gskit is in the path of the DB2 for Linux, UNIX, and Windows installation copy. For example, if two or more instances are created for any DB2 copy, the symbolic link changes affect all the instances.

You can also modify a Domino Server GSKit in a similar manner. A Domino server does not have a GSKit folder, but it has folders C and N, and a library libgsk8iccs_64.so. You can first create soft links for these folders, and files to point to the corresponding folders on the GSKit package, where the Tivoli Storage Manager backup-archive client V8.1.2 is installed, as follows:

- `ln -s /usr/local/ibm/gsk8_64/lib64/C /opt/ibm/lotus/notes/90010/zlinux`
- `ln -s /usr/local/ibm/gsk8_64/lib64/N /opt/ibm/lotus/notes/90010/zlinux`
- `ln -s /usr/local/ibm/gsk8_64/lib64/libgsk8iccs_64.so /opt/ibm/lotus/notes/90010/zlinux`

Next, change the DPD node's password to domdsmc CHANGEADSMpwd
tv1054_domnote2 tv1054_domnote2 tv1054_domnote2. Finally, run domdsmc query
adsm.

Procedure

1. Create a symbolic link on Windows, if you have administrator privileges.
Rename the DB2 GSKit copy of the lib64 directory that is located in the default location, C:\Program Files\ibm\gsk8. Start a DOS shell, navigate to the DB2 GSKit location, and rename the directory as follows:

```
cd C:\Program Files\ibm\gsk8  
rename lib64 lib64-db2
```
2. Create a symbolic link in the location of the DB2 GSKit copy and point to the location of the TSM GSKit copy by running the following commands in the DOS shell. Navigate to the location of the DB2 GSKit copy and then create the symbolic link as follows:

```
cd C:\Program Files\ibm\gsk8  
mklink /d lib64 "c:\Program Files\Common Files\Tivoli\TSM\api64\gsk8\lib64"
```
3. Restart DB2 for changes to take effect. On startup, DB2 loads GSKit from the new location, which points to the Tivoli Storage Manager copy of GSKit. In the DB2 command prompt, enter these commands as follows:

```
db2stop  
db2start
```

Certificate Authorities root certificates

The backup-archive client includes a list of root certificates for a number of common Certificate Authorities.

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

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

- Thawte Server CA
- RSA Secure Server Certification Authority

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

Configure your system for journal-based backup

You must install and configure the journal daemon (Linux) or journal engine service (Windows) before you can perform journal-based backups.

Journal daemon configuration

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

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

After you configured the `tsmjbbd.ini` file, start the journal daemon by starting the **tsmjbbd** executable file.

To start the journal daemon after you restart your system on AIX, run the `jbbinittab` script file to add an entry to the `/etc/inittab` file. The `tsmjbbd` executable file and the `jbbinittab` script file should be in the default installation directory.

To stop the journal daemon on AIX, issue the `kill nnnn` command, where `nnnn` is the process ID of **tsmjbbd**. Before the journal daemon process (**tsmjbbd**) shuts down, it notifies the filepath kernel extension to stop buffering file changes.

Important: Do not use the `kill -9 nnnn` command, because the `kill -9` command immediately ends the process without notifying filepath to stop buffering file changes.

On Linux, the installer creates the `tsmjbbd` service in `/etc/init.d`. To control the service, run the following command as root to stop, start, or restart the service, or to check its status:



If the Linux operating system runs the `systemd` initialization service, complete the following steps to start the journal daemon:

1. Copy the provided `systemd` unit file `/opt/tivoli/tsm/client/ba/bin/tsmjbbd.service` to the `/etc/systemd/system/` directory.
2. Run the following command to refresh the `systemd` unit list:
`systemctl daemon-reload`
3. Run the following command to start the journal daemon at system boot time:
`systemctl enable tsmjbbd.service`

4. Run the following command to start the journal daemon:
`systemctl start tsmjbbd.service`

Note:

1. Network and removable file systems are not supported.
2. Periodic full incremental backups should be performed to complement daily journal-based backups. Full progressive incremental backups can take longer to run than a journal-based backup. Take this information into account when you schedule them, perhaps scheduling the incremental backups during off-peak times. Balance these two backup techniques according to your business needs. For example, you might decide to schedule nightly journal-based backups and also schedule a weekly full progressive incremental backup.
3. Journal-based backup uses the filepath kernel extension to monitor file system changes. To improve the performance of journal-based backups, directories that do not contain user files are not monitored for changes and are not included in journal-based backups. The following lists the directories that are not included in journal-based backups on AIX and Linux systems. Changes to these directories are processed if you perform periodic full incremental backups by using the **incremental** command with the **-nojournal** option.

AIX	Linux
/bin /dev /etc /lib /usr/bin /usr/lib /usr/share	/bin /boot /dev /etc /lib /proc /sbin /sys /usr/bin /usr/lib /usr/share /var

The journal daemon configuration file is periodically checked for updates to the list of journaled file systems. You can add or remove file systems from the list of monitored file systems without stopping the journal daemon.

Attention: If you bring a file system that is being monitored by the journal daemon offline, the journal database for that file system is deleted. To preserve the database, set `PreserveDbOnExit=1` in the journaled file systems settings stanza. This setting preserves the journal database when it is taken offline and ensures that the journal database is valid when the file system comes back online. For more information, see “JournaledFilesystemSettings stanza” on page 76.

The syntax for stanza and stanza settings is as follows:

Syntax for stanzas:

`[StanzaName]`

Syntax for stanza settings:

`stanzaSetting=value`

Note:

1. You can specify comments in the file by beginning the line with a semicolon.
2. Stanza and value names are not case-sensitive.

3. Numeric values can be specified in hexadecimal by preceding the value with 0x; otherwise, they are interpreted as decimal.
4. These journaled file system settings do not correlate to any settings in the client options file. The journal daemon is an independent process; it does not process any options in the client options file.

JournalSettings stanza

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

The following is the syntax for the JournalSettings stanza:

Syntax for JournalSettings stanza:

```
[JournalSettings]
```

Syntax for stanza settings:

```
JournalSettings=value
```

You can specify the following JournalSettings values:

ErrorLog

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

```
ErrorLog=/logs/jbberror.log
```

JournalDir

Directory where journal database files are stored and written.

If the path given is an absolute (for example, it begins with a dir delimiter) pathname, this is the directory used. If the path given is a relative directory name, then this path is appended to each file system name and the resulting path name is used.

The default is a directory named .tSm_JoUrNaL (used within each file system being journaled).

The advantage of having the journal database on the file system being monitored is that the database stays with the file system. The disadvantage is that the updates to the database must be processed and discarded.

Important: Directing the database to a non-journaled file system, unless this file system is shared in a cluster environment.

This setting applies to all journaled file systems but can be overridden with an override stanza for each journal file system.

JournalExcludeList stanza

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

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

Note:

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

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

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

The following pattern matching meta characters are supported:

- % Matches exactly one character.
- * Matches zero or more characters.
- %EnvVar% Expands environment variable.

The following is an exclude statement syntax example:

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

JournaledFileSystemSettings stanza

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

File systems that you specify in the JournalFileSystems.Extended stanza override any file systems specified in the list of journaled file systems that you might have previously specified in the JournaledFileSystemSettings stanza. Any other options that you have specified in the JournaledFileSystemsSettings stanza are preserved.

The syntax for the JournaledFileSystemSettings stanza is as follows:

Syntax for JournaledFileSystemSettings stanza:

[JournaledFileSystemSettings]

Syntax for stanza settings:

JournaledFileSystemSetting=value

You can specify the following *JournaledFileSystemSettings* values:

JournaledFileSystems

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

JournaledFileSystems=/home /other

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

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

JournalDbSize

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

```
JournalDBSize=0x10000000
```

NotifyBufferSize, DirNotifyBufferSize

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

```
NotifyBufferSize=0x00200000
```

PreserveDbOnExit setting

This setting allows a journal to remain valid when a journaled file system goes offline and comes back online. This is useful for preserving the journal during system reboots, and resource movement.

This setting allows a journal-based backup to continue processing when the daemon is restarted (or the file system comes back online) without performing a full incremental backup.

Note: Any change activity which occurs while the journal daemon is not running (or the file system is offline) is not recorded in the journal.

A value of 1 specifies that the journaled file system journal database is not deleted when the journal file system goes offline. The database is also valid when the journal file system comes back online. This value should be used with caution because any file system change activity which occurs while the journaled file system is offline is not reflected in the journal database. The default setting of 0 deletes the journaled file system journal database.

Note: The journal is only preserved when a journaled file system comes offline normally or is brought offline when the resource is no longer available and you specify the `deferFsMonStart` setting. If a file system comes offline due to an error such as a notification buffer overrun, the journal is not preserved.

Note: Set `PreserveDBonExit` only when you can ensure that there is a controlled shutdown of the journal service. The scope of "controlled shutdown" includes stopping the journal service in order to reboot the system, failing over a cluster resource, or moving a cluster resource. The journal database can become corrupted if the shutdown is not controlled. Therefore, perform the following steps if the journal service was not shut down in a controlled manner or if the journal database was otherwise taken offline in an uncontrolled manner.

1. Stop the journal service (if it is running)
2. Delete the corrupted journal databases
3. Restart the journal service
4. Perform an incremental backup

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

preservedBOnExit=1

deferFSMonStart setting

This setting defers an attempt to begin monitoring a file system in the following cases:

- When the specified journaled file system is not valid or available
- The journal directory for the specified journaled file system cannot be accessed or created

Resources are checked at the interval you specify using the *deferRetryInterval* setting.

A value of 1 indicates that the setting is on. A value of 0 indicates that the setting is off. The default value is off (set to 0) .

deferRetryInterval setting

This setting specifies the value in seconds that deferred file systems with the *deferRetryInterval* setting enabled are checked for availability and brought online. The default value is 5 seconds.

logFSErrors setting

A value of 1 indicates that all errors encountered accessing a journaled file system or journal directory should be logged. A value of zero indicates that logging of errors encountered while checking deferred file systems and journal directories is suppressed. This is usually used in conjunction with the *deferFSMonStart* setting to eliminate excessive File System Unavailable messages from being written to the logs when bringing a journaled file system online is deferred. The default value is 1 (log all errors).

Related concepts:

“Overriding stanzas” on page 79

“JournaledFileSystems.Extended stanza”

JournaledFileSystems.Extended stanza:

The JournaledFileSystems.Extended stanza overrides any file systems that are included in the JournaledFileSystems stanza. It also removes the 1023 character limitation imposed by the JournaledFileSystems stanza.

If you include file systems in the JournaledFileSystems stanza, the total number of characters allowed in that stanza is 1023 characters. For large configurations with many file systems, the 1023 character limit is too small to specify all file systems. If you must use more than 1023 characters to include all file systems that you want included in journal-based backups, specify the file systems in the JournaledFileSystems.Extended stanza. This extended stanza does not impose the 1023 character limitation. Values in JournaledFileSystems.Extended override any value specified in the other stanza. If a file system is specified in both the JournaledFileSystems stanza and the JournaledFileSystems.Extended stanza, the file system specified in the JournaledFileSystems stanza is ignored.

The syntax for JournaledFileSystems.Extended has a simple list form. The file systems that you want to be included in journal-based backups by editing the journal daemon configuration file (the default name is tmsjbbd.ini).

Syntax for JournaledFileSystems.Extended stanza:

[JournaledFileSystems.Extended]

Syntax for stanza settings:

/filesystem_1

```
/filesystem_2
.
.
/filesystem_n
```

List each file system that you want included in journal-based backups.

Overriding stanzas

Any setting in the **JournaledFileSystemSettings** stanza, except for the buffer sizes, can be overridden for a particular journaled file system by creating an override stanza.

HookFileName

In order for the journal to begin monitoring a file system, it must know the name of an existing file in that file system. This setting specifies an existing file. Access to this file is then used as a test of whether or not this file system is online. (The system definition of mounted cannot be used because we allow the use of virtual mount points in the backup-archive client. This means that the backup-archive client system can treat a directory as a (virtual) file system).

Therefore, if this file system can be mounted and unmounted, a **HookFileName** needs to be provided.

If a **HookFileName** is not entered, the journal daemon attempts to create a temporary file in the highest directory, use it to begin monitoring, and then delete it.

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

Syntax for JournaledFileSystemSettings stanza:
[JournaledFileSystemSettings.fs]

Syntax for stanza settings:
JournaledFileSystemSetting=override value

For example, the override stanza name for /home would be:

```
JournaledFileSystemSettings./home
HookFileName=/home/doNotDeleteThisFile
```

Client-side data deduplication

Data deduplication is a method of reducing storage needs by eliminating redundant data.

Overview

Two types of data deduplication are available: *client-side data deduplication* and *server-side data deduplication*.

Client-side data deduplication is a data deduplication technique that is used on the backup-archive client to remove redundant data during backup and archive processing before the data is transferred to the Tivoli Storage Manager server. Using client-side data deduplication can reduce the amount of data that is sent over a local area network.

Server-side data deduplication is a data deduplication technique that is done by the server. The Tivoli Storage Manager administrator can specify the data

deduplication location (client or server) to use with the **DEDUP** parameter on the **REGISTER NODE** or **UPDATE NODE** server command.

Enhancements

With client-side data deduplication, you can:

- Exclude specific files on a client from data deduplication.
- Enable a data deduplication cache that reduces network traffic between the client and the server. The cache contains extents that were sent to the server in previous incremental backup operations. Instead of querying the server for the existence of an extent, the client queries its cache.

Specify a size and location for a client cache. If an inconsistency between the server and the local cache is detected, the local cache is removed and repopulated.

Note: For applications that use the Tivoli Storage Manager API, the data deduplication cache must not be used because of the potential for backup failures caused by the cache being out of sync with the Tivoli Storage Manager server. If multiple, concurrent backup-archive client sessions are configured, there must be a separate cache configured for each session.

- Enable both client-side data deduplication and compression to reduce the amount of data that is stored by the server. Each extent is compressed before it is sent to the server. The trade-off is between storage savings and the processing power that is required to compress client data. In general, if you compress and deduplicate data on the client system, you are using approximately twice as much processing power as data deduplication alone.

The server can work with deduplicated, compressed data. In addition, backup-archive clients earlier than V6.2 can restore deduplicated, compressed data.

Client-side data deduplication uses the following process:

- The client creates extents. *Extents* are parts of files that are compared with other file extents to identify duplicates.
- The client and server work together to identify duplicate extents. The client sends non-duplicate extents to the server.
- Subsequent client data-deduplication operations create new extents. Some or all of those extents might match the extents that were created in previous data-deduplication operations and sent to the server. Matching extents are not sent to the server again.

Benefits

Client-side data deduplication provides several advantages:

- It can reduce the amount of data that is sent over the local area network (LAN).
- The processing power that is required to identify duplicate data is offloaded from the server to client nodes. Server-side data deduplication is always enabled for deduplication-enabled storage pools. However, files that are in the deduplication-enabled storage pools and that were deduplicated by the client, do not require additional processing.
- The processing power that is required to remove duplicate data on the server is eliminated, allowing space savings on the server to occur immediately.

Client-side data deduplication has a possible disadvantage. The server does not have whole copies of client files *until* you back up the primary storage pools that contain client extents to a non-deduplicated copy storage pool. (*Extents* are parts of a file that are created during the data-deduplication process.) During storage pool backup to a non-deduplicated storage pool, client extents are reassembled into contiguous files.

By default, primary sequential-access storage pools that are set up for data deduplication must be backed up to non-deduplicated copy storage pools before they can be reclaimed and before duplicate data can be removed. The default ensures that the server has copies of whole files at all times, in either a primary storage pool or a copy storage pool.

Important: For further data reduction, you can enable client-side data deduplication and compression together. Each extent is compressed before it is sent to the server. Compression saves space, but it increases the processing time on the client workstation.

In a data deduplication-enabled storage pool (file pool) only one instance of a data extent is retained. Other instances of the same data extent are replaced with a pointer to the retained instance.

When client-side data deduplication is enabled, and the server has run out of storage in the destination pool, but there is a next pool defined, the server will stop the transaction. The backup-archive client retries the transaction without client-side data deduplication. To recover, the Tivoli Storage Manager administrator must add more scratch volumes to the original file pool, or retry the operation with deduplication disabled.

For client-side data deduplication, the Tivoli Storage Manager server must be Version 6.2 or higher.

Prerequisites

When configuring client-side data deduplication, the following requirements must be met:

- The client and server must be at version 6.2.0 or later. The latest maintenance version should always be used.
- When a client backs up or archives a file, the data is written to the primary storage pool that is specified by the copy group of the management class that is bound to the data. To deduplicate the client data, the primary storage pool must be a sequential-access disk (FILE) storage pool that is enabled for data deduplication.
- The value of the DEDUPLICATION option on the client must be set to YES. You can set the DEDUPLICATION option in the client options file, in the preference editor of the backup-archive client GUI, or in the client option set on the Tivoli Storage Manager server. Use the **DEFINE CLIENTOPT** command to set the DEDUPLICATION option in a client option set. To prevent the client from overriding the value in the client option set, specify **FORCE=YES**.
- Client-side data deduplication must be enabled on the server. To enable client-side data deduplication, use the **DEDUPLICATION** parameter on the **REGISTER NODE** or **UPDATE NODE** server command. Set the value of the parameter to CLIENTORSERVER.

- Ensure files on the client are not excluded from client-side data deduplication processing. By default, all files are included. You can optionally exclude specific files from client-side data deduplication with the `exclude.dedup` client option.
- Files on the client must not be encrypted. Encrypted files and files from encrypted file systems cannot be deduplicated.
- Files must be larger than 2 KB and transactions must be below the value that is specified by the `CLIENTDEDUPTXNLIMIT` option. Files that are 2 KB or smaller are not deduplicated.

The server can limit the maximum transaction size for data deduplication by setting the `CLIENTDEDUPTXNLIMIT` option on the server. For more information about this option, see the Tivoli Storage Manager server documentation.

The following operations take precedence over client-side data deduplication:

- LAN-free data movement
- Simultaneous-write operations
- Data encryption

Important: Do not schedule or enable any of those operations during client-side data deduplication. If any of those operations occur during client-side data deduplication, client-side data deduplication is turned off, and a message is written to the error log.

The setting on the server ultimately determines whether client-side data deduplication is enabled. See Table 25.

Table 25. Data deduplication settings: Client and server

Value of the client DEDUPLICATION option	Setting on the server	Data deduplication location
Yes	On either the server or the client	Client
Yes	On the server only	Server
No	On either the server or the client	Server
No	On the server only	Server

Encrypted files

The Tivoli Storage Manager server and the backup-archive client cannot deduplicate encrypted files. If an encrypted file is encountered during data deduplication processing, the file is not deduplicated, and a message is logged.

Tip: You do not have to process encrypted files separately from files that are eligible for client-side data deduplication. Both types of files can be processed in the same operation. However, they are sent to the server in different transactions.

As a security precaution, you can take one or more of the following steps:

- Enable storage-device encryption together with client-side data deduplication.
- Use client-side data deduplication only for nodes that are secure.
- If you are uncertain about network security, enable Secure Sockets Layer (SSL).
- If you do not want certain objects (for example, image objects) to be processed by client-side data deduplication, you can exclude them on the client. If an

object is excluded from client-side data deduplication and it is sent to a storage pool that is set up for data deduplication, the object is deduplicated on server.

- Use the **SET DEDUPVERIFICATIONLEVEL** command to detect possible security attacks on the server during client-side data deduplication. Using this command, you can specify a percentage of client extents for the server to verify. If the server detects a possible security attack, a message is displayed.

Related tasks:

“Configuring the client for data deduplication”

Related reference:

“Deduplication” on page 349

“Exclude options” on page 389

“Dedupcachepath” on page 348

“Dedupcachesize” on page 348

“Enablededupcache” on page 378

“Iobjtype” on page 416

Configuring the client for data deduplication

Configure the client so that you can use data deduplication to back up or archive your files.

Before you begin

Before you configure your client to use data deduplication, ensure that the requirements listed in “Client-side data deduplication” on page 79 are met:

- The server must enable the client for client-side data deduplication with the **DEDUP=CLIENTORSERVER** parameter on either the **REGISTER NODE** or **UPDATE NODE** command.
- The storage pool destination for the data must be a data deduplication-enabled storage pool.
- Ensure that your files are bound to the correct management class.
- Files must be larger than 2 KB.

A file can be excluded from client-side data deduplication processing. By default, all files are included. Refer to the `exclude.dedup` option for details.

The server can limit the maximum transaction size for data deduplication by setting the `CLIENTDEDUPTXNLIMIT` option on the server.

Procedure

Use one of the following methods to enable data deduplication on the client:

Option	Description
Edit the client options file	• Add the deduplication yes option to the <code>dsm.sys</code> file.

Option	Description
Preferences editor	1. From the Tivoli Storage Manager window, click Edit > Client Preferences. 2. Click Deduplication. 3. Select the Enable Deduplication check box. 4. Click OK to save your selections and close the Preferences Editor.

Results

After you have configured the client for data deduplication, start a backup or archive operation. When the operation completes, the backup or archive report shows the amount of data that was deduplicated in this operation, and how many files were processed by client-side data deduplication.

If you do not have enough disk space for the backup or archive operation, you can enable client-side data deduplication without local data deduplication cache on the client by using these steps:

1. Add the deduplication yes option to the client options file.
 - Add the deduplication yes option to the dsm.sys file. You can also set this option in the GUI.
2. Turn off the local data deduplication cache by completing one of the following steps:
 - Add the ENABLEDEDUPCACHE NO option to the dsm.sys file.

You can also set this option in the backup-archive client preferences editor by clearing the **Enable Deduplication Cache** check box.

Example

The following example uses the query session command to show the type of data that was processed for data deduplication:

```
tsm> q sess
TSM Server Connection Information

Server Name.....: SERVER1
Server Type.....: Windows
Archive Retain Protect..: "No"
Server Version.....: Ver. 6, Rel. 2, Lev. 0.0
Last Access Date.....: 08/25/2009 13:38:18
Delete Backup Files.....: "No"
Delete Archive Files.....: "Yes"
Deduplication.....: "Client Or Server"

Node Name.....: AVI
User Name.....:
```

The following example uses the query management class command to show the type of data that was processed for data deduplication:

```
tsm> q mgmt -det
Domain Name : DEDUP
Activated Policy Set Name : DEDUP
Activation date/time : 08/24/2009 07:26:09
Default Mgmt Class Name : DEDUP
Grace Period Backup Retn. : 30 day(s)
Grace Period Archive Retn.: 365 day(s)
```

MgmtClass Name : DEDUP
Description : dedup - values like standard
 Space Management Technique : None
 Auto Migrate on Non-Usage : 0
 Backup Required Before Migration: YES
 Destination for Migrated Files : SPACEMGP00L
 Copy Group
 Copy Group Name.....: STANDARD
 Copy Type.....: Backup
 Copy Frequency.....: 0 day(s)
 Versions Data Exists...: 2 version(s)
 Versions Data Deleted...: 1 version(s)
 Retain Extra Versions...: 30 day(s)
 Retain Only Version....: 60 day(s)
 Copy Serialization.....: Shared Static
 Copy Mode.....: Modified
 Copy Destination.....: AVIFILEPOOL
 Lan Free Destination...: NO
Deduplicate Data.....: YES

Copy Group Name.....: STANDARD
 Copy Type.....: Archive
 Copy Frequency.....: Cmd
 Retain Version.....: 365 day(s)
 Copy Serialization.....: Shared Static
 Copy Mode.....: Absolute
 Retain Initiation.....: Create
 Retain Minimum.....: 65534 day(s)
 Copy Destination.....: FILEPOOL
 Lan Free Destination...: NO
Deduplicate Data.....: YES

ANS1900I Return code is 0.

Related concepts:

“Client-side data deduplication” on page 79

Related reference:

“Deduplication” on page 349

“Enablededupcache” on page 378

“Exclude options” on page 389

 CLIENTDEDUPTXNLIMIT option

 REGISTER NODE command

 UPDATE NODE command

Excluding files from data deduplication

You can exclude a file from data deduplication during backup or archive processing.

About this task

You can exclude only files for archive data deduplication. You can exclude files and images (where applicable) for backup data deduplication.

Procedure

If you do not want certain files to be processed by client-side data deduplication, you can exclude files from data deduplication processing using the GUI:

1. Click **Edit > Client Preferences**.
2. Click the **Include-Exclude** tab.
3. Click **Add** to open the Define Include-Exclude Options window.
4. Select a category for processing.
 - To exclude a file from data deduplication during archive processing, select **Archive** in the **Category** list.
 - To exclude a file from data deduplication during backup processing, select **Backup** in the **Category** list.
5. Select **Exclude.Dedup** in the **Type** list.
6. Select an item from the **Object Type** list.
 - For archive processing, only the **File** object type is available.
 - For backup processing, select one of the following object types:
 - **File**
 - **Image**
7. Specify a file or pattern in the **File or Pattern** field. You can use wildcard characters. If you do not want to type a file or pattern, click **Browse** to open a selection window and select a file. For mounted file spaces, you can choose the directory mount point from the selection window.
8. Click **OK** to close the Define Include-Exclude Options window. The exclude options that you defined are in an exclude statement at the bottom of the Statements list box in the **Include-Exclude Preferences** tab.
9. Click **OK** to save your selections and close the Preferences Editor.

What to do next

You can also exclude files from data deduplication processing by editing the `dsm.sys` file:

1. Add the deduplication yes option.
2. Exclude the files in a directory from data deduplication. For example, to exclude the files in the `/Users/Administrator/Documents/Taxes/` directory, add the following statement: `EXCLUDE.dedup /Users/Administrator/Documents/Taxes/.../*`
3. Exclude client-side data deduplication for image backup of a file system. For example, to exclude the `/home` file system, add the following statement: `EXCLUDE.DEDUP /home/*/* IEOBJTYPE=Image`.

Important: If an object is sent to a data deduplication pool, data deduplication occurs on the server, even if the object is excluded from client-side data deduplication.

Related concepts:

“Client-side data deduplication” on page 79

Related reference:

“Deduplication” on page 349

“Enablededupcache” on page 378

“Exclude options” on page 389

Automated client failover configuration and use

The backup-archive client can automatically fail over to a secondary server for data recovery when the Tivoli Storage Manager server is unavailable. You can configure the client for automated failover or prevent the client from failing over. You can also determine the replication status of your data on the secondary server before you restore or retrieve the replicated data.

Related tasks:

“Restoring or retrieving data during a failover” on page 229

Automated client failover overview

When there is an outage on the Tivoli Storage Manager server, the backup-archive client can automatically fail over to a secondary server for data recovery.

The Tivoli Storage Manager server that the client connects to during normal production processes is called the *primary server*. When the primary server and client nodes are set up for node replication, that server is also known as the *source replication server*.

The client data on the source replication server can be replicated to another Tivoli Storage Manager server, which is the *target replication server*. This server is also known as the *secondary server*, and is the server that the client automatically fails over to when the primary server fails.

For the client to automatically fail over to the secondary server, the connection information for this server must be made available to the client. During normal operations, the connection information for the secondary server is automatically sent to the client from the primary server during the logon process. The secondary server information is automatically saved to the client options file. No manual intervention is required by you to add the information for this server.

Each time the client logs on to the Tivoli Storage Manager server, the client attempts to contact the primary server. If the server is unavailable, the client automatically fails over to the secondary server, according to the secondary server information in the client options file.

In this failover mode, you can restore or retrieve any replicated client data. When the primary server is online again, the client automatically fails back to the primary server the next time the client is started.

For example, the following sample text is the connection information that secondary server is sent to the client and saved to the client system options file (dsm.sys):

```
*** These options should not be changed manually
REPLSERVERNAME          TARGET
REPLTCPSERVERADDRESS 192.0.2.9
REPLTCPSPORT           1501
REPLSSLPORT            1502
REPLSERVERGUID         60.4a.c3.e1.85.ba.11.e2.af.ce.00.0c.29.2f.07.d3

MYREPLICATIONServer TARGET
*** end of automatically updated options
```

Requirements for automated client failover

Before you configure or use the client for automated client failover, the backup-archive client and Tivoli Storage Manager server must meet several requirements.

Ensure that the client meets the following requirements for automated client failover:

- The primary server, secondary server, and backup-archive client must be running Tivoli Storage Manager Version 7.1, or a later version.
- The primary and secondary servers must be set up for node replication.
- The client node must be configured for node replication on the source replication server by using the REGISTER NODE REPLSTATE=ENABLED or UPDATE NODE REPLSTATE=ENABLED server commands.
- By default, the client is enabled for automated client failover. However, if the `usereplicationfailover` no option is specified in the client options file, either change the value to `yes`, or remove the option.
- Valid connection information for the secondary server must exist in the client options file. During normal operations, this information is automatically sent to the client from the primary server.
- To save the secondary server connection information that is sent from the primary server, the client must have write access to the `dsm.opt` file on Windows clients, and the `dsm.sys` file on AIX, HP-UX, Linux, Mac OS X, and Oracle Solaris clients. If the client does not have write access to these files, the secondary server information is not saved to the client options file, and an error is added to the error log.
- Non-root users cannot use the default location for the node replication table. You must specify a different location by adding the `nrtablepath` option to the `dsm.sys` file. For more information, see “Nrtablepath” on page 464.
- The following processes must occur before the connection information for the secondary server is sent to the options file:
 - The client must be backed up to the source replication server at least one time.
 - The client node must be replicated to the target replication server at least one time.
- Failover occurs for client nodes that are backed up with client-node proxy support when both the target and agent nodes are configured for replication to the target replication server. When the target node is explicitly replicated, the agent node is implicitly replicated to the target replication server as well, along with the proxy relationship.

For example, Node_B is granted authority to perform client operations on behalf of Node_A with the following server command:

```
grant proxynode target=Node_A agent=Node_B
```

If both nodes are configured for replication with the `replstate=enabled` option in the node definition, when Node_A is replicated, Node_B and the proxy relationship are replicated as well.

Restrictions for automated client failover

Review the following information to better understand the process and the restrictions that apply to automated client failover.

The following restrictions apply for automated client failover:

- When the client is in failover mode, you cannot use any functions that require data to be stored on the secondary server, such as backup or archive operations. You can use only data recovery functions, such as restore, retrieve, or query operations. You can also edit client options and change the Tivoli Storage Manager client password.
- Schedules are not replicated to the secondary server. Therefore, schedules are not run while the primary server is unavailable.
- After the client connects to the secondary server in failover mode, it does not attempt to connect to the primary server until the next initial logon to the server. The client attempts to fail over to the secondary server only when the initial connection to the primary server fails. The initial connection is the first connection that the client makes with the server.

If the primary server becomes unavailable during a client operation, the client does not fail over to the secondary server, and the operation fails. You must restart the client so that it can fail over to the secondary server, and then run the client operation again.

Restore operations that are interrupted when the primary server goes down cannot be restarted after the client fails over. You must run the whole restore operation again after the client fails over to the secondary server.

- If the Tivoli Storage Manager password is changed before the client node is replicated, the password will not be synchronized between the primary and secondary servers. If a failover occurs during this time, you must manually reset the password on the secondary server and the client. When the primary server is online again, the password must be reset for the client to connect to the primary server.

If the password is reset while the client is connected to the secondary server, the password must be reset on the primary server before the client can log on to the primary server. This restriction is true if the **passwordaccess** option is set to **generate** or if the password is manually reset.

- If you backed up or archived client data, but the primary server goes down before it replicates the client node, the most recent backup or archive data is not replicated to the secondary server. The replication status of the file space is not current. If you attempt to restore or retrieve the data in failover mode and the replication status is not current, a message is displayed that indicates that the data you are about to recover is out-of-date. You can decide whether to proceed with the recovery or wait until the primary server comes back online.
- If an administrative user ID with client owner authority exists on the source replication server, and the user ID has the same name as the client node, the administrative user ID is replicated during the node replication process on the server. If such a user ID does not exist on the source replication server, the replication process does not create this administrator definition on the target replication server.

If other administrative user IDs are assigned to the node, the Tivoli Storage Manager administrator must manually configure the administrative user IDs on the target replication server. Otherwise, the administrative user cannot connect to the target replication server (secondary server) with the web client.

- If you restore a file from the Tivoli Storage Manager, and the file system is managed by Tivoli Storage Manager for Space Management, you must not restore the file as a stub file. You must restore the complete file. Use the **restoremigstate=no** option to restore the complete file. If you restore the file as a stub from the target server, the following consequences can occur:
 - You cannot recall the file from the Tivoli Storage Manager source server by using the Tivoli Storage Manager for Space Management client.

- A Tivoli Storage Manager for Space Management reconciliation process that runs against the Tivoli Storage Manager source server expires the file. If the file is expired by a reconciliation process, you can restore the complete file with the backup-archive client and the `restoremigstate=no` option.

Failover capabilities of Tivoli Storage Manager components

Tivoli Storage Manager components and products rely on the backup-archive client or API to back up data to the primary Tivoli Storage Manager server. When the primary server becomes unavailable, some of these products and components can fail over to the secondary server, while others are not capable of failover.

To learn more about the failover capabilities of Tivoli Storage Manager components and products, see technote 1649484.

Related tasks:

“Determining the status of replicated client data” on page 92

Configuring the client for automated failover

You can manually configure the client to automatically fail over to the secondary server.

Before you begin

Before you begin the configuration:

- Ensure that the client node participates in node replication on the primary server.
- Ensure that the client meets the requirements for automated client failover.
- Use this procedure only if the connection information for the secondary server is not current or if it is not in the client options file.

About this task

You might manually configure the client for automated failover in the following situations:

- The secondary server configuration was changed and the primary server is down before the client logs on to the server. When you manually add the connection information, the client is enabled for failover to the secondary server.
- You accidentally erased some or all of the secondary server connection information in the client options file.

Tip: Instead of manually configuring the client options file, you can run the **dsmtc q session** command, which prompts you to log on to the primary server. The connection information for the secondary server is sent automatically to the client options file.

Procedure

To manually configure the client for automated failover, complete the following steps:

1. Ensure that the client is enabled for automated client failover by verifying that the `usereplicationfailover` option is either not in the client options file or is set to `yes`. By default, the client is enabled for automated client failover so the `usereplicationfailover` is not required in the client options file.

2. Obtain the connection information about the secondary server from the Tivoli Storage Manager server administrator and add the information to the beginning of the client options file. Group the statements into a stanza under the **replservername** statement.

For example, add the following statements to the `dsm.sys` file:

```
REPLSERVERNAME      TARGET
REPLTCPSERVERADDRESS 192.0.2.9
REPLTCPSPORT         1501
REPLSSLPORT          1502
REPLSERVERGUID       60.4a.c3.e1.85.ba.11.e2.af.ce.00.0c.29.2f.07.d3
```

```
Servname      server_a
COMMMethod    TCPip
TCPPort       1500
TCPServeraddress server_hostname1.example.com
PASSWORDAccess prompt
MYREPLICATIONServer TARGET
```

3. Non-root users must specify a location for the node replication table by adding the **nrtablepath** option to the `dsm.sys` file. The backup-archive client uses this table to store information about each backup or archive operation to the Tivoli Storage Manager server.

You must specify a location that your user ID has write access to. For example:
`nrtablepath /Volumes/nrtbl`

Restriction: Do not specify the root directory (/) for the location of the node replication table.

4. Save and close the client options file.
5. Restart the backup-archive client GUI or log on to the Tivoli Storage Manager server from the command-line interface. The client is connected to the secondary server.

Example

After you configured the client for automated client failover, and the client attempts to log on to the server, the following sample command output is displayed:

```
IBM Tivoli Storage Manager
Command Line Backup-Archive Client Interface
  Client Version 7, Release 1, Level 0.0
  Client date/time: 04/29/2013 11:18:51
(c) Copyright by IBM Corporation and other(s) 1990, 2013. All Rights Reserved.

Node Name: MY_NODE_NAME
ANS2106I Connection to primary TSM server 192.0.2.1 failed

ANS2107I Attempting to connect to secondary server TARGET at 192.0.2.9 : 1501

Node Name: MY_NODE_NAME
Session established with server TARGET: Windows
  Server Version 7, Release 1, Level 0.0
  Server date/time: 04/29/2013 11:18:51 Last access: 04/29/2013 09:55:56

  Session established in failover mode to secondary server
ANS2108I Connected to secondary server TARGET.
```

What to do next

You can restore or retrieve any replicated data in failover mode.

Related concepts:

“Automated client failover overview” on page 87

Related tasks:

“Restoring or retrieving data during a failover” on page 229

Related reference:

“Forcefailover” on page 408

“Myreplicationserver” on page 458

“Nrtablepath” on page 464

“Replserverguid” on page 490

“Replservername” on page 492

“Replsslport” on page 493

“Repltcpport” on page 495

“Repltcpserveraddress” on page 497

“Usereplicationfailover” on page 559

Determining the status of replicated client data

You can verify whether the most recent backup of the client was replicated to the secondary server before you restore or retrieve client data from the secondary server.

About this task

You can obtain the status of replicated client data to determine whether the most recent client backup was replicated to the secondary server.

If the time stamp of the most recent backup operation on the client matches the time stamp of the backup on the secondary server, the replication status is current.

If the time stamp of the most recent backup operation is different from the time stamp of the backup on the secondary server, the replication status is not current. This situation can occur if you backed up the client, but before the client node can be replicated, the primary server goes down.

Procedure

To determine the status of replicated client data, issue the following command at the command prompt:

```
dsmc query filespace -detail
```

The following sample output shows that the time stamps on the server and the client match, therefore the replication status is current:

#	Last Incr Date	Type	fsID	Unicode	Replication	File Space Name
---	-----	----	----	-----	-----	-----
1	00/00/0000 00:00:00	HFS	9	Yes	Current	/
	Last Store Date	Server		Local		
	-----	-----		----		
	Backup Data :	04/22/2013 19:39:17		04/22/2013 19:39:17		
	Archive Data :	No Date Available		No Date Available		

The following sample output shows that time stamps on the server and the client do not match, therefore the replication status is not current:

#	Last Incr Date	Type	fsID	Unicode	Replication	File Space Name
1	00/00/0000 00:00:00	HFS	9	Yes	Not Current	/
	Last Store Date	Server		Local		
	Backup Data :	04/22/2013 19:39:17		04/24/2013 19:35:41		
	Archive Data :	No Date Available		No Date Available		

What to do next

If you attempt to restore the data in failover mode and the replication status is not current, a message is displayed that indicates that the data you are about to restore is old. You can decide whether to proceed with the restore or wait until the primary server is online.

Related tasks:

“Restoring or retrieving data during a failover” on page 229

Related reference:

“Nrtablepath” on page 464

Preventing automated client failover

You can configure the client to prevent automated client failover to the secondary server.

About this task

You might want to prevent automated client failover, for example, if you know that the data on the client node was not replicated to the secondary server before the primary server went offline. In this case, you do not want to recover any replicated data from the secondary server that might be old.

Procedure

To prevent the client node from failing over to the secondary server, add the following statement to the client options file:

```
usereplicationfailover no
```

This setting overrides the configuration that is provided by the Tivoli Storage Manager server administrator on the primary server.

Results

The client node does not automatically fail over to the secondary server the next time it tries to connect to the offline primary server.

Related tasks:

“Determining the status of replicated client data” on page 92

Related reference:

“Usereplicationfailover” on page 559

Forcing the client to fail over

The client can immediately fail over to the secondary server even if the primary server is operational. For example, you can use this technique to verify that the client is failing over to the expected secondary server.

Procedure

To force the client to immediately fail over to the secondary server, complete the following steps:

1. Add the **forcefailover yes** option in the client-system options file (`dsm.sys`).
2. Connect to the secondary server by restarting the backup-archive client GUI or by starting a command session with the **dsmc** command.
3. Optional: Instead of updating the options file, you can establish a connection with the secondary server by specifying the `-forcefailover=yes` option with a command. For example:

```
dsmc q sess -forcefailover=yes
```

What to do next

You can verify that you are connected to the secondary server with one of the following methods:

- Check the **Secondary Server Information** field in the Connection Information window in the backup-archive client GUI.
- Check the command output when you start a command session. The status of the secondary server is displayed in the output.

Related reference:

“Forcefailover” on page 408

Configuring the client to back up and archive Tivoli Storage Manager FastBack data

Before you can back up or archive Tivoli Storage Manager FastBack client data, you must complete configuration tasks.

First ensure that you have configured the backup-archive client and that you installed the Tivoli Storage Manager FastBack client.

Install the FastBack client by using the information at IBM Tivoli Storage Manager Fastback.

After you install the FastBack client, complete the following tasks:

1. Register a node for each FastBack client where data is backed up or archived. The node name must be the short host name of the FastBack client.

This is a one-time configuration performed once for each FastBack client whose volumes need to be backed up or archived.

This registration step must be performed manually only when the Tivoli Storage Manager client is used as a stand-alone application.

The Administration Center does this node registration automatically when the user creates schedules for archiving or backing up FastBack data using the Administration Center. Starting with Tivoli Storage Manager Version 7.1, the Administration Center component is no longer included in Tivoli Storage Manager distributions. FastBack users who have an Administration Center from a previous server release can continue to use it to create and modify FastBack schedules. If you do not already have an Administration Center installed, you can download the previously-released version from <ftp://public.dhe.ibm.com/storage/tivoli-storage-management/maintenance/admincenter/v6r3/>. If you do not have an Administration Center installed, you must create and modify

FastBack schedules on the Tivoli Storage Manager server. For information about creating schedules on the server, see the Tivoli Storage Manager server documentation.

2. Use the server **GRANT PROXY** command to grant proxy authority to your current Tivoli Storage Manager client node on each node representing a FastBack client created in step 1. The FastBack node should be the target, and the current client node should be the proxy.

This is a one-time configuration, and is performed by the Administration Center if the backup or archive is initiated by the Administration Center.

3. Run the **set password** command to store the credentials of the FastBack repositories where the backup-archive client connects. Run the **set password -type=fastback** command once for each repository where the backup-archive client is expected to connect.

The credentials that are stored depends on these configurations:

- Backup-archive client on a FastBack server
- Backup-archive client on a FastBack DR Hub
- Backup-archive client on a dedicated proxy workstation

For information about integrating Tivoli Storage Manager and Tivoli Storage Manager FastBack, see Integrating Tivoli Storage Manager FastBack and Tivoli Storage Manager.

Related concepts:

“Installation requirements for backing up and archiving Tivoli Storage Manager FastBack client data” on page 9

Related reference:

“Set Password” on page 734

Cluster environment configuration and use

The term *cluster* has different meanings in different environments. It can mean highly available, high performance, load balancing, grid computing, or some combination of all of these terms.

There are currently several clustering products available for UNIX and Linux, and this section defines those aspects of a clustering environment that need to exist in order for this backup methodology to work correctly. A basic understanding of how your cluster software functions is needed. Cluster software related activities such as the development of application start and stop scripts are not described in this section.

A cluster environment refers to a UNIX or a Linux environment which exhibits the following characteristics:

- Disks are shared between physical workstations, either in an exclusive fashion (only one host has access to the logical disk at any one time) or in a concurrent fashion.
- Disks appear as local disks to the host and not as network resources.

Important: Mount the file systems locally to the system, not through a LAN-based file share protocol such as network file system (NFS).

- Mount points of local disks are identical on each physical host in the environment (if file system /group1_disk1 fails from NodeA to NodeB, it is mounted on NodeB as /group1_disk1).

Overview of cluster environments

Cluster environments can be set up in many different configurations. This section describes the most popular cluster configurations.

Active/Active: Pool cluster resources

In an active/active configuration, each node is actively managing at least one resource and is configured as a backup for one or more resources in the cluster. Active/active is the most common form of a cluster environment.

Active/Passive: Fault tolerant

In an active/passive configuration, one node actively manages the resource.

The other node is only used if the primary node experiences a fault and the resource needs to failover. An active/passive cluster is a subtype of an active/active cluster.

Concurrent access

In a concurrent configuration, more than one node manages a resource. When a fault occurs, the resource continues to be managed by the other nodes.

Configuring the backup-archive client in a cluster environment

The backup-archive client is designed to manage the backup of cluster drives by placing the backup-archive client within the context of the cluster's resource groups.

About this task

This gives the advantage of backing up data from local resources (as opposed to accessing the data across the network) to maximize the performance of the backup operation and to manage the backup data relative to the resource group. Therefore, the backup-archive client can always back up data on cluster resources as if the data were local data and maximize backup performance. This ensures that critical data is getting backed up across system failures.

For example, an active/active cluster environment has three physical hosts in the cluster named NodeA, NodeB, and NodeC.

The nodes have the following qualities:

- NodeA owns the cluster resource with file systems /A1 and /A2
- NodeB owns the cluster resources with file systems /B1 and /B2
- NodeC owns the cluster resources with file systems /C1 and /C2

Note: NodeA might also have two non-clustered volumes, /fs1 and /fs2, that must be backed up.

For best backup performance, you might want all nodes in the cluster to perform the backups of the shared file systems that they own. When a node failover occurs, the backup tasks of the failed node shift to the node to which the failover occurred. For example, when NodeA fails over to NodeB, the backup of /A1 and /A2 moves to NodeB.

The following are prerequisites before configuring the backup-archive client to back up cluster and non-cluster volumes:

- A separate backup-archive client scheduler process must be run for each resource group being protected. In normal conditions, each node would have two scheduler processes: one for the cluster resources, and one for the local file systems. After a failure, additional scheduler processes are started on a node in order to protect the resources that have moved over from another node.
- The backup-archive client password files must be stored on cluster disks so that after a failure, the generated backup-archive client password is available to the takeover node.
- The file systems to be protected as part of a resource group are defined using the backup-archive client domain option. The domain option is specified in the `dsm.sys` file, which should also be stored on a cluster disk so that it can be accessed by the takeover node.

Follow the steps below to configure the backup-archive client in a cluster environment.

Procedure

1. Register backup-archive client node definitions on the Tivoli Storage Manager server. All nodes in the cluster must be defined on the Tivoli Storage Manager server. If you are defining multiple cluster resources in a cluster environment to failover independently, then unique node names must be defined per resource group. For the above sample three-way active/active cluster configuration, define three nodes (one per resource), as follows: (1) `tsm: IBM>register node nodeA nodeApw domain=standard`, (2) `tsm: IBM>register node nodeB nodeBpw domain=standard`, (3) `tsm: IBM>register node nodeC nodeCpw domain=standard`.
2. Configure the backup-archive client system-options file. Each node in the cluster must have separate server stanzas for each cluster resource group in order to be backed up in each respective `dsm.sys` file. You must ensure that the server stanzas are identical in the system option files on each node. Alternatively, you can place the `dsm.sys` file on a shared cluster location. The server stanzas defined to back up clustered volumes must have the following special characteristics:
 - The `nodename` option must refer to the client node name registered on the Tivoli Storage Manager server. If the client node name is not defined, the node name defaults to the host name of the node, which might conflict with other node names used for the same client system.

Important: Use the `nodename` option to explicitly define the client node.

- The `tcpclientaddress` option must refer to the service IP address of the cluster node.
- The `passworddir` option must refer to a directory on the shared volumes that are part of the cluster resource group.
- The `errorlogname` and `schedlogname` options must refer to files on the shared volumes that are part of the cluster resource group to maintain a single continuous log file.
- All include exclude statements must refer to files on the shared volumes that are part of the cluster resource group.
- If you use the `incl excl` option, it must refer to a file path on the shared volumes that are part of the cluster group.
- The stanza names identified with the `servername` option must be identical on all systems.

3. Other backup-archive client options can be set as needed. In the following example, all three nodes, NodeA, NodeB, and NodeC, must have the following three server stanzas in their dsm.sys file:

```
Servename      server1_nodeA
nodename       NodeA
commmethod     tcpip
tcpport        1500
tcpserveraddress server1.example.com
tcpclientaddress nodeA.example.com
passwordaccess generate
passworddir    /A1/tsm/pwd
manageservices schedule
schedlogname   /A1/tsm/dsmsched.log
errorlogname   /A1/tsm/errorlog.log
```

```
Servename      server1_nodeB
nodename       NodeB
commmethod     tcpip
tcpport        1500
tcpserveraddress server1.example.com
tcpclientaddress nodeB.example.com
passwordaccess generate
passworddir    /B1/tsm/pwd
manageservices schedule
schedlogname   /B1/tsm/dsmsched.log
errorlogname   /B1/tsm/errorlog.log
```

```
Servename      server1_nodeC
nodename       NodeC
commmethod     tcpip
tcpport        1500
tcpserveraddress server1.example.com
tcpclientaddress nodeC.example.com
passwordaccess generate
passworddir    /C1/tsm/pwd
manageservices schedule
schedlogname   /C1/tsm/dsmsched.log
errorlogname   /C1/tsm/errorlog.log
```

4. Configure the backup-archive client user-options file. The options file (dsm.opt) must reside on the shared volumes in the cluster resource group. Define the DSM_CONFIG environment variable to refer to this file. Ensure that the dsm.opt file contains the following settings:
- The value of the servename option must be the server stanza in the dsm.sys file which defines parameters for backing up clustered volumes.
 - Define the clustered file systems to be backed up with the domain option.

Note: Ensure that you define the domain option in the dsm.opt file or specify the option in the schedule or on the backup-archive client command line. This is to restrict clustered operations to cluster resources and non-clustered operations to non-clustered resources.

In the example, nodes NodeA, NodeB, and NodeC set up their corresponding dsm.opt file and DSM_CONFIG environment variable as follows:

NodeA:

- 1) Set up the /A1/tsm/dsm.opt file:

```
servename server1_nodeA
domain    /A1 /A2
```

- 2) Issue the following command or include it in your user profile:

```
export DSM_CONFIG=/A1/tsm/dsm.opt
```


NodeB:

- 1) Set up the /B1/tsm/dsm.opt file:

```
servername server1_nodeB
domain      /B1 /B2
```

- 2) Issue the following command or include it in your user profile:

```
export DSM_CONFIG=/B1/tsm/dsm.opt
```

NodeC:

- 1) Set up the /C1/tsm/dsm.opt file:

```
servername server1_nodeC
domain      /C1 /C2
```

- 2) Issue the following command or include it in your user profile:

```
export DSM_CONFIG=/C1/tsm/dsm.opt
```

5. Set up the schedule definitions for each cluster resource group. After the basic setup is completed, define the automated schedules to back up cluster resources to meet the backup requirements. The procedure illustrates the schedule setup by using the built-in Tivoli Storage Manager scheduler. If you are using a vendor-acquired scheduler, refer to the documentation provided by the scheduler vendor.

- Define a schedule in the policy domain where cluster nodes are defined. Ensure that the schedule's startup window is large enough to restart the schedule on the failover node in case of a failure and fallback event. This means that the schedule's duration must be set to longer than the time it takes to complete the backup of the cluster data for that node, under normal conditions.

If the reconnection occurs within the start window for that event, the scheduled command is restarted. This scheduled incremental backup reexamines files sent to the server before the failover. The backup then "catches up" to where it stopped before the failover situation.

In the following example, the `clus_backup` schedule is defined in the standard domain to start the backup at 12:30 A.M. every day with the duration set to two hours (which is the normal backup time for each node's data).

```
tsm: IBM>define schedule standard clus_backup action=incr
starttime=00:30 startdate=TODAY Duration=2
```

- Associate the schedule with the all of the backup-archive client nodes defined to backup cluster resources, as follows: (1) `tsm: IBM>define association standard clus_backup nodeA`, (2) `tsm: IBM>define association standard clus_backup nodeB`, (3) `tsm: IBM>define association standard clus_backup nodeC`.

6. Set up the scheduler service for backup. On each client node, a scheduler service must be configured for each resource that the node is responsible for backing up, under normal conditions. The `DSM_CONFIG` environment variable for each resource scheduler service must be set to refer to the corresponding `dsm.opt` file for that resource. For the sample configuration, the following shell scripts must be created to allow `dsmcad` processes to be started, as needed, from any node in the cluster.

```

NodeA: /A1/tsm/startsched
#!/bin/ksh
export DSM_CONFIG=/A1/tsm/dsm.opt
dsmcad
NodeB: /B1/tsm/startsched
#!/bin/ksh
export DSM_CONFIG=/B1/tsm/dsm.opt
dsmcad
NodeC: /C1/tsm/startsched
#!/bin/ksh
export DSM_CONFIG=/C1/tsm/dsm.opt
dsmcad

```

7. Define the backup-archive client to the cluster application. To continue the backup of the failed resource after a failover condition, the Tivoli Storage Manager scheduler service (for each cluster client node) must be defined as a resource to the cluster application in order to participate in the failover processing. This is required in order to continue the backup of the failed resources from the node that takes over the resource. Failure to do so would result in the incomplete backup of the failed resource. The sample scripts in step 5 can be associated with the cluster resources to ensure that they are started on nodes in the cluster while the disk resources being protected move from one node to another. The actual steps required to set up the scheduler service as a cluster resource are specific to the cluster software. Refer to your cluster application documentation for additional information.
8. Ensure that the password for each node is generated and cached correctly in the location specified using the `passworddir` option. This can be validated by performing the following steps:
 - a. Validate that each node can connect to the Tivoli Storage Manager server without the password prompt. You can do this by running the backup-archive client command line interface and issuing the following command on each node:

```
#dsmc query session
```

If you are prompted to submit your password, enter the password to run the command successfully and rerun the command. The second time, the command should run without the prompt for the password. If you get prompted for the password, check your configuration.
 - b. Validate that the other nodes in the cluster can start sessions to the Tivoli Storage Manager server for the failed-over node. This can be done by running the same commands, as described in the step above, on the backup nodes. For example, to validate if NodeB and NodeC can start a session as NodeA in the failover event without prompting for the password, perform the following commands on NodeB and NodeC

```
#export DSM_CONFIG=/A1/tsm/dsm.opt
#dsmc query session
```

The prompt for the password might appear at this time, but this is unlikely. If you are prompted, the password was not stored in the shared location correctly. Check the `passworddir` option setting used for NodeA and follow the configuration steps again.
 - c. Ensure that the schedules are run correctly by each node. You can trigger a schedule by setting the schedule's start time to now. Remember to reset the start time after testing is complete.

```
tsm: IBM>update sched standard clus_backup starttime=now
```
 - d. Failover and fallback between nodeA and nodeB, while nodeA is in the middle of the backup and the schedule's start window, is still valid. Verify that the incremental backup continues to run and finish successfully after failover and fallback.

- e. Issue the command below to cause a node's (nodeA) password to expire. Ensure that backup continues normally under normal cluster operations, as well as failover and fallback:

```
tsm: IBM>update node nodeA forcep=yes
```

9. Configure the backup-archive client to back up local resources.

- a. Define client nodes on the Tivoli Storage Manager server. Local resources should never be backed up or archived using node names defined to back up cluster data. If local volumes that are not defined as cluster resources are backed up, separate node names (and separate client instances) must be used for both non-clustered and clustered volumes.

In the following example, assume that only NodeA has local file systems /fs1 and /fs2 to be backed up. In order to manage the local resources, register a node NodeA_local on the Tivoli Storage Manager server: tsm: IBM>register node nodeA_local nodeA_localpw domain=standard.

- b. Add a separate stanza in each node's system options file dsm.sys that must back up local resources with the following special characteristics:

- The value of the tcpclientaddress option must be the local host name or IP address. This is the IP address used for primary traffic to and from the node.
- If the client backs up and restores non-clustered volumes without being connected to the cluster, the value of the tcpclientaddress option must be the boot IP address. This is the IP address used to start the system (node) before it rejoins the cluster:

Example stanza for NodeA_local:

```
Servername      server1_nodeA_local
nodename        nodeA_local
commethod       tcpip
tcpport         1500
tcpserveraddress server1.example.com
tcpclientaddress nodeA_host.example.com
passwordaccess  generate
managedservices schedule
```

- c. Define the user options file dsm.opt in a path that is on a non-clustered resource.

- The value of the servername option must be the server stanza in the dsm.sys file which defines parameters for backing up non-clustered volumes.
- Use the domain option to define the non-clustered file systems to be backed up.

Note: Ensure that you define the domain option in the dsm.opt file or specify the option in the schedule or on the backup-archive client command line, in order to restrict the backup-archive operations to non-clustered volumes.

In the following example, nodeA uses the following /home/admin/dsm.opt file and sets up the DSM_CONFIG environment to refer to /home/admin/A1.dsm.opt.

Contents of /home/admin/A1.dsm.opt

```
servername ibm_nodeA_local
domain     /fs1 /fs2
```

```
export DSM_CONFIG=/home/admin/A1.dsm.opt
```

- d. Define and set up a schedule to perform the incremental backup for non-clustered file systems.

```
tsm: IBM>define schedule standard local_backup action=incr  
starttime=00:30 startdate=TODAY Duration=2
```

Associate the schedule with all of the backup-archive client nodes that are defined to backup non-clustered resources.

```
tsm: IBM>define association standard nodeA_local
```

10. Restore cluster file system data. All volumes in a cluster resource are backed up under the target node defined for that cluster resource. If you need to restore the data that resides on a cluster volume, it can be restored from the client node that owns the cluster resource at the time of the restore. The backup-archive client must use the same user options file (dsm.opt) that was used during the backup to restore the data. There are no additional setup requirements necessary to restore data on cluster volumes.
11. Restore local file system data. The non-clustered volumes are backed up under the separate node name setup for non-clustered operations. In order to restore this data, the backup-archive client must use the same user options file dsm.opt that was used during the backup. In the example, set environment variable DSM_CONFIG to refer to /home/admin/A1.dsm.opt prior to performing a client restore for the local node nodeA_local.

Related concepts:

Chapter 5, “Restoring your data,” on page 221

Enabling web client access in a Cluster Environment

If Tivoli Storage Manager web client access is needed during a failover condition, you must configure the web client acceptor daemon that is associated with the cluster to failover along with the cluster resource.

Before you begin

Beginning with Tivoli Storage Manager Version 7.1.8, you can no longer use the web client to connect to the Tivoli Storage Manager V7.1.8 or later or V8.1.2 or later server. However, you can still use the web client to connect to Tivoli Storage Manager V7.1.7 and earlier servers, or V8.1.1 and earlier V8 servers.. For more information, see “Using the web client in the new security environment” on page 139.

About this task

After you have completed the configuration steps described in the *Configuring the backup-archive client in a cluster environment* section, perform the additional steps described below to complete the web client access setup:

Procedure

1. Set up the client acceptor daemon to manage the web client and scheduler. The client acceptor daemon should be set up to manage schedulers as well as web client access. This reduces the number of daemons that need to be configured as cluster applications and thus simplifies the configuration and administration. When a failover occurs, the client acceptor starts on the node that is managing the takeover.
2. Update the managedservices option in the system-options file dsm.sys on each node for each server stanza, as shown below for NodeA

```
Servername      server1_NodeA
nodename        NodeA
commethod       tcpip
tcpp            1500
tcps            server1.example.com
tcpclientaddress nodeA.example.com
passwordaccess  generate
passworddir     /A1/tsm/pwd
schedlogn       /A1/tsm/dsmsched.log
errorlogname    /A1/tsm/errorlog.log
managedservices webclient schedule
```

3. Set up the client acceptor daemon to use a known HTTP port. By default, the client acceptor daemon uses HTTP port 1581, when available, for the web client access. If this port is not available, the client acceptor finds the first available port, starting with 1581. In a failover condition of an active-active cluster configuration, a failover cluster host system is probably running multiple instances of the client acceptor. If default settings are used for the HTTP port, the failover node uses any available port for the client acceptor being failed over, since the default port is probably in use by the current client acceptor processes of the failover host. This causes problems for the web client associated with the client acceptor that failed over, as the new HTTP port is not known to the web client users. You might use the `httpport` option to specify the specific ports for the web client access for each resource. This allows you to always use the same port when connecting from a web browser, independent of the node serving the cluster resource. Add the `httpport` option in the system-options file (`dsm.sys`) on each node for each server stanza as follows, making sure that each stanza uses a unique value:

```

Servername      server1_NodeA
nodename        NodeA
commmethod      tcpip
tcpp            1500
tcps            server1.example.com
tcpclientaddress nodeA.example.com
passwordaccess  generate
passworddir     /A1/tsm/pwd
managedservices webclient schedule
schedlogn       /A1/tsm/dsmsched.log
errorlogname    /A1/tsm/errorlog.log
httpport        1510

Servername      server1_NodeB
nodename        NodeB
commmethod      tcpip
tcpp            1500
tcps            server1.example.com
tcpclientaddress nodeB.example.com
passwordaccess  generate
passworddir     /B1/tsm/pwd
managedservices webclient schedule
schedlogn       /B1/tsm/dsmsched.log
errorlogname    /B1/tsm/errorlog.log
httpport        1511

Servername      server1_NodeC
nodename        NodeC
commmethod      tcpip
tcpp            1500
tcps            server1.example.com
tcpclientaddress nodeC.example.com
passwordaccess  generate
passworddir     /C1/tsm/pwd
managedservices webclient schedule
schedlogn       /C1/tsm/dsmsched.log
errorlogname    /C1/tsm/errorlog.log
httpport        1512

```

Migrating legacy AIXIBM PowerHA SystemMirror setups

If you are currently using the backup-archive client in an IBM PowerHA SystemMirror environment using the `clusternode` option, you must update your current configurations. The `clusternode` option is no longer supported.

About this task

Perform the following steps to update your current configurations:

Procedure

1. Update the backup-archive client system-options file. As with the `clusternode` option, each node in the cluster must continue to have separate server stanzas for each cluster resource group to be backed up in each respective `dsm.sys` file. The existing `dsm.sys` file for NodeA might appear as follows:

Servername	server1_nodeA
commmethod	tcpip
tcpp	1500
tcps	server1.example.com
tcpclientaddress	nodeA.example.com
passwordaccess	generate
passworddir	/A1
clusternode	yes
managedservices	schedule
schedlogn	/A1/dsmsched.log
errorlogname	/A1/errorlog.log

2. Notice that no nodename option is used in this sample. Make the following changes to the existing dsm.sys file for NodeA.

- Remove the clusternode option.
- Specify a nodename option if you do not have one already specified.

3. The new dsm.sys file for NodeA should appear as follows:

Servername	server1_nodeA
commmethod	tcpip
nodename	myclus (myclus is the existing cluster name)
tcpp	1500
tcps	server1.example.com
tcpclientaddress	nodeA.example.com
passwordaccess	generate
passworddir	/A1
managedservices	schedule
schedlogn	/A1/dsmsched.log
errorlogname	/A1/errorlog.log

4. Register backup-archive client nodes on the Tivoli Storage Manager server. If new backup-archive client nodes are added in the first step to replace the current default value of the cluster node name, register those nodes on the Tivoli Storage Manager server.
5. Update schedule definitions. If new backup-archive client nodes are added in the previous step, ensure that the backup schedule definitions used earlier to back up this node's data are now associated with the new client node names.
6. Validate the setup. For more information, see “Configuring the backup-archive client in a cluster environment” on page 96

AIX configuration considerations prior to performing snapshot-based file backups and archives

If you are configuring your Tivoli Storage Manager AIX client to perform snapshot-based file backups and archives, there are some items that you need to consider.

- Ensure that the volume group containing the file system to be snapshot has sufficient disk space to allow JFS2 external snapshots to be created for the file system.
- The client uses a default size of 100 percent of the file system size for the snapshot size. This value was found to be most appropriate for file systems with even moderate file system activity. If you need to lower this value based on your experience with your own file system activity, you can use the snapshotcachesize option to fine-tune this value.
- Do not enable internal snapshots when creating new JFS2 file systems on AIX 6.1 or later for all file systems managed by Tivoli Storage Manager. The client uses external snapshots and JFS2 does not allow the creation of external and internal snapshots concurrently for the same file system.

Related reference:
“Snapshotcachesize” on page 525

Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups

You must configure the NetApp file server connection information to run the snapshot difference incremental backup command on the backup-archive client. You must also use the **set password** command to specify the file server host name, and the password and user name that is used to access the file server.

Procedure

1. Establish a console session on the NetApp filer and define a new user on the file server by using the following steps:
 - a. Add the user ID to a group that permits users to log in to the file server with http and running API commands.
 - b. From the file server, enter the following command to list the user ID to verify the settings and verify that the output is similar:

```
useradmin user list snapdiff_user
```

```
Name: snapdiff_user
Info:
Rid: 131077
Groups: snapdiff_group
Full Name:
```

For 7-mode NetApp filers:

```
Allowed Capabilities: login-http-admin,api-*
```

For clustered-data ONTAP NetApp filers, the only capability that is required is `ontapapi` with the `admin` role.

- c. If the **security.passwd.firstlogin.enable** option for the user ID on the NetApp server is set to on, ensure that all groups have the **login-telnet** and **cli-passwd*** capabilities.

Tip: When **security.passwd.firstlogin.enable** option is enabled, the user ID is set to expired when created. The user cannot run any commands, including snapshot differential incremental, until their password is changed. Users in groups that do not have these capabilities cannot log in to the storage system. For information about defining a user ID and a password on the NetApp file server, see the NetApp documentation.

2. Configure the NetApp Data ONTAP built-in HTTP server to allow remote administrative sessions to the NetApp filer.
 - a. If you plan to use a plain HTTP connection for snapshot differential backups, turn on the **httpd.admin.enable** option on the NetApp filer.
 - b. If you plan to use a secure HTTPS connection for snapshot differential backups (by specifying the **-snapdiffhttps** option), turn on the **httpd.admin.ssl.enable** option on the NetApp filer.
 - c. From the Tivoli Storage Manager client node, test the connection between the Tivoli Storage Manager client computer and the NetApp ONTAP server to ensure that firewalls or other NetApp configuration options do not prevent you from connecting to the NetApp server.

Tip: See the NetApp ONTAP documentation for instructions on how to test the connection.

3. Export the NetApp volumes and consider the following settings:

Tip: See the NetApp documentation for details on exporting the NetApp volumes for use with AIX, or Linux hosts.

- Map the NetApp volumes by using an NFS mount.
 - Ensure the NetApp volumes have the UNIX security setting
4. Set the user ID, and password on the backup-archive client for the user ID that you created in step 1 on page 106 using the following steps:
 - a. Log in as the root user ID.

- b. From the backup-archive client command line, enter the following command:

```
dsmc set password -type=filer my_file_server snapdiff_user newPassword
```

Substitute the following values:

my_file_server

This value is the fully qualified host name of your NetApp file server.

snapdiff_user

This value is the user ID that you created in step 1 on page 106.

newPassword

This value is the password for the user ID that you created in step 1 on page 106.

Related reference:

“Snapdiff” on page 519

“Snapdiffhttps” on page 524

“Createnewbase” on page 342

Protecting clustered-data ONTAP NetApp file server volumes

You can create a snapshot differential incremental backup of a volume on a NetApp file server that is part of a clustered-data ONTAP configuration (c-mode file server).

Before you begin

- Complete the procedure in “Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 106.
- Ensure that the clustered-data ONTAP environment is correctly set up by the NetApp storage virtual machine administrator.

Restriction: Tivoli Storage Manager support for snapshot differential incremental backups of clustered-data ONTAP volumes is supported only on NetApp ONTAP 8.2.1 and later versions.

About this task

In a clustered-data ONTAP environment, storage virtual machines (also referred to as data vServers) contain data volumes that can be protected by the backup-archive client.

A storage virtual machine consists of a single infinite volume or one or more flex volumes. Volumes are accessed remotely using file sharing (CIFS on Windows operating systems, NFS on AIX and Linux operating systems).

The storage virtual machines are managed by the cluster management filer, which is the physical filer (the c-mode filer) on which the storage virtual machines reside. The backup client is installed on the remote machine that accesses the volumes.

The backup-archive client must be configured with credentials for the NetApp c-mode filers that are being accessed for backup operations.

Requirements: The following information is required for this procedure:

- The host name or IP address of the cluster management filer.
- The host name or IP address of the storage virtual machine.
- The storage virtual machine name.
- The cluster management filer credentials (user name and password).
- The cluster management filer user that is configured by the client must be assigned the `ontapapi` capability with the role of `admin`.

The `ontapapi` capability does not allow interactive access to the filer with methods such as `telnet`, `ssh`, or `http/https`. No other user capabilities are required to run snapshot differential incremental backups.

Procedure

Complete the following steps on the remote machine where the backup-archive client is installed:

1. Configure the backup-archive client with the cluster management filer credentials. Use the **`dsmc set password`** command to store the credentials of the management filer that is associated with the storage virtual machine. For example, enter the following command:

```
dsmc set password -type=filer management_filer_hostname
management_filer_username management_filer_password
```

Where:

management_filer_hostname

The host name or IP address of the cluster management filer.

management_filer_username

The user name of the cluster management filer.

management_filer_password

The password for user of the management filer.

Tip: The cluster management filer password is encrypted when it is stored by the backup-archive client.

2. Associate each storage virtual machine with the management filer with the **`dsmc set netappsvm`** command. For example, enter the following command:

```
dsmc set netappsvm storage_virtual_machine_hostname
management_filer_hostname storage_virtual_machine_name
```

Where:

storage_virtual_machine_hostname

The host name or IP address of the storage virtual machine that is used to mount volumes to back up.

management_filer_hostname

The host name or IP address of the cluster management filer.

storage_virtual_machine_name

The name of the storage virtual machine.

Note: The host name or IP address of the storage virtual machine that is used to mount volumes must be consistent with what is specified in the **dsmc set** commands. For example, if the volumes are mounted with a storage virtual machine IP address, the IP address (not the host name) must be used in the **dsmc set** commands. Otherwise, client authentication with the cluster management filer fails.

You need only to specify the **dsmc set netappsvm** command once for each storage virtual machine. If the storage virtual machine is moved to a different cluster management filer, you must use the command to update the associated cluster management filer host name.

3. Mount the remote storage virtual machine to a local file system. For example, enter the following command for each storage virtual machine:

```
mount storage_virtual_machine_hostname /tmp/fs1
```

Where:

storage_virtual_machine_hostname

The host name or IP address of the storage virtual machine.

/tmp/fs1

An example of a file system to mount the storage virtual machine volume to.

4. Start a full progressive incremental backup of a flex or infinite volume.

By default, HTTP access to the NetApp file server is not enabled. If you did not configure your file server to allow access by using HTTP, use the backup-archive client **snapdiffhttps** option to enable access to the cluster management server with the HTTPS protocol.

For example, on AIX or Linux clients, enter the following command:

```
dsmc incr /tmp/fs1 -snapdiff -snapdiffhttps
```

Tip: You need only to run the full progressive incremental backup once. After this backup is successfully completed, run differential backups in future backup operations.

5. Start a snapshot differential backup of the flex or infinite volume.

For example, on AIX or Linux clients, enter the following command:

```
dsmc incr /tmp/fs1 -snapdiff -snapdiffhttps
```

Example

A backup-archive client user wants to complete a snapshot differential incremental backup of the volumes on a c-mode file server. The user is using a Windows backup-archive client to complete the backup and the volumes are mounted as CIFS shares. The c-mode filer configuration is as follows:

ONTAP 8.31 management filer

Hostname: netapplmgmt.example.com

User: netapplmgmt_user

Password: pass4netapplmgmt

CIFS Domain Controller: WINDC

Domain User: domainuser

Flex volume storage virtual machine

Hostname: netapp1-v1.example.com
Storage virtual machine name: netapp1-client1
CIFS share: demovol
Volume name: demovol

Infinite volume storage virtual machine

Hostname: netapp1-v4.example.com
Storage virtual machine name: netapp1-infiniteVolume1
CIFS Share: InfiniteVol

The user completes the following steps on the backup-archive client:

1. Configure the client with the management filer credentials by issuing the following command:

```
dsmc set password -type=filer netapplmgmt.example.com netapplmgmt_user  
pass4netapplmgmt
```
2. Define storage virtual machine associations for each storage virtual machine with the following commands:

```
dsmc set netappsvm netapp1-v1.example.com netapplmgmt.example.com  
netapp1-client1  
  
dsmc set netappsvm netapp1-v4.example.com netapplmgmt.example.com  
netapp1-infiniteVolume1
```
3. Map remote volumes to drive letters for each storage virtual machine:

```
net use y: \\netapp1-v1.example.com\demovol WINDC\domainuser  
net use z: \\netapp1-v4.example.com\InfiniteVol WINDC\domainuser
```
4. Run a full progressive incremental backup of the flex volume and infinite volume:

```
dsmc incr y: -snapdiff -snapdiffhttps  
dsmc incr z: -snapdiff -snapdiffhttps
```

You need only to run the full progressive incremental backup once. After this backup is successfully completed, run differential backups in future backup operations.
5. Run a snapshot differential backup of the flex volume and infinite volume:

```
dsmc incr y: -snapdiff -snapdiffhttps  
dsmc incr z: -snapdiff -snapdiffhttps
```

SnapMirror support for NetApp snapshot-assisted progressive incremental backup (snapdiff)

You can use NetApp's SnapDiff backup processing in conjunction with NetApp's SnapMirror replication to back up NetApp source or destination filer volumes.

In a NetApp SnapMirror environment, data that is on volumes attached to the primary data center are mirrored to volumes attached to a remote server at a disaster recovery site. The NetApp filer in the primary data center is called the source filer; the NetApp filer at the disaster recovery site is called the destination filer. You can use the backup-archive client to create snapshot differential backups of the source or destination filer volumes.

Scenario: Back up data on a source filer volume

You can configure the backup archive client to back up data from the source filer volumes. This scenario requires you to configure a backup-archive client node such that it has access to the NetApp source filer volumes by using NFS-exported shares to mount the filer volumes.

For example, assume a configuration where the source filer is named ProdFiler. Assume that a volume named UserDataVol exists on ProdFiler filer and that the volume is accessible by using NFS from a backup-archive client node. Assume that the share is mounted as UserDataVol_Share.

When you initiate a snapshot differential backup, the NetApp filer creates a new differential snapshot on the volume that is being backed up. That differential snapshot is compared with the base (previous) snapshot. The base snapshot name was registered on the Tivoli Storage Manager server when the previous backup was completed. The contents of that base snapshot are compared to the differential snapshot that is created on the source filer volume. Differences between the two snapshots are backed up to the server.

The following command is used to initiate the snapshot differential backup. The command is entered on the console of a client node that is configured to access and protect the source filer volumes. Because this command is issued to back up volumes on a source filer, a new snapshot (the differential snapshot) is created and the snapshot registered on the Tivoli Storage Manager server is used as the base snapshot. Creating both the differential and base snapshots is the default behavior; the `-diffsnapshot=create` option is a default value, and it does not need to be explicitly specified on this command.

```
dsmc incr \\ProdFiler\UserDataVol_Share -snapdiff -diffsnapshot=create
```

Back up data on a destination filer

A more typical configuration is to offload the backups from the source filer by creating backups of the source volumes by using the replicated volume snapshots stored on the destination filer. Ordinarily, backing up a destination filer presents a problem because creating a snapshot differential backup requires that a new snapshot must be created on the volume that you are backing up. The destination filer volumes that mirror the contents of the source volumes are read only volumes, so snapshots cannot be created on them.

To overcome this read-only restriction, client configuration options are provided to allow you to use the existing base and differential snapshots on the read-only destination volume to back up changes to the Tivoli Storage Manager server.

Like in the source filer scenario, the destination filer volumes are accessed by using NFS-exported shares.

Snapshot differential options summary

The `useexistingbase` option causes the most recent snapshot on the volume to be used as the base snapshot, when a base snapshot must be established. A new base snapshot is established when any of the following conditions are true:

- When this backup is the initial backup.
- When `createneibase=yes` is specified.

- When the base snapshot that was registered by a previous differential snapshot no longer exists, and an existing snapshot that is older than the missing base snapshot does not exist.

If this option is not specified, a new snapshot is created on the volume that is being backed up. Because destination filer volumes are read-only volumes, `useexistingbase` must be specified when creating snapshot differential backups of destination filer volumes. If `useexistingbase` is not specified, snapshot differential backups of a destination filer volume fail because the new snapshot cannot be created on the read-only volume.

When backing up destination filer volumes, use both the `useexistingbase` option and the `diffsnapshot=latest` option to ensure that the most recent base and most recent differential snapshots are used during the volume backup.

You use the `basesnapshotname` option to specify which snapshot, on the destination filer volume, to use as the base snapshot. If you do not specify this option, the most recent snapshot on the destination filer volume is used as the base snapshot. You can use wildcards to specify the name of the base snapshot.

You use the `diffsnapshotname` option to specify which differential snapshot, on the destination filer volume, to use during a snapshot differential backup. This option is only specified if you also specify `diffsnapshot=latest`. You can use wildcards to specify the name of the differential snapshot.

The `diffsnapshot=latest` option specifies that you want to use the latest snapshot that is found on the file server as the source snapshot.

Additional information about each of these options is provided in the *Client options reference* topics.

Snapshot differential backup command examples

In the examples that follow, assume that volumes on a source filer are replicated, by using NetApp's SnapMirror technology, to a disaster recovery filer (host name is DRFiler). Because the DRFiler volumes are read only, you use the options to specify which of the replicated snapshots that you want to use as the base snapshot, and which of the snapshots you want to use as the differential snapshot. By specifying the snapshots to use when creating a snapshot differential backup of a destination filer, no attempt is made to create a snapshot on the read-only volumes.

The following commands are used to initiate snapshot differential backups. Most of these commands create snapshot differential backups by using snapshots stored on the destination filer volumes. When backing up from a destination filer volume, be sure to include the `-useexistingbase` option, because that option prevents attempts to create a new snapshot on the read-only destination filer volumes.

Example 1: Back up a destination filer by using default nightly backups that were created by the NetApp snapshot scheduler

```
dsmc incr \\DRFiler\UserDataVol_Share -snapdiff -useexistingbase
-diffsnapshot=latest -basesnapshotname="nightly.?"
```

You can use a question mark (?) to match a single character. In this example, `-basesnapshotname=nightly.?` uses the latest base snapshot that is named "nightly.", followed by a single character (for example: `nightly.0`, `nightly.1`, and so on).

Example 2. Back up a destination filer volume by using snapshots created manually (not created by the NetApp snapshot scheduler)

```
dsmc incr \\DRFiler\UserDataVol_Share -snapdiff -useexistingbase  
-diffsnapshot=latest -basesnapshotname="share_vol_base?"  
-diffsnapshotname="share_vol_diff?"
```

This example also uses the question mark (?) wildcard to illustrate the syntax if the base and differential snapshot names have different numbers as part of the name.

Example 3. Back up a destination filer volume, and specify which snapshots to use for the base and differential snapshots

```
dsmc incr \\DRFiler\UserDataVol_Share -snapdiff -useexistingbase  
-diffsnapshot=latest -basesnapshotname="share_vol_base"  
-diffsnapshotname="share_vol_diff_snap"
```

Example 4: Back up script-generated snapshots that use a naming convention

In this example, a script that is running on the NetApp filer adds a date and time stamp to the snapshot names. For example, a snapshot created on November 3, 2012 at 11:36:33 PM is named

UserDataVol_20121103233633_snapshot. You can use wildcards with the options to select the most recent base and differential snapshots. For example:

```
dsmc incr \\DRFiler\UserDataVol_Share -snapdiff -useexistingbase  
-basesnapshotname="UserDataVol_Share_*_snapshot" -diffsnapshot=latest  
-diffsnapshotname="UserDataVol_Share_*_snapshot"
```

-useexistingbase selects the most recent base snapshot. Adding an asterisk (*) wildcard to -basesnapshotname selects the most recent base snapshot that follows the script-naming convention. The -diffsnapshot=latest option suppresses the creating of a new differential snapshot and -diffsnapshotname= selects the most recent existing differential snapshot that follows the script-naming convention. (The asterisks wildcards match any string).

Example 5: Perform a snapshot differential backup by using an existing differential snapshot that exists on the source filer

To use an existing differential snapshot that exists on the source filer, use the -diffsnapshot=latest to prevent the creation of a new differential snapshot. Also use the -diffsnapshotname option to specify which existing differential snapshot to use. The snapshot you specify is compared to the base snapshot, which was registered in the Tivoli Storage Manager server database when the last backup was created. For example:

```
dsmc incr \\ProdFiler\UserDataVol_Share -snapdiff -diffsnapshot=latest  
-diffsnapshotname="share_vol_diff_snap"
```

Register your workstation with a server

Before you can use Tivoli Storage Manager, you must set up a node name and password and your node must be registered with the server.

The process of setting up a node name and password is called *registration*. There are two types of registration: *open* and *closed*.

Your Tivoli Storage Manager administrator chooses the type of registration for your site.

Restriction: Beginning with the Tivoli Storage Manager Version 7.1.8 server, open registration is no longer available. You must use closed registration. Open registration is available only for the Tivoli Storage Manager V7.1.7 or earlier server, or the V8.1.1 or V8.1.0 server.

You must be a root user or authorized user to perform this required task.

The web client interface is available only for connections to Tivoli Storage Manager V7.1.7 and earlier servers, or V8.1.1 and earlier V8 servers. 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.

The Tivoli Storage Manager server administrator must specify the `userid` parameter with the **REGISTER NODE** server command:

```
REGISTER NODE node_name password userid=user_id
```

where the node name and the administrative user ID must be the same. For example:

```
REGISTER NODE node_a mypassw0rd userid=node_a
```

Closed registration

With closed registration, a Tivoli Storage Manager administrator must register your workstation as a client node with the server. If your enterprise uses closed registration, you must provide some information to your Tivoli Storage Manager administrator.

About this task

You must provide the following items to your Tivoli Storage Manager administrator:

- Your node name (the value returned by the **hostname** command, the name of your workstation, or the node name you specified with the *nodename* option). If you do not specify a node name with the *nodename* option, the default login ID is the name that the **hostname** command returns.
- The initial password you want to use, if required.
- Contact information, such as your name, user ID, and phone number.

Your Tivoli Storage Manager administrator defines the following for you:

- The policy domain to which your client node belongs. A policy domain contains policy sets and management classes that control how Tivoli Storage Manager manages the files you back up and archive.
- Whether you can compress files before sending them to the server.
- Whether you can delete backup and archive data from server storage.

Open registration

With open registration, a system administrator can register your workstation as a client node with the Tivoli Storage Manager Version 7.1.7 or earlier server, or the V8.1.1 or V8.1.0 server.

About this task

The first time you start a session, you are prompted for information necessary to register your workstation with the Tivoli Storage Manager server that is identified in your client options file. You need to supply your node name, a password, and contact information.

When you use open registration:

- Your client node is assigned to a policy domain named **standard**.
- You can delete archived copies of files from server storage, but not backup versions of files.

If necessary, your Tivoli Storage Manager administrator can change these defaults later.

Creating an include-exclude list

If you do not create an include-exclude list, the backup-archive client considers all files for backup services and uses the default management class for backup and archive services.

About this task

This is an optional task, but an important one.

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

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

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

Procedure

You can use the following methods to create an include-exclude list or specify an include-exclude file:

- You can add include-exclude statements in the backup-archive client GUI or web client directory tree. The online help provides detailed instructions.
 1. Open the **Edit** menu and select **Client Preferences**. In the Preferences dialog, select the **Include/Exclude** tab. You can specify an INCLEXCL file using the Preferences editor. However, you cannot create the INCLEXCL file using the Preferences editor.
 2. Create the include-exclude list manually, following the steps listed.
- You can create an include-exclude list manually by performing the following steps:
 1. Determine your include and exclude requirements.

2. Locate the server stanza in your dsm.sys file. Each server stanza must have its own include-exclude list.
3. Enter your include and exclude statements. The client evaluates all exclude.fs and exclude.dir statements *first* (regardless of their position within the include-exclude list), and removes the excluded file spaces, directories, and files from the list of objects available for processing. All other include-exclude statements are processed from the bottom of the list up. Therefore, it is important to enter all your include-exclude statements in the proper order. For example, in the following include-exclude list the includefile.cpp file *is not* backed up:

```
include /Users/user01/Documents/includefile.cpp
exclude /Users/user01/Documents/.../*
```

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

```
exclude /Users/user01/Documents/.../*
include /Users/user01/Documents/includefile.cpp
```

4. Save the file and close it.
For Mac OS X, ensure that you save the file as plain text encoded as Unicode (UTF-8 or UTF-16). Do not add the .txt extension.
5. Restart the client to enable your include-exclude list.

Related concepts:

“Considerations for Unicode-enabled clients” on page 421

“System files to exclude” on page 119

Chapter 9, “Storage management policies,” on page 275

Related reference:

“Incl excl” on page 420

Include-exclude options

This topic provides brief descriptions of the include and exclude options that you can specify in your client options file, a minimum include-exclude list that excludes system files, a list of supported wildcard characters, and examples of how you might use wildcard characters with include and exclude patterns.

Exclude file spaces and directories

Use exclude.dir statements to exclude all files and subdirectories in the specified directory from processing.

The backup-archive client 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 26 lists the options you can use to exclude file spaces and directories from processing.

Table 26. Options for excluding file spaces and directories

Option	Description
exclude.fs “Exclude options” on page 389	Excludes file spaces matching the pattern. The client does not consider the specified file space for processing and the usual deleted-file expiration process cannot occur. If you exclude a file space that was previously included, existing backup versions remain on the server subject to retention rules specified in the associated management class definition.

Table 26. Options for excluding file spaces and directories (continued)

Option	Description
exclude.dir "Exclude options" on page 389	Excludes a directory, its files, and all its subdirectories and their files from backup processing. For example, the statement <code>exclude.dir /test/dan/data1</code> excludes the <code>/test/dan/data1</code> directory, its files, and all its subdirectories and their files. Using the <code>exclude.dir</code> option is preferable over the standard <code>exclude</code> option to exclude large directories containing many files that you do not want to back up. You cannot use include options to override an <code>exclude.dir</code> statement. Only use <code>exclude.dir</code> when excluding an entire directory branch. - Use the following statements to exclude volumes <code>/Volumes/disk2</code> altogether from backup processing. Note that the volume (<code>/Volumes/disk2</code>) is backed up, but all other directories on <code>/Volumes/disk2</code> is excluded. <pre>exclude /Volumes/disk2/* exclude.dir /Volumes/disk2/*</pre> - An alternative method for excluding an entire volume from domain incremental backup is to use a domain statement to exclude the volume. For example: <pre>domain "-/Volumes/disk2"</pre> This alternative still permits selective backup processing of files on <code>/Volumes/disk2</code>.

Exclude files and directories from a journal-based backup

There are two methods of excluding files and directories from a journal-based backup.

- On AIX and Linux, 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.
- On AIX and Linux the other method is to add exclude statements to the journal configuration file `tsmjbbd.ini`, to prevent journal entries from being added for the files or directories, which prevents them from being processed during a journal-based backup.

If you are running AIX Version 6.1 or later, add an `exclude .snapshot` statement to the `tsmjbbd.ini` file to prevent JFS2 internal snapshot directories from being monitored by the journal-based backup daemon.

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

Control processing with exclude statements

After the client evaluates all exclude statements, the following options are evaluated against the remaining list of objects available for processing.

Table 27 lists the options that you can use to control processing with include and exclude statements.

Table 27. Options for controlling processing using include and exclude statements

Option	Description	Page
Back up processing		

Table 27. Options for controlling processing using include and exclude statements (continued)

Option	Description	Page
exclude exclude.backup exclude.file exclude.file.backup	<i>These options are equivalent.</i> Use these options to exclude a file or group of files from backup services and space management services (if the HSM client is installed). The exclude.backup option only excludes files from normal backup, but not from HSM.	"Exclude options" on page 389
include include.backup include.file	Use these options to include files or assign management classes for backup processing.	"Include options" on page 421
include.fs	Controls how the client processes your file space for incremental backups.	"Include options" on page 421
Archive processing		
exclude.archive	Excludes a file or group of files from archive services.	"Exclude options" on page 389
include include.archive	<i>These options are equivalent.</i> Use these options to include files or assign management classes for archive processing.	"Include options" on page 421
Image processing		
exclude.fs.nas	Excludes file systems on the NAS file server from an image backup when used with the backup nas command. If you do not specify a NAS node name, the file system identified applies to all NAS file servers. The backup nas command ignores all other exclude statements including exclude.fs and exclude.dir statements. This option is for AIX and Solaris clients <i>only</i> .	"Exclude options" on page 389
exclude.image	Excludes mounted file systems and raw logical volumes that match the specified pattern from full image backup operations. Incremental image backup operations are unaffected by exclude.image. This option is valid for AIX, HP-UX, Solaris, and all Linux clients.	"Exclude options" on page 389
include.fs.nas	Use the include.fs.nas option to bind a management class to Network Attached Storage (NAS) file systems. To specify whether the client saves Table of Contents (TOC) information during a NAS file system image backup, use the toc option with the include.fs.nas option in your dsm.sys file. For more information, see "Toc" on page 552. This option is valid only for AIX and Solaris clients.	"Include options" on page 421

Table 27. Options for controlling processing using include and exclude statements (continued)

Option	Description	Page
include.image	Includes a file space or logical volume, assigns a management class, or allows you to assign one of several image backup processing options to a specific logical volume when used with the backup image command. The backup image command ignores all other include options. This option is valid for AIX, HP-UX, Solaris, Linux x86_64, and Linux on POWER® only.	"Include options" on page 421

System files to exclude

There are some system files that should be placed in the client options file so that they are excluded.

Attention: These system files are either locked by the operating system or they can cause problems during restore. These are system files that cannot be recovered without the possibility of corrupting the operating system, or temporary files with data that you can easily recreate.

Note: This section applies to Mac OS X only.

The implicitly generated statements can be seen in the lines of output of the **query indexcl** command with the source "operating system".

The backup-archive client adds the following exclude statements to the include-exclude list from your dsm.sys file. Do not include any of these statements in the dsm.sys file, or duplicate entries occurs.

```
EXCLUDE.ARCHIVE "/.../Desktop DB"
EXCLUDE.BACKUP "/.../Desktop DB"
EXCLUDE.ARCHIVE "/.../Desktop DF"
EXCLUDE.BACKUP "/.../Desktop DF"
EXCLUDE.ARCHIVE /.vol
EXCLUDE.BACKUP /.vol
EXCLUDE.ARCHIVE /automount
EXCLUDE.BACKUP /automount
EXCLUDE.ARCHIVE /Network
EXCLUDE.BACKUP /Network
EXCLUDE.ARCHIVE /dev
EXCLUDE.BACKUP /dev
EXCLUDE.BACKUP /.vol/.../*
EXCLUDE.ARCHIVE /.vol/.../*
EXCLUDE.BACKUP /automount/.../*
EXCLUDE.ARCHIVE /automount/.../*
EXCLUDE.BACKUP /Network/.../*
EXCLUDE.ARCHIVE /Network/.../*
EXCLUDE.BACKUP /dev/.../*
EXCLUDE.ARCHIVE /dev/.../*
EXCLUDE.DIR /.vol
EXCLUDE.DIR /automount
EXCLUDE.DIR /Network
EXCLUDE.DIR /dev
```

Note:

1. Do not specify volumes with periods in the name (...). The backup-archive client uses the sequence of periods as part of include-exclude processing. The

client reports an invalid include-exclude statement if a volume has a sequence of periods in the name. The volume *must* be renamed.

2. Objects that have a type of rhap and a creator of lcmt are excluded from processing. Generally, these are special file-system objects that can also be created with the **mknod** command or are UNIX mount points. The objects or mount points must be manually recreated as part of a full system restore.

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

```
EXCLUDE    /.../dsmsched.log
EXCLUDE    /.../dsmprune.log
EXCLUDE    /.../dsmj.log
EXCLUDE    /.../dsmerror.log
EXCLUDE    /.../.hotfiles.bTree

EXCLUDE.DIR /private/tmp
EXCLUDE.DIR /private/var/vm
EXCLUDE.DIR /private/var/tmp
EXCLUDE.DIR /private/var/db/netinfo/local.nidb

EXCLUDE.DIR /.../.Trashes
EXCLUDE.DIR /.../.Spotlight-*
EXCLUDE.DIR /.../Library/Caches
EXCLUDE.DIR /.../.fsevents
```

Include and exclude files that contain wildcard characters

You must use special escape characters when including or excluding files and directories that contain wildcard characters.

The backup-archive client treats wildcard characters in different ways on different platforms.

The names of directories and files can contain different symbols. The types of symbols that are allowed depend on the operating system.

For example, on AIX, the names of directories or files can contain:

* ? : []

To specify files and directories in include and exclude statements, you must use the escape character "\" to specify the wildcards. However, the escape character can only be used inside the character classes "[]".

The following examples illustrate how to specify files and directories that contain wildcard characters using the escape character and character classes in include-exclude statements.

To exclude the single directory /usr1/[dir2] from backup processing, enter the following in the dsm.sys file or the include-exclude file:

```
exclude.dir "/usr1/[\[ ]dir2[\ ]]"
```

To exclude the single file /usr1/fi*le1 from backup processing, enter the following statement in the dsm.sys file or the include-exclude file:

```
exclude "/usr1/fi[\*]le1"
```

Tip: If you use the Preferences Editor to include or exclude a single file or directory that contains wildcard characters, you must manually edit the include or exclude statement to escape the wildcard characters. The Preferences Editor does

not automatically escape the wildcard characters. Follow the previous examples to edit the include or exclude statements in the dsm.sys file or the include-exclude file.

Related concepts:

“Wildcard characters” on page 610

Include and exclude groups of files with wildcard characters

You can use wildcard characters to include or exclude groups of files.

To specify groups of files that you want to include or exclude, use the wildcard characters listed in the following table. This table applies to include and exclude statements *only*.

A very large include-exclude list can decrease backup performance. Use wildcards and eliminate unnecessary include statements to keep the list as short as possible.

Table 28. Wildcard and other special characters

Character	Function
?	The match one character matches any single character <i>except</i> the directory separator; it does not match the end of the string. For example: - The pattern ab?, matches abc, but does not match ab, abab, or abzzz. - The pattern ab?rs, matches abfrs, but does not match abrs, or abllrs. - The pattern ab?ef?rs, matches abdefjrs, but does not match abefrs, abdefrs, or abefjrs. - The pattern ab??rs, matches abcdrs, abzzrs, but does not match abrs, abjrs, or abkkrs.
*	The match-all character. For example: - The pattern ab*, matches ab, abb, abxxx, but does not match a, b, aa, bb. - The pattern ab*rs, matches abrs, abtrs, abrsrs, but does not match ars, or aabrs, abrrs. - The pattern ab*ef*rs, matches abefrs, abefghrs, but does not match abefr, abers. - The pattern abcd.*, matches abcd.c, abcd.txt, but does not match abcd, abcdc, or abcdtxt.
/...	The match- <i>n</i> character matches zero or more directories.
[	The open character-class character begins the enumeration of a character class. For example: xxx[abc] matches xxxa, xxxb, or xxxc.
–	The character-class range includes characters from the first character to the last character specified. For example: xxx[a-z] matches xxxa, xxxb, xxxc, ... xxxz.
\	The literal escape character. When used within a character class, it treats the next character literally. When used outside of a character class, it is not treated in this way. For example, if you want to include the ']' in a character class, enter [...\]]...]. The escape character removes the usual meaning of ']' as the close character-class character.
]	The close character-class character ends the enumeration of a character class.

Related concepts:

“Wildcard characters” on page 610

Examples using wildcards with include and exclude patterns

The backup-archive client accepts the `exclude.dir` option, which can be used to exclude directory entries. However, the `include` and `exclude.dir` options cannot be used together.

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

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

enter:

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

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

Table 29 shows how to use wildcard characters to include or exclude files.

Table 29. Using wildcard characters with include and exclude patterns

Task	Pattern
Exclude all files that end with <code>.doc</code> , except those found in the home directory of aleko, Documents directory.	EXCLUDE <code>/.../*.doc</code> INCLUDE <code>"/home/aleko/Documents/*.doc"</code>
Exclude all files during backup with an extension of <code>.bak</code> , except those found on the <code>/usr</code> file system in the <code>dev</code> directory.	exclude <code>/.../*.bak</code> include <code>/usr/dev/*.bak</code>
Exclude all files and directories under any Documents directory that might exist, <i>except</i> for the Current file of user aleko.	EXCLUDE <code>/.../Documents/.../*</code> INCLUDE <code>"/home/aleko/Documents/Current"</code>
Exclude all files in any directory named "tmp" and its subdirectories, <i>except</i> for the file <code>/home/tmp/save.fil</code> .	exclude <code>/.../tmp/.../*</code> include <code>/home/tmp/save.fil</code>
Exclude any <code>.cpp</code> file in any directory on the Vol1, Vol2, Vol3, and Vol4 volumes.	EXCLUDE <code>/Volumes/Vol[1-4]/.../*.cpp</code>
Exclude any <code>.cpp</code> file in any directory on the Vol1, Vol2, Vol3, and Vol4 volumes.	EXCLUDE <code>/Volumes/Vol[1-4]/.../*.cpp</code>
Exclude any <code>.cpp</code> file in any directory on the /fs1, /fs2, /fs3 and /fs4 file systems.	EXCLUDE <code>/fs[1-4]/.../*.cpp</code>
Exclude the <code>.cpp</code> files found in the <code>/fs2/source</code> directory.	EXCLUDE <code>/fs2/source/*.cpp</code>
Exclude any <code>.o</code> file in any directory on the <code>/usr1</code> , <code>/usr2</code> , and <code>/usr3</code> file systems.	exclude <code>/usr[1-3]/.../*.o</code>
Exclude the <code>.o</code> files found in the root directory in the <code>usr2</code> file system <i>only</i> .	exclude <code>/usr2/*.o</code>
Exclude any file that resides under the tmp directory found in any file system.	exclude <code>/.../tmp/.../*</code>
Exclude the entire directory structure <code>/var/spool</code> from all processing.	exclude.dir <code>/var/spool</code>
Exclude a single file system from backup processing.	exclude.fs <code>/fs1</code> exclude.fs <code>home:</code>

Table 29. Using wildcard characters with include and exclude patterns (continued)

Task	Pattern
Exclude all file systems mounted anywhere in the /test/myfs/fs01 and /test/myfs/fs02 directory tree from backup processing.	exclude.fs /test/myfs/fs01/.../* exclude.fs /test/myfs/fs02/*
Exclude the /home/mydir/test1 directory and any files and subdirectories under it.	exclude.dir /home/mydir/test1
Exclude all directories under the /home/mydir directory with names beginning with test.	exclude.dir /home/mydir/test*
Exclude all directories directly under the /mydir directory with names beginning with test, on any file system.	exclude.dir /.../mydir/test*
Exclude the raw logical volume from image backup.	exclude.image /dev/hd0
Exclude all symbolic links or aliases (aliases apply to Mac OS X) from backup processing, except for the Docs directory for user1.	EXCLUDE.ATTRIBUTE.SYMLINK /.../* INCLUDE.ATTRIBUTE.SYMLINK /Users/ user1/Docs/*

Related concepts:

“Examples using wildcards with include and exclude patterns” on page 122

Related reference:

“Exclude options” on page 389

Symbolic link and alias processing

The backup-archive client evaluates all `exclude.fs` and `exclude.dir` statements and removes the excluded file spaces and directories.

After this initial evaluation, the client evaluates any include-exclude statements for controlling symbolic link and alias processing (`exclude.attribute.symlink` and `include.attribute.symlink`) against the remaining list of objects available for processing.

Alias processing applies to Mac OS X.

Table 30 defines options for controlling symbolic link and alias processing.

Table 30. Options for controlling symbolic link and alias processing

Option	Description	Page
<code>exclude.attribute.symlink</code>	Excludes a file or a group of files that are symbolic links or aliases from backup processing only.	“Exclude options” on page 389
<code>include.attribute.symlink</code>	Includes a file or a group of files that are symbolic links or aliases within broad group of excluded files for backup processing only.	“Include options” on page 421

Determine compression and encryption processing

The backup-archive client evaluates `exclude.dir` and any other include-exclude options controlling backup and archive processing, and then determines which files undergo compression and encryption processing.

The following options determine which files undergo compression and encryption processing.

Table 31. Options for controlling compression and encryption

Option	Description	Page
Compression processing		
exclude.compression	Excludes files from compression processing if <code>compression=yes</code> is specified. This option applies to backups and archives.	"Exclude options" on page 389
include.compression	Includes files for compression processing if <code>compression=yes</code> is specified. This option applies to backups and archives.	"Include options" on page 421
Encryption processing		
exclude.encrypt	Excludes files from encryption processing.	"Exclude options" on page 389
include.encrypt	Includes files for encryption processing. The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received. Important: The <code>include.encrypt</code> option is the only way to enable encryption on the Backup-Archive client. If no <code>include.encrypt</code> statements are used encryption will not occur.	"Include options" on page 421

Preview include-exclude list files

You can preview the list of objects to be backed up or archived according to the include-exclude list, prior to sending any data to the server.

The backup-archive client GUI directory tree shows detailed information of included and excluded objects. The directory tree windows in the backup-archive client GUI allow you to select files and directories to include or exclude. You should use this **preview** command to make sure that you include and exclude the correct files. The following is a sample scenario for using the include-exclude preview function.

For example, follow these steps to back up the files on your `/Users/home` file space:

1. Start the backup-archive client GUI and open the Backup tree. You can see all of the directories and files that have been excluded by your options file and other sources.
2. Scroll down the tree and notice that all of the `*.o` files in your `/Volumes/home/mary/myobjdir` are backed up.
3. You do not want to back up all of the `*.o` files, so you right click a `.o` file, and choose "View File Details" from the popup menu.
4. The dialog shows that these files are included, so click the "Advanced" button and create a rule to exclude all `.o` files from the `DATA:\home` file space.

5. A rule is created at the bottom of your options file. The current directory is refreshed in the Backup tree, and the .o files have the red 'X', meaning they are excluded.
6. When you look at other directories, they show the new excludes that you have added. Press "Backup" and back up the files on your /home file space.

Related reference:

"Preview Archive" on page 663

"Preview Backup" on page 664

Include and exclude option processing

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. The server include-exclude statements are always enforced and placed at the bottom of the include-exclude list and evaluated before the client include-exclude statements.

If the `dsm.sys` file include-exclude list contains one or more `incl excl` options that specify include-exclude files, the include-exclude statements in these files are placed in the list position occupied by the `incl excl` option and processed accordingly.

A very large include-exclude list can decrease backup performance. Use wildcards and eliminate unnecessary include statements to keep the list as short as possible.

When performing an incremental backup, the client evaluates all `exclude.fs` and `exclude.dir` statements first, and removes the excluded file spaces, directories, and files from the list of objects available for processing.

After evaluating all `exclude.fs` and `exclude.dir` statements, the client evaluates the include-exclude statements for controlling symbolic link or alias processing (`exclude.attribute.symlink` and `include.attribute.symlink`) from the bottom up and stops if it finds an include or exclude statement that matches the file it is processing. After the include-exclude statements for controlling symbolic link or alias processing are processed, the client evaluates the remaining include-exclude list from the bottom up and stops when it finds an include or exclude statement that matches the file it is processing. The order in which the include and exclude options are entered therefore affects which files are included and excluded.

To display a list of all include-exclude statements in effect on your client workstation in the actual order they are processed, use the **query incl excl** command.

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

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

Note: If a match is not found, files are implicitly included and backed up.

3. When a file is backed up, it is bound to the default management class unless it matched an include statement that specified a different management class name, in which case the file is bound to that management class.

The following examples demonstrate bottom up processing.

Example 1

Assume that La Pomme is not the startup disk.

```
EXCLUDE ../../*.cpp
INCLUDE "/Volumes/La Pomme/Foo/../../*.cpp"
EXCLUDE "/Volumes/La Pomme/Foo/Junk/*.cpp"
```

The file being processed is: /Volumes/La Pomme/Foo/Dev/test.cpp.

Processing follows these steps:

1. Rule 3 (the last include or exclude statement defined) is checked first because of bottom-up processing. The pattern /Volumes/La Pomme/Foo/Junk/*.cpp does not match the file name that is being processed.
2. Processing moves to Rule 2 and checks. This time, pattern /Volumes/La Pomme/Foo/../../*.cpp matches the file name that is being processed. Processing stops, the option is checked, and it is included.
3. File /Volumes/La Pomme/Foo/Dev/test.cpp is backed up.

Example 2

Assume that La Pomme is not the startup disk.

```
EXCLUDE ../../*.cpp
INCLUDE "/Volumes/La Pomme/Foo/../../*.cpp"
EXCLUDE "/Volumes/La Pomme/Foo/Junk/*.cpp"
```

The file being processed is: /Volumes/La Pomme/Widget/Sample File.

Processing follows these steps:

1. Rule 3 is checked and finds no match.
2. Rule 2 is checked and finds no match.
3. Rule 1 is checked and finds no match.
4. Because a match is not found, Volumes/La Pomme/Widget/Sample File is implicitly included and is backed up.

Example 3

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

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

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

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

Example 4

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

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

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

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

Example 5

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

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

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

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

Example 6

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

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

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

1. The exclude.attribute.symlink statement is checked first. If the printf.o file is a symbolic link it is excluded, otherwise proceed to the next step. Note that the exclude.attribute.symlink statements are always processed before the other include-exclude statements, regardless of their position in the include-exclude list.
2. Rule 3 is checked and finds no match.
3. Rule 2 is checked and finds no match.
4. Rule 1 is checked and a match is found.
5. Processing stops, the option is checked, and it is excluded.
6. File /home/lib/objs/printf.o is not backed up.

Related concepts:

“Exclude file spaces and directories” on page 116

Chapter 10, “Processing options,” on page 287

Related reference:

"Exclude options" on page 389
"Query Inclexcl" on page 679

Chapter 3. Getting started

Before you can use the Tivoli Storage Manager backup-archive client, you must learn how to start a GUI or command-line session, and how to start the client scheduler automatically. You can also learn about other commonly used tasks.

Before you use the backup-archive client, complete the following tasks:

- “Starting a Java GUI session” on page 135
- “Starting a command-line session” on page 137
- “Starting a web client session” on page 140
- “Start the client scheduler automatically” on page 141
- “Changing your password” on page 141

You can also complete the following tasks:

- “Sorting file lists using the backup-archive client GUI” on page 143
- “Displaying online help” on page 144
- “Ending a session” on page 144

Configuring security settings to connect the client to the Tivoli Storage Manager server

There are several configuration options that pertain to the Tivoli Storage Manager client security settings when connecting to the Tivoli Storage Manager server V7.1.8 or V8.1.2 or later. Accepting the default values for those options transparently configures the client for enhanced security, and is recommended for most use cases.

Configuring by using the default security settings (fast path)

Fast path details the configuration options that impact the security of the client connection to the server and the behavior for various use cases when default values are accepted. The fast path scenario minimizes the steps in the configuration process at endpoints.

This scenario automatically obtains certificates from the server when the client connects the first time, assuming that the Tivoli Storage Manager server **SESSIONSECURITY** parameter is set to **TRANSITIONAL**, which is the default value at first connection. You can follow this scenario whether you first upgrade the Tivoli Storage Manager server to V7.1.8 and later V7 levels, or V8.1.2 and later V8 levels, and then upgrade the client to these levels, or vice versa.

Attention: This scenario cannot be used if the Tivoli Storage Manager server is configured for LDAP authentication. If LDAP is used, you can manually import the certificates necessary by using the `dsmcert` utility. For more information, see “Configuring without automatic certificate distribution” on page 131.

Client options that affect session security

The following `dsmc` options specify security settings for the client. For more information about these options, see Client options reference.

- **SSLREQUIRED.** The default value `Default` enables existing session-security connections to servers earlier than V7.1.8 or V8.1.2, and automatically configures the client to securely connect to a V7.1.8 or V8.1.2 or later server by using TLS for authentication.
- **SSLACCEPTCERTFROMSERV.** The default value `Yes` enables the client to automatically accept a self-signed public certificate from the server, and to automatically configure the client to use that certificate when the client connects to a V7.1.8 or V8.1.2 or later server.
- **SSL.** The default value `No` indicates that encryption is not used when data is transferred between the client and a server earlier than V7.1.8 or V8.1.2. When the client connects to a V7.1.8 or V8.1.2 or later server, the default value `No` indicates that object data is not encrypted. All other information is encrypted, when the client communicates with the server. The value `Yes` indicates that TLS is used to encrypt all information, including object data, when the client communicates with the server.
- **SSLFIPSMODE.** The default value `No` indicates that a Federal Information Processing Standards (FIPS) certified TLS library is not required.

In addition, the following options apply only when the client uses TLS connection to a server earlier than V7.1.8 or V8.1.2. They are ignored when the client connects to a later server.

- **SSLDISABLELEGACYTLS.** A value of `No` indicates that the client does not require TLS 1.2 for SSL sessions. It allows connection at TLS 1.1 and lower SSL protocols. When the client communicates with a Tivoli Storage Manager server that is V7.1.7 or V8.1.1 or earlier, `No` is the default.
- **LANFREESSL.** The default value `No` indicates that the client does not use TLS when communicating with the Storage Agent when LAN-free data transfer is configured.
- **REPLSSLPORT.** Specifies the TCP/IP port address that is enabled for TLS when the client communicates with the replication target server.

Uses cases for default security settings

- First, the server is upgraded to V7.1.8 or V8.1.2 or later. Then, the client is upgraded. The existing client *is not* using SSL communications:
 - No changes are required to the security options for the client.
 - The configuration is automatically updated to use TLS when the client authenticates with the server.
- First, the server is upgraded to V7.1.8 or V8.1.2 or later. Then, the client is upgraded. The existing client *is* using SSL communications:
 - No changes are required to the security options for the client.
 - SSL communication with existing server public certificate continues to be used.
 - SSL communication is automatically enhanced to use the TLS level that is required by the server.
- First, the client is upgraded to V7.1.8 or V8.1.2 or later. Then, the server is upgraded later. The existing client *is not* using SSL communications:
 - No changes are required to the security options for the client.
 - Existing authentication protocol continues to be used to servers at levels earlier than V7.1.8 or V8.1.2 .
 - The configuration is automatically updated to use TLS when the client authenticates with the server after the server is updated to V7.1.8 or V8.1.2 or later.

- First, the client is upgraded to V7.1.8 or V8.1.2 or later. Then, the server is upgraded later. The existing client *is* using SSL communications:
 - No changes are required to the security options for the client.
 - SSL communication with existing server public certificate continues to be used with servers at levels earlier than V7.1.8 or V8.1.2.
 - SSL communication is automatically enhanced to use the TLS level that is required by the server after the server is updated to V7.1.8 or V8.1.2 or later.
- First, the client is upgraded to V7.1.8 or V8.1.2 or later. Then, the client connects to multiple servers. The servers are upgraded at different times:
 - No changes are required to the security options for the client.
 - The client uses existing authentication and session security protocol to servers at versions earlier than V7.1.8 or V8.1.2 , and automatically upgrade to use TLS authentication when initially connecting to a server at V7.1.8 or V8.1.2 or later. Session security is managed per server.
- New client installation, server is at V7.1.8 or V8.1.2 or later:
 - Configure the client according to a new installation.
 - Default values for the security options automatically configure the client for TLS-encrypted session authentication.
 - Set the SSL parameter to the Yes value if encryption of all data transfers between the client and the server is required.
- New client installation, server is at a version earlier than V7.1.8 or V8.1.2 :
 - Configure the client according to a new client installation.
 - Accept the default values for client session-security parameters if SSL encryption of all data transfers is not required.
 - Non-SSL authentication protocol is used until the server is upgraded to V7.1.8 or V8.1.2 or later.
 - Set the SSL parameter to the Yes value if encryption of all data transfers between the client and the server is required, and proceed with the manual configuration for SSL.
 - See Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer for configuration instructions.
 - SSL communication is automatically enhanced to use the TLS level that is required by the server after the server is updated to V7.1.8 or V8.1.2 or later.

Configuring without automatic certificate distribution

This scenario details the configuration options that impact the security of the client when automatic distribution of certificates from the server is not acceptable. For example, automatic distribution of certificates from the server is not acceptable if the server is configured to use LDAP authentication or it is necessary that certificates are signed by a certificate authority (CA).

Options that affect session security

The options for security settings are the same as those described in “Configuring by using the default security settings (fast path)” on page 129, with the exception that you must set the SSLACCEPTCERTFROMSERV option to No to ensure that the client does not automatically accept a self-signed public certificate from the server when the client first connects to a V7.1.8 or V8.1.2 or later server.

Uses cases for configuring the client without automatic certificate distribution

If automatic certificate distribution is not possible or wanted, use the `dsmcert` utility to import the certificate. Obtain the necessary certificate from the Tivoli Storage Manager server or from a CA. The CA can be from a company such as VeriSign or Thawte, or an internal CA that is maintained within your company.

- First, the server is upgraded to V7.1.8 or V8.1.2. Then, the client is upgraded. The existing client *is not* using SSL communications:
 - Set the `SSLACCEPTCERTFROMSERV` option with the value `No`.
 - Obtain the necessary certificate from the Tivoli Storage Manager server or from a CA and use the `dsmcert` utility to import the certificate. See *Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer* for configuration instructions.
- First, the server is upgraded to V7.1.8 or V8.1.2. Then, the client is upgraded. The existing client *is* using SSL communications:
 - No changes are required to the security options for the client. If the client already has a server certificate for SSL communication, the `SSLACCEPTCERTFROMSERV` option does not apply.
 - SSL communication with existing server public certificate continues to be used.
 - SSL communication is automatically enhanced to use the TLS level that is required by the server.
- First, the client is upgraded to V7.1.8 or V8.1.2. Then, the server is upgraded later. The existing client *is not* using SSL communications:
 - Set the `SSLACCEPTCERTFROMSERV` option with the value `No`.
 - Existing authentication protocol continues to be used to servers at levels earlier than V7.1.8 or V8.1.2.
 - Before the client connects to a V7.1.8 or V8.1.2 or later server:
 - Obtain the necessary certificate from the Tivoli Storage Manager server or from a CA and use the `dsmcert` utility to import the certificate. See *Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer* for configuration instructions.
- First, the client is upgraded to V7.1.8 or V8.1.2. Then, the server is upgraded later. The existing client *is* using SSL communications
 - No changes are required to the security options for the client. If the client already has a server certificate for SSL communication, the `SSLACCEPTCERTFROMSERV` option does not apply.
 - SSL communication with existing server public certificate continues to be used with servers at levels earlier than V7.1.8 or V8.1.2.
 - SSL communication is automatically enhanced to use the TLS level that is required by the server after the server is updated to V7.1.8 or V8.1.2 or later.
- First, the client is upgraded to V7.1.8 or V8.1.2. Then, the client connects to multiple servers. The servers are upgraded at different times:
 - Set the `SSLACCEPTCERTFROMSERV` option with the value `No`.
 - Existing authentication protocol continues to be used to servers at levels earlier than V7.1.8 or V8.1.2.
 - Before the client connects to a V7.1.8 or V8.1.2 or later server, or when SSL communication is required at any server level:
 - Obtain the necessary certificate from the Tivoli Storage Manager server or from a CA and use the `dsmcert` utility to import the certificate. See

Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer for configuration instructions.

- The client uses existing authentication and session security protocol to servers at versions earlier than V7.1.8 or V8.1.2, and automatically upgrade to use TLS authentication when initially connecting to a server at V7.1.8 or V8.1.2 or later. Session security is managed per server.
- New client installation, server is at V7.1.8 or V8.1.2 or later:
 - Configure the client according to a new installation.
 - Set the SSLACCEPTCERTFROMSERV option with the value No.
 - Obtain the necessary certificate from the Tivoli Storage Manager server or from a CA and use the dsmcert utility to import the certificate. See Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer for configuration instructions.
 - Set the SSL parameter to the Yes value if encryption of all data transfers between the client and the server is required.
- New client installation, server is at a version earlier than V7.1.8 or V8.1.2, SSL-encrypted sessions *are* required:
 - Configure the client according to a new installation.
 - Set the SSL parameter to the Yes value.
 - Obtain the necessary certificate from the Tivoli Storage Manager server or from a CA and use the dsmcert utility to import the certificate. See Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer for configuration instructions.
- New client installation, server is at a version earlier than V7.1.8 or V8.1.2, SSL-encrypted sessions *are not* required:
 - Configure the client according to a new installation.
 - Set the SSLACCEPTCERTFROMSERV option with the value No.
 - Non-SSL authentication protocol is used until the server is upgraded to V7.1.8 or V8.1.2 or later.
 - Before the client connects to a V7.1.8 or V8.1.2 or later server:
 - Obtain the necessary certificate from the Tivoli Storage Manager server or from a CA and use the dsmcert utility to import the certificate. See Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer for configuration instructions.

Secure password storage

Beginning in Tivoli Storage Manager Version 7.1.8, the location of the Tivoli Storage Manager password is changed.

In V8.1.0 and V7.1.6 and earlier clients, the Tivoli Storage Manager password was stored in the Windows registry for Windows clients, and stored in the TSM.PWD file on UNIX and Linux clients.

Beginning in V7.1.8, the IBM Global Security Kit (GSKit) keystores are used to store all Tivoli Storage Manager passwords. The process of importing server certificates is simplified. For information about importing server certificates, see “Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer” on page 68.

When you upgrade to Tivoli Storage Manager V7.1.8 backup-archive client, the existing passwords are migrated to the following files in the new password store:

TSM.KDB

The file that stores the encrypted passwords.

TSM.sth

The file that stores the random encryption key that is used to encrypt passwords in the TSM.KDB file. This file is protected by the file system. This file is needed for automated operations.

TSM.IDX

An index file that is used to track the passwords in the TSM.KDB file.

For Data Protection for VMware clients, the Data Protection for VMware GUI server administration password is migrated to a keystore.

Password locations on UNIX and Linux clients

On UNIX and Linux clients, the existing passwords in the TSM.PWD files are migrated to the new password store in the same location. For root users, the default location for the password store is /etc/adsm. For non-root users, the location of the password store is specified by the passworddir option.

The TSM.PWD file is deleted after the migration.

Note: The new password store will not be in the default location (/etc/adsm) in the following situations:

- The TSM.PWD file did not exist in the /etc/adsm directory.
- The options file specifies a passworddir option that points to a different location.

The trusted communications agent is no longer available

The trusted communications agent (TCA), previously used by non-root users in V8.1.0 and V7.1.6 and older clients, is no longer available. Root users can use the following methods to allow non-root users to manage their files:

Help desk method

With the help desk method, the root user runs all backup and restore operations. The non-root user must contact the root user to request certain files to be backed up or restored.

Authorized user method

With the authorized user method, a non-root user is given read/write access to the password store by using the passworddir option to point to a password location that is readable and writable by the non-root user. This method allows non-root users to back up and restore their own files, use encryption, and manage their passwords with the passwordaccess generate option.

For more information, see “Enable non-root users to manage their own data” on page 49.

If neither of these methods are satisfactory, you must use the earlier clients that included the TCA.

Tivoli Storage Manager client authentication

When using the graphical user interface or command line interface of a Tivoli Storage Manager client, you can log on using a node name and password *or* administrative user ID and password.

The client prompts for your user ID and compares it to the configured node name. If they match, the client attempts to authenticate the user ID as a node name. If the authentication fails or if the user ID does not match the configured node name, the client attempts to authenticate the user ID as an administrative user ID.

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

System privilege

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

Policy privilege

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

Client owner

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

Client access

To use the web client to back up and restore files on a remote client system, you must have an administrative user ID with client access authority over the node name for the remote client system. If you do not want Tivoli Storage Manager administrators with client access authority over your node name to be able to back up and restore files on your system, specify the `revokeremoteaccess` option in your client options file.

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

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

- From the main Tivoli Storage Manager GUI window, select **File** → **Connection Information**.
- Use the Tivoli Storage Manager server `QUERY ADMIN` command from the administrative command-line client.

Related reference:

“`Revokeremoteaccess`” on page 502

 `QUERY ADMIN` command

Starting a Java GUI session

The steps that are used to start the backup-archive client graphical interface (GUI) program depend on the operating system.

Procedure

Complete the procedure that is appropriate for your operating system to start the Java GUI.

Operating System	Procedure
Mac OS X	• Double-click the Tivoli Storage Manager application to start the backup-archive client without system administrator privileges. When you run the client without system administrator privileges, you can manage files that are owned by the current user. • Double-click TSM Tools for Administrators and select Tivoli Storage Manager. After you enter a system administrator user name and password, the client starts with system administrator privileges. When you run the client with system administrator privileges, you can manage files that are owned by all users on the system. • You can also start the backup-archive client by using the dsmj command. You can run the client as either a foreground or background process. The dsmj script is installed in /Library/Application Support/tivoli/tsm/client/ba/bin.
AIX, HP-UX, Linux, Solaris	On UNIX systems other than Mac OS X, the backup-archive client GUI must be run from the X Window System. If you see the Tivoli Storage Manager icon on your desktop, the client is already running. Double-click the icon to open the Tivoli Storage Manager window. If the Tivoli Storage Manager icon is not displayed on your desktop, start the backup-archive client graphical interface by using the dsmj command. You can run the client as either a foreground or background process.

The backup-archive client locates and uses the options that are specified in the client system options file (dsm.sys) and the client options files (dsm.opt).

Related concepts:

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

Tivoli Storage Manager password

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

The Tivoli Storage Manager client prompts you for the password if one is required. Contact your Tivoli Storage Manager administrator if you do not know your password.

Related tasks:

“Changing your password” on page 141

Setup wizard

When the client GUI starts, it checks to see whether a client options file exists.

If the client options file does not exist (which usually happens after you have installed the client for the first time on your system), the setup wizard automatically starts and guides you through the configuration process.

The client options file is `dsm.sys`.

Starting a command-line session

You can start a command-line session by invoking the **dsmc** command.

Note: If the `/usr/bin` directory contains a symbolic link to the Tivoli Storage Manager executable, and all DSM environment variables are set, you can enter the **dsmc** command from any directory. Otherwise, enter the fully qualified path of the command.

Note: On Mac OS X, system administrators can use the **sudo** command to gain additional authority so the backup-archive client can access files for all users on the system.

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

Your Tivoli Storage Manager administrator can require you to use a password to connect to the server. The client prompts you for a password, if it is required. Contact your Tivoli Storage Manager administrator if you do not know your password.

Related concepts:

“Options in interactive mode” on page 608

“UNIX and Linux client root and authorized user tasks” on page 47

Chapter 11, “Using commands,” on page 603

Using batch mode

Use *batch* mode to enter a single client command. When you use batch mode, you must precede the command with **dsmc**.

About this task

For example, to issue the **incremental** command, enter the following at the command prompt:

```
dsmc incremental
```

Some commands require one or more arguments. For example, to archive a file:

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

Depending upon the current setting of your `passwordaccess` option, the client might prompt you for your password before the command is processed in a batch mode session.

When you enter your password, the password is not displayed on your screen.

Related reference:

"Passwordaccess" on page 470

Issuing a series of commands by using interactive mode

Use *interactive* mode when you want to issue a series of commands.

About this task

The connection to the server is established only once for interactive mode, so you can process a series of commands more quickly in interactive mode than in batch mode.

To start a client command session in interactive mode, enter either of the following commands:

- `dsmc`
- `dsmc loop`

The following prompt is displayed on your screen:

```
tsm>
```

When you are in interactive mode, do not precede commands with **dsmc**. For example, instead of typing **dsmc archive** to archive a file, type only **archive**.

For example, to archive a file, enter the command with the file specification:

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

Depending upon the current setting of the passwordaccess option, the client might prompt you for your password before you are allowed to enter a command in an interactive session.

When you enter your password, the password is not displayed on your screen.

Specifying input strings that contain blank spaces or quotation marks

You must follow certain rules when you specify an input string that has blanks or quotation marks.

Follow these rules when you specify an input string that has blank spaces or quotation marks:

- If the input string has one or more spaces, enclose the string with either single or double quotation marks. You can use single or double quotation marks, as long as they match.
- If the input string has a single quotation mark, enclose the string within double quotation marks, as in this example:
`-description="Annual backup of the accounting department's monthly reports"`
- If the input string has a double quotation mark, enclose the string within single quotation marks, as in this example:
`-description='New translations of "The Odyssey" and "The Iliad"'`
- If the input string has spaces and quotation marks, enclose the string in quotation marks. The outer quotation marks must not be the same as the quotation marks within the string.

Restriction: An input string that has single and double quotation marks is not a valid input string.

The following rules apply to these types of data:

- Fully qualified names
- The description that you specify in the **archive** command
- Any value for an option value where the character string can include spaces or quotation marks

Important: You cannot use escape characters in input strings. Escape characters are treated the same as any other characters. Here are some examples where escape characters are not recognized:

- If the character string is in an option file
- If the character string is in a list file
- If the character string is entered in interactive mode

Starting: Additional considerations

You can include options as arguments to **dsmj** and **dsmc** commands. For example, you can use options to modify the format that displays dates, times, and numbers, or to include your password so that the backup-archive client does not prompt for it.

About this task

In addition, if you have more than one server defined in `dsm.sys` and you want to contact a different server for backup-archive services (other than the one specified in your client user-options file `dsm.opt`), specify the server with the `servername` option.

For example:

```
dsmj -servername=server_b
```

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

```
dsmj -Xmx512M
```

Using the web client in the new security environment

Beginning with Tivoli Storage Manager Version 7.1.8, you can no longer use the web client to connect to the Tivoli Storage Manager V7.1.8 or later server, or V8.1.2 or later server.

However, you can still use the web client to connect to Tivoli Storage Manager V7.1.7 and earlier servers, or V8.1.1 and earlier V8 servers.

If you are connected to the V7.1.8 or later server or V8.1.2 or later server, use the following alternatives to the web client:

- To restore data that was backed up with the backup-archive client, use the backup-archive client Java GUI (**dsmj**) or the command-line client (**dsmc**). For more information, see:
 - Chapter 4, “Backing up your data,” on page 147

- Chapter 5, “Restoring your data,” on page 221
- To back up and restore NAS file servers using Network Data Management Protocol (NDMP), use the Tivoli Storage Manager server commands on the administrative command-line client (**dsmadm**).

For more information, see the following server documentation:

- Protecting NAS file servers
- Backing up and restoring NAS file servers using NDMP
- File-level backup and restore for NDMP operations

Tip: If you already upgraded the backup-archive client to V7.1.8, you can uninstall it and reinstall the V7.1.6 client to continue to use the web client. The Tivoli Storage Manager server administrator needs to set the SESSIONSECURITY parameter on the node back to the TRANSITIONAL value. For more information, see UPDATE NODE (Update node attributes).

Starting a web client session

After you install the web client on your workstation you can use a browser to backup or restore, or archive or retrieve, data on the node. The web client facilitates the use of assistive devices for users with disabilities and contains improved keyboard navigation.

Before you begin

Refer to the software requirements topic for your operating system to determine which browsers are supported by this software.

To run the web client from Firefox browsers, the option for **Enable JavaScript** must be enabled. By default, this option is enabled in Firefox.

If your browser does not have the correct JRE level, the web client notifies you and if possible, tries to automatically install the correct JRE for you.

Procedure

Specify the URL of the client workstation that you installed the web client on, in your web browser. Also, specify the HTTP port number that is defined on the client workstation for the web client. The default port number is 1581. The following example shows the syntax of a web client URL:

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

If you enter a different URL or click the **Back** button during an operation, the web client is disconnected and the current operation ends.

Related concepts:

“Web client configuration overview” on page 59

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.

The Tivoli Storage Manager server administrator must specify the `userid` parameter with the **REGISTER NODE** server command:

```
REGISTER NODE node_name password userid=user_id
```

where the node name and the administrative user ID must be the same. For example:

```
REGISTER NODE node_a mypassw0rd userid=node_a
```

Tip: You can use the `revokeremoteaccess` option to prevent a Tivoli Storage Manager administrator with client access privilege from performing client operations on your workstation through the web client. However, Tivoli Storage Manager administrators with client owner privilege, system privilege, or policy privilege can still perform client operations on your workstation through the web client.

Related concepts:

“Tivoli Storage Manager client authentication” on page 134

Related reference:

“Revokeremoteaccess” on page 502

Start the client scheduler automatically

You can start the client scheduler automatically when you start your workstation.

If the Tivoli Storage Manager administrator has defined schedules for your node, starting the client scheduler permits you to automatically back up your workstation (or perform other scheduled actions).

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

Related tasks:

“Setting the client scheduler process to run as a background task and start automatically at startup” on page 264

Changing your password

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

About this task

The backup-archive client prompts you for the password if one is required. Contact your Tivoli Storage Manager administrator if you do not know your password.

Important: The password discussed in this topic is different than the password used for encrypting files.

To change your password from the GUI:

Procedure

1. On Mac OS X clients, start the backup-archive client with TSM Tools for Administrators.
2. From the main window, open the **Utilities** menu and select **Change password**.

3. Enter your current and new passwords, and enter your new password again in the **Verify password** field.
4. Click **Change**.

Results

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

For UNIX, Linux, and Windows clients:

```
dsmc set password
```

For Mac OS X clients, enter this command to change your password from the command-line client:

```
sudo dsmc set password
```

Then, enter your old and new passwords when prompted.

Passwords can be up to 63 character in length. Password constraints vary, depending on where the passwords are stored and managed, and depending on the version of the Tivoli Storage Manager server that your client connects to.

If your Tivoli Storage Manager server is at version 6.3.3 or later, and if you use an LDAP directory server to authenticate passwords

Use any of the following characters to create a password:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9 . ! @ # $ % ^ & * _ - + = ` ( )
| { } [ ] : ; < > , ? / ~
```

Passwords are case-sensitive and are subject to more restrictions that can be imposed by LDAP policies.

If your Tivoli Storage Manager server is at version 6.3.3 or later, and if you do not use an LDAP directory server to authenticate passwords

Use any of the following characters to create a password:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9 . ! @ # $ % ^ & * _ - + = ` ( )
| { } [ ] : ; < > , ? / ~
```

Passwords are stored in the Tivoli Storage Manager server database and are not case-sensitive.

If your Tivoli Storage Manager server is earlier than version 6.3.3

Use any of the following characters to create a password:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9
_ - & + .
```

Passwords are stored in the Tivoli Storage Manager server database and are not case-sensitive.

Remember:

On the command line, enclose all parameters that contain one or more special characters in quotation marks. Without quotation marks, the special characters can

be interpreted as shell escape characters, file redirection characters, or other characters that have significance to the operating system.

On AIX, HPUX, Linux, Mac, and Solaris systems:

Enclose the command parameters in single quotation marks (').

Command line example:

```
dsmc set password -type=vmguest 'Win 2012 SQL'  
'ttml2dag\administrator' '7@#$$%^&7'
```

Quotation marks are not required when you type a password with special characters in an options file.

Related concepts:

“Start the client scheduler automatically” on page 141

Related tasks:

“Starting: Additional considerations” on page 139

Related reference:

“Password” on page 468

“Set Password” on page 734

Sorting file lists using the backup-archive client GUI

You can use the backup-archive client GUI to display, sort, or select files.

About this task

Table 32. Working with your files using the backup-archive client GUI

Task	Procedure
Displaying files	To display files in a directory, click the folder icon next to the directory name. The files appear in the File List box on the right.
Sorting the file list	Click the appropriate column heading in the File List box.
Display active and inactive backup versions	- Click the Display Active/Inactive Files option from the View menu. - Click the Display both active and inactive files tool on the tool bar.
Display only active backup versions	Click the Display active files only option from the View menu.
Selecting files to restore or retrieve.	- Click the selection box next to the directory or file name that you want to restore or retrieve. - Highlight the files that you want to restore or retrieve and click the Select Items tool on the tool bar. - Highlight the files that you want to restore or retrieve and click the Select Items option from the Edit menu.
Deselecting files	- Click the checked selection box next to the directory or file name. - Highlight the files that you want to deselect and click the Deselect Items tool on the tool bar. - Highlight the files that you want to deselect and click the Deselect Items option from the Edit menu.
Displaying file information	- Highlight the file name, and click the View File Details button on the tool bar. - Highlight the file name, and select File Details from the View menu.

Note:

1. Unless otherwise noted, the tasks and procedures in the above table apply to all client GUIs.
2. Using the client GUIs, you can sort a list of files by various attributes, such as name, directory, size, or modification date. Sorting files by the last backup date can be useful in determining what date and time to use for the point-in-time function.
3. An *active* file is the most recent backup version of a file that existed on your workstation when you ran your last backup. All other backup versions of that file are *inactive*. Only active backup versions of files are displayed, unless you select the **Display active/inactive files** menu option. If you delete the file from your workstation, the active version becomes inactive the next time you run an incremental backup.

On the command-line client, you can use **query** commands with the **inactive** option to display both active and inactive objects. You can use **restore** commands with the **pick** and **inactive** options to produce the list of active and inactive backups to choose from.

Related reference:

“Inactive” on page 419

“Pick” on page 472

Displaying online help

You can display online help in any of the following ways: On the backup-archive client GUI, from the web client, or from the **dsmc** command line.

About this task

- On the backup-archive client GUI:
 - Open the help menu. Click **Help** or press F1.
 - Click the **Help** button in the current window.
 - On Mac systems, click the GUI question mark (?) icon, which displays online information about the current operation.
- From the **dsmc** command line: Enter the **help** command. The complete table of contents for the available help text is displayed.

On Mac systems, for the GUI-based help, the help pages contain HTML style hyper-links to other topics related to the current operation you are performing. You can select a topic of interest in the Ask field and search the online help for matches. Search results are sorted by relevance. Use the left and right arrow keys in the lower right-hand corner to navigate between help pages.

Related reference:

“Help” on page 652

Ending a session

You can end a client session from the backup-archive client GUI or from the **dsmc** command line.

About this task

- From the backup-archive client GUI:
 - Open the **File** menu and select **Quit**.
 - Press Command+Q.

- Open the **File** menu and select **Exit**.
- Open the **System** menu and select **Close**.
- For the web client: Open a different URL or close the browser.
- From the DSMC command line:
 - In batch mode, each **dsmc** command you enter is a complete session. The client ends the session when it finishes processing the command.
 - To end an interactive session, enter **quit** at the **tsm>** prompt.
 - To interrupt a **dsmc** command before the client has finished processing, enter **QQ** on the Tivoli Storage Manager console. In many cases but not all, this interrupts the command. If the command cannot be interrupted, use the UNIX **kill -9** command from an available command prompt. Do not press Ctrl-C or use the UNIX **kill -15** command because it can lead to unexpected results.

Related reference:

“**Loop**” on page 660

Online forums

To participate in user discussions of Tivoli Storage Manager products, you can subscribe to the ADSM-L list server.

About this task

This is a user forum maintained by Marist College. While not officially supported by IBM, product 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>

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.

All client backup and restore procedures also apply to the web client.

Restriction: The web client does not provide a Preferences Editor for setting client options.

The following is a list of primary backup tasks.

- “Planning your backups”
- “Pre-backup considerations (UNIX and Linux)” on page 149
- “Performing an incremental, selective, or incremental-by-date backup (UNIX and Linux)” on page 160
- “Deleting backup data” on page 173
- “Backing up files from one or more file spaces for a group backup (UNIX and Linux)” on page 174
- “Image backup” on page 181
- “Back up NAS file systems using Network Data Management Protocol” on page 192

Planning your backups

If you are a first-time user, or if you only back up files occasionally, you can use the table in this topic as a checklist of preliminary steps to consider before backing up data.

Read the list of tasks to determine whether you are ready to back up your data.

- Decide whether you want to back up files or archive them. See “When to back up and when to archive files” on page 148 for more information.
- See “Pre-backup considerations (UNIX and Linux)” on page 149 for important considerations before you back up your files and directories.
- Do you need to exclude files from backup services? See “Include-exclude options to control processing” on page 152 for more information.

Related concepts:

Chapter 1, “Installing the Tivoli Storage Manager backup-archive clients,” on page 1

Which files are backed up

When you request a backup, the client backs up a file if certain requirements are met.

To back up a file, the client must meet the following are the requirements:

- The selected management class contains a backup copy group.

- The file meets the serialization requirements that are defined in the backup copy group. If the copy group serialization parameter is `static` or `shrstatic`, and the file changes during backup, the file is not backed up.
- The file meets the **mode** requirements that are defined in the backup copy group. If the copy group **mode** parameter is `modified`, the file must have changed since the last backup. If the **mode** is `absolute`, the file can be backed up even if it does not change.
- The file meets the frequency requirements that are defined in the backup copy group. The specified minimum number of days since the last backup must elapse before a file is backed up.
- The file is not excluded from backup by an exclude statement.
- The file is not excluded from backup by the operating system. These excluded files can be found in registry subkey `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup`.

Files that are part of the Windows system state are eligible for backup only when the system state is backed up. You can back up the system state only as a single entity because of dependencies among the system state components. You cannot back up or restore the files individually. For example, because `C:\windows\system32\ntoskrnl.exe` is part of the Windows system state, it is not backed up during an incremental or selective backup of the `C:\` drive.

Related concepts:

Chapter 9, “Storage management policies,” on page 275

“Management classes and copy groups” on page 276

Related reference:

“Absolute” on page 316

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, backup and archive operations have different results.

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 by using predetermined policy-based rules. Using these rules, the Tivoli Storage Manager administrator can control the following processes:

- The number of versions that are maintained on the Tivoli Storage Manager server
- The number of days each additional backup copy is kept
- What happens to backup versions when the file is deleted on the client system

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

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

Tip: If a file is archived multiple times by using the same archive description, a new copy of the file is added to the archive each time that archive is operation run. To simplify the retrieve operation, store only one copy of a file in each archive.

Backups protect against file damage or loss that can occur through accidental deletion, corruption, or disk crashes. The server maintains one or more backup versions for each file that you back up. Older versions are deleted as newer versions are made. The number of backup versions the server maintains is set by your administrator.

Archive copies are saved for long-term storage. Your administrator can limit how long archive copies are kept. The server can store an unlimited number of archive versions of a file. Archives are useful if you must go back to a particular version of your files, or you want to delete a file from your workstation and retrieve it later, if necessary. For example, you might want to save spreadsheets for tax purposes, but because you are not using them, you do not want to leave them on your workstation.

Related concepts:

Chapter 6, “Archive and retrieve your data (UNIX and Linux),” on page 249
“Restore data from a backup set” on page 224

Pre-backup considerations (UNIX and Linux)

Various factors in your system or environment can affect the way the backup-archive client processes data. Review these considerations before you back up your data.

LAN-free data movement

LAN-free data movement shifts the movement of client data from the communications network to a storage area network (SAN). This decreases the load on the Tivoli Storage Manager server.

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

AIX, HP-UX, Linux and Solaris clients support LAN-free data movement.

LAN-free prerequisites

To enable LAN-free support, you must install and configure the Tivoli Storage Manager for Storage Area Networks storage agent on the client workstation.

IBM Tivoli Storage Manager for Storage Area Networks is a separate product.

For more information about installing and configuring the storage agent, see the documentation for Tivoli Storage Manager for Storage Area Networks.

LAN-free data movement options

To enable LAN-free data movement, you can use several client options. You must first install and configure the IBM Tivoli Storage Manager for Storage Area Networks storage agent on the client workstation.

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.

lanfreecommmethod

Specifies a communication protocol between the client and the Storage Agent.

lanfreeshmport

Specifies the unique number that is used by the client and the storage agent to identify shared memory area used for communications.

lanfreetcpport

Specifies the TCP/IP port number where the Storage Agent is listening.

lanfreetcpserveraddress

Specifies the TCP/IP address for the storage agent.

Related reference:

"Enablelanfree" on page 382

"Lanfreecommmethod" on page 441

"Lanfreeshmport" on page 442

"Lanfreessl" on page 444

"Lanfreetcpport" on page 443

"Lanfreetcpserveraddress" on page 444

Incremental backups on memory-constrained systems

Incremental backup performance suffers if the system has a low amount of memory available before starting the backup.

If your system is memory constrained, specify the *memoryefficientbackup yes* option in your client options file. This option causes the backup-archive client to process only one directory at a time, which reduces memory consumption but increases backup time. When you specify *yes*, the client analyzes only one directory at a time for backup consideration. If performance remains poor, check your communication buffer settings and the communication link between your system and the Tivoli Storage Manager server. If your system is not memory constrained, setting the *memoryefficientbackup* option to *yes* degrades your backup performance.

Related reference:

"Memoryefficientbackup" on page 452

Incremental backups on systems with a large number of files

The client can use large amounts of memory to run 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 client requires, on average, approximately 300 MB of memory. The exact amount of memory that is used per object varies, depending on the length of the object path and name length, or the nesting depth of directories. The number of bytes of data is not an important factor in determining the backup-archive 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 that is needed per object.

The total memory requirement can be reduced by any of the following methods:

- Use the client option **memoryefficientbackup diskcachemethod**. This choice reduces the use of memory to a minimum at the expense of performance and a significant increase in disk space that is required for the backup. The file description data from the server is stored in a disk-resident temporary database, not in memory. As directories on the workstation are scanned, the database is consulted to determine whether to back up, update, or expire each object. At the completion of the backup, the database file is deleted.
- Use the client option **memoryefficientbackup yes**. The average memory that is used by the client then becomes 300 bytes times the number of directories plus 300 bytes per file in the directory that is being processed. For file systems with large numbers (millions) of directories, the client still might not be able to allocate enough memory to perform incremental backup with **memoryefficientbackup yes**.
- UNIX and Linux clients might be able to use the **virtualmountpoint** client option to define multiple virtual mount points within a single file system, each of which can be backed up independently by the client.
- If the client option **resourceutilization** is set to a value greater than 4, and multiple file systems are being backed up, then reducing **resourceutilization** to 4 or lower limits the process to incremental backup of a single file system at a time. This setting reduces the memory requirement. If the backup of multiple file systems in parallel is required for performance reasons, and the combined memory requirements exceed the process limits, then multiple instances of the backup client can be used to back up multiple file systems in parallel. For example, if you want to back up two file systems at the same time but their memory requirements exceed the limits of a single process, then start one instance of the client to back up one of the file systems, and start a second instance of the client to back up the other file system.
- Use the - **incrbydate** client option to perform an "incremental-by-date" backup.
- Use the **exclude.dir** client option to prevent the client from traversing and backing up directories that do not need to be backed up.
- Except for Mac OS X, use the client image backup function to back up the entire volume. An image backup might actually use less system resources and run faster than incremental backup of some file systems with a large number of small files.
- Reduce the number of files per file system by spreading the data across multiple file systems.

Related reference:

"Snapdiff" on page 519

"Exclude options" on page 389

"Incrbydate" on page 436

"Memoryefficientbackup" on page 452

"Resourceutilization" on page 499

"Virtualmountpoint" on page 566

Include-exclude options to control processing

You might have files in your file systems that you do not want to back up. These files might be core files, local caches of network file systems, operating system or application files that could be easily recovered by reinstalling the program, or any other files that you could easily rebuild.

You can use the `exclude` and `include` options in your include-exclude options list to specify which files to exclude from backup processing.

Use the `include` and `exclude` options in `dsm.sys` to define which files to include or exclude from incremental or selective backup processing. A file is eligible for backup unless excluded by an `exclude` option. It is not necessary to use an `include` option to include specific files for backup unless those files are in a directory containing other files you want to exclude.

Tivoli Storage Manager uses management classes to determine how to manage your backups on the server. Every time you back up a file, the file is assigned a management class. The management class is either a default chosen for you, or one you assign to the file using the `include` option in the include-exclude list. If you assign a management class, it must contain a backup copy group for the file to be backed up.

Related tasks:

“Creating an include-exclude list” on page 115

“Setting the client scheduler process to run as a background task and start automatically at startup” on page 264

Data encryption during backup or archive operations

The way to ensure data security is by encrypting data. Use data encryption to protect data during a backup or archive operation. Advanced Encryption Standard (AES) 128-bit encryption is the default encryption option. For the highest level of data encryption, use 256-bit Advanced Encryption Standard (AES) data encryption by specifying the **encrypttype** option.

The data that you include is stored in encrypted form, and encryption does not affect the amount of data that is sent or received.

The **include.encrypt** option is the only way to enable encryption on the backup-archive client. If no `include.encrypt` statements are used encryption cannot occur.

Encryption is not compatible with VMware virtual machine backups that use the incremental forever backup modes (**MODE=IFIncremental** and **MODE=IFFull**). If the client is configured for encryption, you cannot use incremental forever backup. However, you can use the full or incremental backup modes (**MODE=Full** and **MODE=Incremental**).

Use the **include** and **exclude** options in `dsm.sys` to define which files to include or exclude from incremental or selective backup processing. A file is eligible for backup unless excluded by an **exclude** option. It is not necessary to use an **include** option to include specific files for backup unless those files are in a directory that contains other files that you want to exclude.

To encrypt file data, you must select an encryption key password, which the client uses to generate the encryption key for encrypting and decrypting the file data.

Store the encryption key password for later use. You can specify whether to save the encryption key password in a file that is named `TSM.sth` by using the **encryptkey** option.

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

While restoring the encrypted file, the client prompts you for the key password to decrypt the file in the following cases:

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

Related reference:

“Encryptiontype” on page 383

“Encryptkey” on page 384

“Exclude options” on page 389

“Include options” on page 421

File system and ACL support

Special file systems contain dynamic information that is generated by the operating system; they contain no data or files. The UNIX and Linux clients ignore special file systems and their contents.

Special file systems include the following types:

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

The backup-archive client can work on specific file system types that are commonly used. For a list of supported file system types, see Table 33.

Restriction: The table shows full support for NFS on AIX, including preservation of ACLs and extended attributes. On other operating systems, NFS backups are supported, but the backups include only standard POSIX metadata (access permissions, creation date, and so on). For more information about backing up NFS file systems, see “Backup network file systems” on page 197.

Table 33. Supported file systems and ACL support

Platform	File System	ACL Support
AIX	GPFS	yes
	JFS	yes
	JFS2	yes
	JFS2 NFSV4	yes
	VxFS	yes
HP-UX	HFS	no
	VxFS (JFS Veritas)	yes (V3.3.Layout 4)

Table 33. Supported file systems and ACL support (continued)

Platform	File System	ACL Support
Linux x86_64	Btrfs	yes
	XFS	yes
	EXT2	yes
	EXT3	yes
	EXT4	yes
	ReiserFS GPFS	yes
	JFS	yes
	VxFS	no
	NSS	no
Linux on Power Systems Servers	Btrfs	yes
	XFS	yes
	EXT2	yes
	EXT3	yes
	EXT4	yes
	ReiserFS	yes
	JFS GPFS	no
		yes
Linux on z Systems	Btrfs	yes
	EXT2	yes
	EXT3	yes
	EXT4	yes
	ReiserFS	yes
	JFS	no
	GPFS	yes
Mac	HFS Standard (HFS)	yes
	HFS Extended (HFS+)	yes
	HFS Extended case-sensitive (HFSX)	yes
	Xsan (XSAN) UNIX file system (UFS)	yes
	Universal disk format (UDF)	yes
	ISO9660	yes
Solaris	UFS	yes
	VxFS	yes
	QFS	no
	ZFS	yes

With file systems where NFS V4 ACLs are defined and used (Solaris ZFS and AIX JFS2 V2), even if only the standard UNIX permissions or ACLs have changed (such as with the CHMOD command), the file or directory is fully backed up again. With other file systems, this type of change causes only an attribute update on the Tivoli Storage Manager server.

To process all other file systems, use the `virtualmountpoint` option to enable support for the following items:

- To back up, restore, archive, and retrieve file data
- For basic UNIX and Linux permissions
- For change, access, and modification time stamps, and the directory tree structure

No other file system specific attributes, such as the ACL, are valid. The file system type for such file systems is set to "UNKNOWN".

For example, if the /media/abc/DATA1 file system is not supported by the client, add the following statement to dsm.sys to back up or archive the data in this file system:

```
VIRTUALMOUNTPPOINT /media/abc/DATA1
```

This support is only available if the file system can use basic POSIX system calls, such as read or write processing on your system.

Cross-platform backup and restore are not supported. For example, data backed up by an AIX client is not available for restore by a Windows client and vice versa.

Note: Data that is backed up or archived by the Mac OS X client cannot be restored by any other client. Additionally, the Mac OS X client cannot restore or retrieve data from any other client.

You can use the cross-file system type restore or retrieve method for ACL information if both the original file system and the destination file system support compatible ACLs. For example, on Solaris, the ACL information that is backed up from a VxFS file system is restored to a UFS file system because these file systems support compatible ACLs. The ACL information is not restored during cross-file system restore or retrieve operations if the original file system and the destination file system do not support ACLs,

The stand-alone package LSCqfs 3.5.0 is the only supported version of QFS. In addition, the following restrictions also apply to the QFS file system:

- Image backup is not supported on QFS file systems.
- The Solaris backup-archive client does not support the combination of QFS and SAM needed to archive files onto tertiary background storage, such as tapes. Instead, it recalls files from tape to disk automatically if it finds migrated files during a backup.
- A QFS file system contains two hidden system files and a system directory that cannot be backed up; and this is acceptable because a backup of these files is not needed. They contain internal data to manage the file system. The internal data is automatically excluded from a backup and is re-created automatically by the file system itself, if a restore of files in that file system is completed.

Incremental, selective, filelist back up, archive, restore, and retrieve processing of the Veritas file system and its ACLs on AIX are supported. Restore of a Veritas volume on a Logical Volume Manager volume (and vice versa) is allowed, provided both have the same file system type.

The following information pertains only to Mac OS X systems:

- On Mac OS X systems, the UFS and HFSX file systems are case-sensitive whereas the HFS+ file system is not case-sensitive but is case-preserving. Files that you back up from a UFS or HFSX file system (case-sensitive) might not be restored properly to an HFS+ file system (not case-sensitive) file system. For example, on a UFS file system, files Afile and afile are seen as different files. However, on an HFS+ file system the two files are seen as identical.
- On Mac OS X, if case-sensitive HFS+ or UFS file systems are used, it is important that the data from the HFSX or UFS file system is not backed up to an HFS+ file system on the Tivoli Storage Manage server. Either a new name must be used on the system or the existing file space on the Tivoli Storage Manager server must be renamed. For example, consider a system that has a file system named /Volumes/fs2 and this system is repartitioned with a case-sensitive HFS+ file system. Either the /Volumes/fs2 file system on the Tivoli Storage Manager

server must be renamed, or a new name must be used on the local system. If this renaming is not done, the HFSX case-sensitive data is mixed with the HFS+ case-insensitive data that is already stored on the Tivoli Storage Manager server.

- On Mac OS X, aliases and symbolic links are backed up. However, the client does not back up the data to which the symbolic links point.
- On Mac OS X, when files that are backed up from an HFS volume are restored to a UFS volume, the resource forks are not assigned to the correct owner. Correct this problem by using the **chown** command on the resource fork file to change the owner. The resource fork file stores structured data in a file.

On Linux on POWER and Linux on System z, you must install `libacl.so` for the client to back up ACLs.

Important: If you are running GPFS for AIX, GPFS for Linux x86_64, or GPFS for Linux on z Systems in a multinode cluster, and all nodes share a mounted GPFS file system, the client processes this file system as a local file system. The client backs up the file system on each node during an incremental backup. To avoid this, you can do one of the following things:

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

If the GPFS cluster contains different platforms, you must use backup-archive clients on only one platform to protect a single file system. Do not use backup-archive clients on more than one platform to protect a GPFS file system that is shared among more than one platform.

For example, assume that a cluster contains nodes on AIX, Linux x86, and Linux zSeries systems. You can protect file system A with AIX backup-archive clients and protect file system B with Linux zSeries backup-archive clients. Or you can protect file system A and file system B with AIX backup-archive clients. If you protect file system A with an AIX backup-archive client, you must not protect file system A with a backup-archive client on any platform other than AIX.

Support for cross operating system recovery for files stored in IBM Spectrum Scale™ file systems

In an IBM Spectrum Scale cluster with multiple operating system types, a file that holds ACL or extended attribute metadata and was backed up on a source operating system, can be restored on a target operating system. The ACL or extended attribute metadata is correctly restored correctly if both operating system types on the source and the target use the same version of IBM Spectrum Scale.

The following are the supported source-operating-systems types:

- AIX
- Linux for IBM System Power big endian (pBE)
- Linux x86
- Linux for IBM System z

The following are the supported target-operating-system types:

- Linux for IBM System Power little endian (pLE)
- Linux x86

- Linux for IBM System z

The security settings for affected users and groups must match on both the source and the target systems.

Do not mix operating system types for backup activity. Choose only one operating system type available in your IBM Spectrum Scale cluster, and use it for all backup operations.

Maximum file size for operations

The maximum file size depends on the type of a file system. The backup-archive client does not check any file size limit during backup, archive, restore, or retrieve operations.

If the file system allows creation of the file, the client backs up or archives the file.

The following table specifies the maximum file sizes for the native file systems on UNIX and Linux client platforms.

Table 34. Maximum file size

Platform	Max file size (in bytes)
AIX 6.1 (JFS2) size limitations	Maximum JFS2 file system size: 32 TB Maximum JFS2 file size: 16 TB Minimum JFS2 file system size: 16 MB
HP-UX	1 099 511 627 775 (1 TB-1)
All Linux clients	9 223 372 036 854 775 807 (8 EB-1)
Mac OS X	HFS - 2 147 485 648 (2GB) HFS+, HFSX, XSAN, and UFS - 9 223 372 036 854 775 808 (8EB)
Solaris	1 099 511 627 775 (1 TB-1)
Solaris (ZFS)	18 446 744 073 709 551 616 (16 EB)

Long user and group names

Tivoli Storage Manager can handle user and group names that are up to 64 characters without any issues. However, names longer than 64 characters require special handling by Tivoli Storage Manager.

Important: Do not exceed the 64 character limit for user and group names. If you do, Tivoli Storage Manager shortens the name to fall within this limit by using the following transformation: Take the first 53 characters, append a forward slash (/), and then the numeric ID as a character string.

An error message is logged that contains both the long name and the resulting shortened string. For most functions, you do not need to be aware of the shortened name. The exceptions are:

- The **set access** command
- The **fromowner** option
- The **users and groups (authorization)** options

In each of these cases, when you need to enter a name, you either have to find the error message containing the transformation, or construct the name using the rule outlined here.

Mac OS X volume names

The backup-archive client backs up volumes based on their UNIX mount point name.

Tivoli Storage Manager maintains each volume name as a separate restore or retrieve volume. These volume names become the names of file spaces on the server.

If you change the name of a volume you have already backed up, the client sees it as a new volume and does not relate it to the previous one. Any backup of the volume backs up the files under the new name. A mismatch might occur if you rename your volumes, or if you access Tivoli Storage Manager from a different workstation than the one from which you backed up the files.

Mac OS X volume naming precautions

Tivoli Storage Manager creates all new file spaces on the server with the UNIX mount point of the volume.

If there are two volumes with the names such as "La Pomme" and "la pomme", two unique UNIX mount points are created.

The following examples show the two mount points that are created:

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

If duplicate volumes exist on your desktop, it is possible for the UNIX mount points to be different than the last time the client did a backup. The client might not back up the data to the correct file system on the Tivoli Storage Manager server.

You can check the file system where the client backs up the data:

1. In the TSM Backup window, select a file system.
2. Click **File** → **Show Info**.

The UNIX mount point is in the Information dialog.

The best way to avoid any potential naming problems is to ensure that the volume names are unique.

Important:

- The client continues to use the existing file space names on the Tivoli Storage Manager Server. Only new file spaces use the UNIX mount point for the name.
- Do not specify volumes with periods in the name (...). The client uses the sequence of periods as part of include-exclude processing. The client reports an invalid include-exclude statement if a volume has a sequence of periods in the name. The volume *must* be renamed.

Mac OS X volume naming precautions on dual boot systems

If you have more than one version of Mac OS X that you switch between, it is critical that you understand how the client uses the UNIX mount paths for file space names on the Tivoli Storage Manager server.

For example, consider a dual-boot system that has two volumes, Panther and Tiger. The finder and the backup-archive client GUI displays these as Panther and Tiger. However, the UNIX mount points depend upon which version of Mac OS is running. If Panther is the startup disk, the UNIX paths are:

```
/
/Volumes/Tiger
```

If Tiger is the startup disk, the UNIX paths are:

```
/
/Volumes/Panther
```

When a backup or archive operation is run, the file space names also depend on which version of Mac OS X is running.

Both versions of Mac OS X back up to the / file system on the Tivoli Storage Manager server. When this happens, the system files are intermixed.

To avoid potential problems on dual-boot systems, complete one of these tasks:

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

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

If the volume Panther is the startup disk, add this to the system options file:

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

Mac OS X Unicode enablement

The Mac OS X client is Unicode enabled. New clients storing data on the server for the first time require no special set up.

The server automatically stores files and directories as Unicode enabled. However, if you are upgrading to the Unicode-enabled client, you need to plan the migration of existing file spaces so they can support Unicode.

Any file spaces that are already on the server must be renamed so Unicode-enabled file spaces can be created. Use the `autofsrename` option rename existing file spaces.

Related reference:

“Autofsrename” on page 326

Mac OS X Time Machine backup disk

Time Machine is the backup application available with Mac OS X.

Tivoli Storage Manager can be used at the same time as Mac OS X Time Machine application. However, due to the unique nature of how the Mac OS X Time Machine application backs up data, consider the following items before using the backup-archive client to back up the Mac OS X Time Machine data:

- The Mac OS X Time Machine backup disk makes extensive use of both file and directory hard links to minimize disk usage. For example, if the disk backed up

with the Mac OS X Time Machine application is 5 GB, the first backup copies 5 GBs of data to the Mac OS X Time Machine backup disk.

Subsequent backups only copy the files that have changed since the previous backup. All files and directories that have not changed are hard-linked with the version that was copied during the previous backup.

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

All hard-linked objects that are not already on the Tivoli Storage Manager server are backed up.

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

- When files that are restored are hard-linked, the client recreates the original hard link. Recreating the original hard link can only be done if *all* files that are hard-linked are restored at the same time. Restoring all the hard-linked files at the same time is not a practical method for a large backup disk that uses the Mac OS X Time Machine application.
- When the Mac OS X Time Machine application copies files to the backup disk, ACLs are added to the files to protect them from deletion. the backup-archive can back up and restore files with ACLs. However, any files that are restored must have these restrictive ACLs in place.

Tip: For best results, exclude the Time Machine application backup data. All Time Machine application data is in a directory named `Backups.backupdb`.

Related concepts:

“System files to exclude” on page 119

Performing an incremental, selective, or incremental-by-date backup (UNIX and Linux)

Your administrator might have set up schedules to automatically back up files on your workstation. The following sections discuss how to back up files without using a schedule.

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

Related tasks:

“Setting the client scheduler process to run as a background task and start automatically at startup” on page 264

Full and partial incremental backup

An incremental backup backs up only new and changed files. The type of incremental backup depends on what objects you select to be backed up.

If you select entire file systems, the backup is a full incremental backup. If you select a directory tree or individual files, the backup is a partial incremental backup.

The first time that you run a full incremental backup, the backup-archive client backs up all the files and directories that you specify. The backup operation can take a long time if the number of files is large, or if one or more large files must be backed up. Subsequent full incremental backups only back up new and changed

files. The backup server maintains current versions of your files without having to waste time or space by backing up files that exist in Tivoli Storage Manager server storage.

Depending on your storage management policies, the Tivoli Storage Manager server might keep more than one version of your files in storage. The most recently backed up files are active backup versions. Older copies of your backed up files are inactive versions. However, if you delete a file from your workstation, the next full incremental backup causes the active backup version of the file to become inactive. You can restore an inactive version of a file. The number of inactive versions that are maintained by the server and how long they are retained is governed by the management policies that are defined by your Tivoli Storage Manager server administrator. The active versions represent the files that existed on your file system at the time of the last backup.

To start a full or partial incremental backup by using the client GUI, select **Backup**, and then select the **Incremental (complete)** option. From the command line, use the **incremental** command and 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 for the following tasks:

- Back up new files.
- Back up files whose contents changed since the last backup.

Files are backed up when any of the following attributes change:

- File size
- Date or time of last modification
- Extended Attributes
- Access Control List

If only the following attributes change, the attributes are updated on the Tivoli Storage Manager server, but the file is not backed up:

- File owner
- File permissions
- Inode
- Group ID
- Change time (ctime) attribute (for objects in GPFS file systems only and if the **updatectime** option is set to yes). For more details, see the **updatectime** option.
- Icon location (Mac OS X only)
- Type or creator (Mac OS X only)

- Back up directories.

A directory is backed up in any of the following circumstances:

- The directory was not previously backed up
- The directory permissions changed since the last backup
- The directory Access Control List changed since the last backup
- The directory Extended Attributes changed since the last backup
- The change time (ctime) attribute is updated since the last backup (for GPFS file systems only). For more details, see the **updatectime** option.

Directories are counted in the number of objects that are backed up. To exclude directories and their contents from backup, use the **exclude.dir** option.

- Expire backup versions of files on the server that do not have corresponding files on the workstation. The result is that files that no longer exist on your

workstation do not have active backup versions on the server. However, inactive versions are retained according to rules defined by the Tivoli Storage Manager administrator.

- Rebind backup versions if management class assignments change. Only objects that have active backup versions are bound again. Objects for which only inactive backup versions exist are not bound again.

During a partial incremental backup operation, objects are rebound or expired as follows:

If the file specification matches all files in a path:

Rebinding and expiration occurs for all eligible backup versions that match the file specification. This is the case for an incremental command like `dsmc incr c:\mydir\* -subdir=yes`.

If the file specification does not match all files in a path:

Rebinding and expiration occurs for all eligible backup versions that match the file specification. However, eligible backup versions are not expired or rebound if they were in a directory that no longer exists on the client file system.

Consider an incremental command like `dsmc incr c:\mydir\*.txt -subdir=yes`. Assume that some files in `c:\mydir\` do not have the `.txt` file type. Rebinding and expiration occurs only for files that match the `*.txt` specification and whose directories still exist on the client file system.

You can use the `preserveleastaccessdate` option to specify whether to modify the last access date after a backup or archive operation. By default, the access date changes after a backup or archive operation.

Related concepts:

Chapter 9, "Storage management policies," on page 275

Related reference:

"Exclude options" on page 389

"Preserveleastaccessdate" on page 479

"Updatectime" on page 557

Journal-based backup on AIX and Linux

Journal-based backup is an alternate method of backup that uses a change journal maintained by the Tivoli Storage Manager journal daemon process.

On AIX, journal-based backup is supported on JFS and JFS2 file systems.

On Linux, journal-based backup is supported on Ext2, Ext3, Ext4; XFS, ReiserFS, JFS, VxFS, and NSS, and for a local file system shared through NFS. GPFS is not supported for journal-based backups.

To support journal-based backup you must install and configure the Tivoli Storage Manager journal daemon.

A backup of a particular file system will be journal-based when the journal daemon has been installed and configured to journal the particular file system, and a valid journal has been established for the file system. Journal-based backup is enabled by successfully completing a full incremental backup.

The primary difference between traditional incremental backup and journal-based backup is the method used for backup and expiration candidates.

Traditional incremental backup obtains the list of backup and expiration candidates by building comprehensive lists of local objects, and lists of active server objects for the file system being backed up. The local lists are obtained by scanning the entire local file system. The server list is obtained by querying the entire server inventory for all active objects.

The two lists are compared, and candidates are selected according to the following criteria:

- An object is selected as a backup candidate if it exists in the local list, but does not exist in the server list. The object is also a backup candidate if it exists in both lists, but differs according to incremental criteria (for example, attribute changes, date and size changes).
- An object is selected as an expiration candidate if it exists in the server list, but doesn't exist in the local list.

Journal-based backup obtains the candidates list of objects to back up and expire by querying the journal daemon for the contents of the change journal of the file system being backed up.

Change journal entries are cleared (marked as free) after they have been processed by the backup client and committed on the Tivoli Storage Manager server.

You can use journal-based backup when backing up file systems with small or moderate amounts of change activity between backup cycles. If you have many file changes between backup cycles, you will have very large change journals. Large change journals might create memory and performance problems that can negate the benefits of journal-based backup. For example, creating, deleting, renaming, or moving very large directory trees can also negate the benefit of using journal-based backup instead of normal incremental backup.

Journal-based backup is not intended to be a complete replacement for traditional incremental backup. You should supplement journal-based backup with a full progressive incremental backup on a regular basis. For example, perform journal-based backups on a daily basis, and full incremental backups on a weekly basis.

Here are some limitations of journal-based backup:

- Individual server attributes are not available during a journal-based backup. Certain policy settings such as copy frequency and copy mode might not be enforced.
- Other platform-specific behaviors might prevent objects from being processed properly. Other software that changes the default behavior of the file system might prevent file system changes from being detected.
- If the file system is very active when a journal-based backup is in progress, it is possible that a small number of deleted files will not be expired.
- If you restore files to a file system that has an active journal, some of the restored files might get backed up again when the next journal-based backup occurs, even if the files have not changed since they were restored.
- The “`Skipaclupdatecheck`” on page 518 option has no effect during journal-based backups. No matter how this option is set, when performing journal-based backups of a file system, the client always backs up a file if its ACL or extended attributes have been changed since the last backup.

You should perform periodic full incremental backups and more frequent journal backups. Traditional incremental backup compares the entire server inventory of files against the entire local file system. Therefore, incremental backup is always the most comprehensive backup method.

Note: 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.

Restore processing with journal-based backups (AIX and Linux):

The journal service attempts to identify changes that are made to a file as the result of a restore operation. If a file is unchanged since it was restored, it is not backed up again during the next journaled backup. The presumption is that you are restoring a file because it contains the data you need, so there is no point to backing up the file again when the next journal backup occurs. Changes to restored files that occur after the files are restored must be recognized as new changes and the file is processed in the next journal backup.

When an active journal exists for a particular file system, the backup-archive client notifies the journal daemon when a file is about to be restored. Any changes to the file that occur within a short window in time after the journal daemon is notified are assumed to be a result of the file being restored. These changes are not recorded and the file is not included in the next journal backup.

In most cases, journal processing correctly identifies file changes that are generated as the result of the file being restored and prevents the file from being backed up by the next journal backup.

Systemic system delays, whether caused by intensive I/O or file system latency, might prevent a restore operation from starting in the time frame allotted by the journal daemon once it is notified that a restore is about to take place. If such a delay occurs, changes made to the file are assumed to be new changes that occurred after the file was restored. These changes are recorded, and the file is included in the next journal backup. Things like systemic processing delays and file system latency are beyond the control of the backup-archive client and are simply recognized limitations of journal-based backups.

Incremental-by-date backup

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

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

Files that were renamed after the last incremental backup, but otherwise remain unchanged, will not be backed up. Renaming a file does not change the

modification date and time of the file. However, renaming a file does change the modification date of the directory in which it is located. In this case, the directory is backed up, but not the files it contains.

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

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

Related tasks:

“Backing up data using the Java GUI” on page 169

Comparing incremental-by-date, journal-based, and NetApp snapshot difference to full incremental and partial incremental backups

Incremental-by-date, journal-based, and NetApp snapshot difference are alternatives to full incremental and partial incremental back methods.

Incremental-by-date backup

An incremental-by-date backup takes less time to process than a full incremental backup and requires less memory.

An incremental-by-date backup might not place exactly the same backup files into server storage because the incremental-by-date backup:

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

Journal-based backup

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 1 KB) 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.

On AIX and Linux, journal-based backup does have some limitations. See “Journal-based backup on AIX and Linux” on page 162 for information.

NetApp snapshot difference

For NAS and N-Series file servers that are running ONTAP 7.3.0, or later, you can use the `snapdiff` option to invoke the snapshot difference backup from NetApp when running a full-volume incremental backup. Using this option reduces memory usage and is faster.

Consider the following restrictions when running a full-volume incremental backup using the `snapdiff` option, to ensure that data is backed up when it should be.

- A file is excluded due to an exclude rule in the include-exclude file. The client runs a backup of the current snapshot with that exclude rule in effect. This happens when you have not made changes to the file, but you have removed the rule that excluded the file. NetApp will not detect this include-exclude change because it only detects file changes between two snapshots.
- If you added an include statement to the option file, that include option does not take effect unless NetApp detects that the file has changed. The client does not inspect every file on the volume during backup.
- If you used the **`dsmc delete backup`** command to explicitly delete a file from the Tivoli Storage Manager inventory, NetApp cannot detect that a file was manually deleted from Tivoli Storage Manager storage. Therefore, the file remains unprotected in Tivoli Storage Manager storage until it is changed on the volume and the change is detected by NetApp, which signals the client to back it up again.
- Policy changes such as changing the policy from **`mode=modified`** to **`mode=absolute`** are not detected.
- The entire file space is deleted from the Tivoli Storage Manager inventory. This action causes the `snapdiff` option to create a new snapshot to use as the source, and a full incremental backup to be run.

The NetApp software determines what is a changed object, not Tivoli Storage Manager.

If you run a full volume backup of an NFS-mounted NetApp or N-Series volume, all the snapshots under the snapshot directory might also be backed up.

To avoid backing up all snapshots under the snapshot directory, do one of the following actions:

- Run NDMP backups
- Run backups using the `snapshotroot` option
- Run incremental backups using the `snapdiff` option

Tip: If you run an incremental backup using the `snapdiff` option and you schedule periodic incremental backups, use the `createnewbase=yes` option with the `snapdiff` option to create a base snapshot and use it as a source to run an incremental backup.

- Exclude the snapshot directory from backups.

On AIX and Linux systems, the snapshot directory is in `.snapshot`.

Note: The `.snapshot` directory is not backed up for some versions of Red Hat Linux, so you are not required to exclude it.

Snapshot differential backup with an HTTPS connection

You can use a secure HTTPS connection for the backup-archive client to communicate with a NetApp filer during a snapshot differential backup.

The HTTPS protocol is enabled on NetApp filers by default and cannot be disabled.

When you run a snapshot differential backup, the backup-archive client establishes an administrative session with a NetApp filer. The filer credentials, such as the filer host name or IP address, the user name that is used to connect to the filer, and the filer password, are stored locally on the backup-archive client. This information must be transmitted to the filer to establish the authenticated administrative session. It is important to use a secure connection because authenticating the administrative filer session requires the client to transmit the filer password in clear text.

To establish a secure connection by using the HTTPS communication protocol, you must use the **snappdiffhttps** option whenever you run a snapshot differential backup. Without the **snappdiffhttps** option, the backup-archive client can establish filer sessions only with the HTTP protocol, which would require HTTP administrative access to be enabled on the filer. With the **snappdiffhttps** option, you can establish a secure administrative session with the NetApp filer regardless of whether HTTP administrative access is enabled on the NetApp filer.

Restrictions:

The following restrictions apply to snapshot differential backups with HTTPS:

- The HTTPS connection is used only to securely transmit data over the administrative session between the backup-archive client and the NetApp filer. The administrative session data includes information such as filer credentials, snapshot information, and file names and attributes that are generated by the snapshot differencing process. The HTTPS connection is not used to transmit normal file data that is accessed on the filer by the client through file sharing. The HTTPS connection also does not apply to normal file data transmitted by the client to the Tivoli Storage Manager server through the normal Tivoli Storage Manager client/server protocol.
- The **snappdiffhttps** option does not apply to vFilers because the HTTPS protocol is not supported on the NetApp vFiler.
- The **snappdiffhttps** option is available only by using the command-line interface. It is not available for use with the backup-archive client GUI.

Related tasks:

“Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 106

“Running a snapshot differential backup with an HTTPS connection”

Related reference:

“Snappdiffhttps” on page 524

“Snappdiff” on page 519

Running a snapshot differential backup with an HTTPS connection

When you run a snapshot differential backup, you can use the **snappdiffhttps** option to create a secure HTTPS connection between the backup-archive client and the NetApp filer.

Before you begin

Before you begin a snapshot differential backup over an HTTPS connection, ensure that you configured the client as described in “Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 106.

This method is available only at the command-line interface.

Procedure

To start a snapshot differential backup operation over an HTTPS connection, specify the **incremental** command with the **snappdiff** and **snappdiffhttps** options at the command-line interface.

For example, you are using an AIX or Linux system with an NFS mounted file system `/vol/vol1` hosted on the file server `homestore.example.com`. The `/net/home1` directory is the mount point of `/vol/vol1`. Issue the following command:

```
dsmc incr /net/home1 -snappdiff -snappdiffhttps
```

Related concepts:

“Snapshot differential backup with an HTTPS connection” on page 167

Related reference:

“Snappdiffhttps” on page 524

Selective backup

Use a selective backup when you want to back up specific files or directories regardless of whether a current copy of those files exists on the server.

Incremental backups are generally part of an automated system to back up entire file systems. In contrast, selective backups allow you to manually select a set of files to back up regardless of whether they have changed since your last incremental backup.

Unlike incremental backups, a selective backup provides the following:

- Does not cause the server to update the date and time of the last incremental.
- Backs up directory and file entries even if their size, modification timestamp, or permissions have not changed.
- Does not expire deleted files.
- Does not rebind backup versions to a new management class if you change the management class.

Related tasks:

“Backing up data using the Java GUI” on page 169

Related reference:

“Selective” on page 725

Solaris global zone and non-global zones backups

For Solaris zones, perform incremental and selective backups of file systems within the zone where these file systems were created.

Treat each non-global zone as a separate system that has its own Tivoli Storage Manager node name and run backups from within each of the zones.

If you run incremental or selective backups of non-global zones from the global zone, the global-zone administrator must decide which files in the non-global zone are included or excluded in backups. For example, device, system and kernel files

of the non-global zones are not automatically excluded from backups, but they must not be backed up. Restoring such files can make a non-global zone unusable.

Saving access permissions

When you back up your files, the backup-archive client also saves standard UNIX access permissions assigned to the files.

Depending on your operating system, it also saves extended permissions. For example, for files on an AIX workstation, the client saves access control lists.

It is possible for an authorized user to back up files for another user, but this should not cause ownership conflicts. The backup server properly records that the file belongs to the original owner. The authorized user does not need to grant the original owner access to the backup versions.

Setting a virtual mount point

If you are an authorized user and you want to back up files beginning with a specific directory within a file system, you can define that directory as a virtual mount point.

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

Related reference:

“Virtualmountpoint” on page 566

Backing up data using the Java GUI

You can back up specific files, entire directories, or entire file systems from the directory tree.

About this task

You can locate the files you want to back up by searching or filtering. Filtering displays only the files matching the filter criteria for your backup.

Use the backup-archive client Java GUI to back up your data as follows:

Procedure

1. Click **Backup** in the Tivoli Storage Manager window. The Backup window appears.
2. Expand the directory tree if necessary. Click on the selection boxes next to the object or objects you want to back up. To search or filter files, click the **Find** icon on the tool bar.
3. Enter your search criteria in the Find Files (Backup) window.
4. Click the **Search** button. The Matching Files (Backup) window appears.
5. Click the selection boxes next to the files you want to back up and close the Matching Files (Backup) window.
6. Enter your filter criteria in the Find Files (Backup) window.
7. Click the **Filter** button. The Backup window displays the filtered files.

8. Click the selection boxes next to the filtered files or directories you want to back up.
9. Select one of the following backup types from the pull-down menu: (1) To run an incremental backup, click **Incremental (complete)**, (2) To run an incremental-by-date backup, click **Incremental (date only)**, (3) To run a selective backup, click **Always backup**.
10. Click **Backup**. The Backup **Task List** window displays the backup processing status.

Results

Consider the following items when you back up your data using the Java GUI.

- To modify specific backup options, click the **Options** button. The options you select are effective during the current session *only*.
- Tivoli Storage Manager uses management classes to determine how to manage your backups on the server. Every time you back up a file, the file is assigned a management class. The management class used is either a default selected for you, or one that you assign to the file using an *include* option in the include-exclude options list. Select **Utilities** → **View Policy Information** from the backup-archive client Java GUI or web client GUI to view the backup policies defined by the Tivoli Storage Manager server for your client node.
- To perform an automatic incremental backup of your default domain, select **Actions** → **Backup Domain**. Your default domain is set with the *domain* option in your client user-options file (dsm.opt). If you do not have the *domain* option set, the default domain is *all local file systems*.
- You can use the Preferences editor to exclude file systems in your default domain from backup processing.

Related concepts:

Chapter 9, “Storage management policies,” on page 275

Related reference:

“Domain” on page 361

Backing up data using the command line

You can use the **incremental** or **selective** commands to perform backups.

The following table shows examples of using these commands to perform different tasks.

Table 35. 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 654 for more information about the incremental command.
Back up the /fs1 and /fs2 file systems in addition to the /home, /usr, and /datasave file systems defined in your client domain.	dsmc incremental -domain="/fs1 /fs2"	See “Domain” on page 361 for more information about the domain option.

Table 35. Command-line backup examples (continued)

Task	Command	Considerations
Back up the /Volumes/fs1 and /Volumes/fs2 file systems in addition to the volumes defined in your client domain.	<code>dsmc incremental -domain="/Volumes/fs1 /Volumes/fs2"</code>	See "Domain" on page 361 for more information about the domain option.
Back up all local file systems defined in your client domain except for the /home file system.	<code>dsmc incremental -domain="all-local -/home"</code>	You cannot use the (-) operator in front of the domain keyword all-local. See "Domain" on page 361 for more information. For Windows clients, you can also exclude the system state domain from backup processing in this way.
Back up only the /fs1 and /fs2 file systems.	<code>dsmc incremental /fs1 /fs2</code>	None
Back up all files in the /home directory and all its subdirectories.	<code>dsmc incremental /home/ -subdir=yes</code>	See "Subdir" on page 539 for more information about the subdir option.
Back up all files in the /Users directory and all its subdirectories.	<code>dsmc incremental /Users/ -subdir=yes</code>	See "Subdir" on page 539 for more information about the subdir option.
Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, run an incremental backup of all files and directories under the local snapshot and manage them on the Tivoli Storage Manager server under the file space name /usr.	<code>dsmc incremental /usr -snapshotroot=/snapshot/day1</code>	The backup-archive client considers the snapshotroot value as a file space name. See "Snapshotroot" on page 528 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 incrbydate option with the incremental command to back up new and changed files with a modification date later than the last incremental backup stored at the server. See "Incrbydate" on page 436 for more information about the incrbydate option.
<i>Selective backups</i>		
Back up all files in the /home/proj or /Users/van/Documents directory.	<code>dsmc selective /home/proj/ or dsmc selective /Users/van/Documents/</code>	Use the selective command to back up specific files or directories regardless of whether they have changed since your last incremental backup. You can use wildcards to back up multiple files at once. See "Selective" on page 725 for more information about the selective command.

Table 35. Command-line backup examples (continued)

Task	Command	Considerations
Back up all files in the /home/proj directory and all its subdirectories.	<code>dsmc selective /home/proj/ -subdir=yes</code>	If you specify <code>-subdir=yes</code> when backing up a specific path and file, the client recursively backs up all subdirectories under that path, and any instances of the specified file that exist under any of those subdirectories. If a subdirectory is a mounted file system, the client does not back up the files in that subdirectory when you use the <code>subdir=yes</code> option. See “Subdir” on page 539 for more information about the <code>subdir</code> option.
Back up all files in the /Users/van/Documents directory and all its subdirectories.	<code>dsmc selective /Users/van/Documents/ -subdir=yes</code>	If you specify <code>-subdir=yes</code> when backing up a specific path and file, the client recursively backs up all subdirectories under that path, and any instances of the specified file that exist under any of those subdirectories. If a subdirectory is a mounted file system, the client does not back up the files in that subdirectory when you use the <code>subdir=yes</code> option. See “Subdir” on page 539 for more information about the <code>subdir</code> option.
Back up the /home/dir1/h1.doc and /home/dir1/test.doc files.	<code>dsmc selective /home/dir1/h1.doc /home/dir1/test.doc</code>	If you specify the <code>removeoperandlimit</code> option with the incremental or selective commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits. This allows you to specify more than 20 files on a single command. See “Removeoperandlimit” on page 488 for more information about this option.
Back up the /Users/ann/Documents/h1.doc and /Users/ann/Documents/test.doc files.	<code>dsmc selective /Users/ann/Documents/h1.doc /Users/ann/Documents/test.doc</code>	If you specify the <code>removeoperandlimit</code> option with the incremental or selective commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits. This allows you to specify more than 20 files on a single command. See “Removeoperandlimit” on page 488 for more information about this option.
Back up a list of files in the /home/filelist.txt file.	<code>selective -filelist=/home/filelist.txt</code>	Use the <code>filelist</code> option to process a list of files. See “Filelist” on page 402 for more information.
Back up all files listed in the /Users/filelist.txt file.	<code>dsmc selective -filelist=/Users/filelist.txt</code>	Use the <code>filelist</code> option to process a list of files. See “Filelist” on page 402 for more information.

Table 35. Command-line backup examples (continued)

Task	Command	Considerations
Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, run a selective backup of the /usr/dir1/sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name /usr.	<code>dsmc selective /usr/dir1/sub1/ -subdir=yes -snapshotroot=/snapshot/ day1</code>	The client considers the snapshotroot value as a file space name. See “Snapshotroot” on page 528 for more information.

Related reference:

“Incremental” on page 654

“Selective” on page 725

Deleting backup data

If your administrator has given you authority, you can delete individual backup copies from the Tivoli Storage Manager server without deleting the entire file space. To determine if you have this authority, select **File > Connection Information** from the backup-archive client GUI or web client main menu. Your authority status is provided in the **Delete Backup Files** field.

About this task

Important: When you delete backup files, *you cannot restore them*. Verify that the backup files are no longer needed before you delete them. The client prompts whether you want to continue with the delete. If you specify *yes*, the specified backup files are immediately deleted and removed from Tivoli Storage Manager server storage.

To delete backup copies using the backup-archive client GUI or web client:

Procedure

1. Select **Delete Backup Data** from the **Utilities** menu. The Backup Delete window appears.
2. Expand the Directory tree by clicking the plus sign (+) or folder icon next to the object you want to expand.
3. Click the selection boxes next to objects that you want to delete.
4. Select an item from the drop-down list near the top of the **Backup Delete** window to specify the type of backup delete to perform. You can delete active backup versions, inactive backup versions, or all objects that you have selected in the tree.

Results

Note:

1. A directory is deleted only if you select **Delete All Objects**.
2. To delete backup copies using the command line client, use the **delete backup** command.

Related reference:

“Delete Backup” on page 645

Deleting file spaces

If your Tivoli Storage Manager administrator gives you authority, you can delete entire file spaces from the server. When you delete a file space, you delete all the files and images, both backup versions and archive copies, that are contained within the file space. For example, if you delete the /tmp file space, you are deleting every backup for every file in that file system and every file you archived from that file system. Carefully consider whether you want to delete a file space.

About this task

You can also delete a file space using the **delete filespace** command. Use the **class** option with the **delete filespace** command to delete NAS file spaces.

You can delete individual backup versions by using the **delete backup** command.

You can delete file spaces using the backup-archive client GUI or command line clients. To delete NAS file spaces, use the web client or command line client.

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

Procedure

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

Related reference:

“Class” on page 333

“Delete Backup” on page 645

“Delete Filespace” on page 648

Backing up files from one or more file spaces for a group backup (UNIX and Linux)

You can use the **backup group** command to create and back up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.

Restriction: The **backup group** command does not apply to Mac OS X.

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

- All objects in the group are assigned to the same management class.
- Existing *exclude* statements for any files in the group are ignored.
- All objects in the group are exported together.
- All objects in the group are expired together as specified in the management class. No objects in a group are expired until all other objects in the group are expired, even when another group they belong to gets expired.

A group backup can be added to a backup set.

You can perform a full or differential backup using the mode option.

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

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

Related concepts:

“Restore data from a backup set” on page 224

Related reference:

“Backup Group” on page 618

“Include options” on page 421

“Mode” on page 454

Backing up data with client-node proxy support (UNIX and Linux)

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

About this task

Consolidating backups from multiple nodes to a common target node name on the server is helpful in configurations where the workstation that is responsible for performing the backups can change over time, such as within a cluster.

An agent node is a client node that is 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.

Use the asnodename option with the appropriate command to back up, archive, restore, and retrieve data under the target node name on the server.

The asnodename option also allows data to be restored from a different system than the one that performed the backup.

Consider the following features when you use a proxy node to back up or restore data on other nodes:

- A proxy operation uses the settings for the target node (such as **maxnummp** and **deduplication**) and schedules that are defined on the Tivoli Storage Manager server. The Tivoli Storage Manager server node settings and schedules for the agent node are ignored.
- All of the agent nodes in the multiple node environment must be running the same operating system type.
- Do not use target nodes as traditional nodes, especially if you encrypt your files before you back them up to the server.
- You cannot access another node (either from the GUI drop-down or by using the fromnode option).
- You cannot perform NAS backup or restore.

Procedure

1. Install the backup-archive client on all nodes in a shared data environment.

2. Register each node with the Tivoli Storage Manager server. Register the common target node name to be shared by each of the agent nodes that are used in your shared data environment.
3. Register each of the nodes in the shared data environment with the Tivoli Storage Manager server. Register the agent node name that is used for authentication purposes. Data is not stored on the server, under that node name, when the `asnodename` option is used.
4. The Tivoli Storage Manager server administrator must grant proxy authority to all nodes in the shared environment to access the target node name by using the `GRANT PROXYNODE` command.
5. Use the `QUERY PROXYNODE` administrative client command to display the client nodes of the authorized user that was granted by the `GRANT PROXYNODE` command.

Related reference:

"Asnodename" on page 319

Enabling multiple node operations from the GUI

To enable multinode operations in the GUI, use the Preferences editor to specify the name of the target node to which you have been granted proxy authority.

Procedure

1. Verify that the client node has proxy authority to a target node (or authorized to act as the target node) by using the `QUERY PROXYNODE` administrative client command.
2. Select **Edit > Client Preferences** to open the preferences window.
3. Select the **General** tab and fill in the **As Node Name** field with the name of the target node.
4. Click **Apply** and then **OK** to close the preferences window.

What to do next

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

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

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

Setting up encryption

This topic lists the steps that you must follow to set up encryption with the `encryptkey` option.

Procedure

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.

Results

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

1. Specify **encryptkey=prompt** in the options file.
2. Ensure that users of the agent nodes in the multiple node environment are using the same encryption key.

Important:

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

Scheduling backups with client-node proxy support

Multiple nodes can be used to perform backup operations using the scheduler.

About this task

When you grant proxy authority to the agent nodes, they perform scheduled backup operations on behalf of the target node. Each agent node must use the `asnodename` option within their schedule to perform multiple node backup for the agent node.

Start the schedules using `dsmc sched client` command:

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

- The administrator registers all of the nodes to be used by issuing the following commands:
 - `register node NODE-A`
 - `register node NODE-B`
 - `register node NODE-C`
- The administrator grants proxy authority to each agent node using the following commands:
 - `grant proxynode target=NODE-Z agent=NODE-A`
 - `grant proxynode target=NODE-Z agent=NODE-B`
 - `grant proxynode target=NODE-Z agent=NODE-C`
- The administrator defines the schedules using the following commands:
 - `define schedule standard proxy1 description="NODE-A proxy schedule" action=incremental options="-asnode=NODE-Z" objects=/Volumes/Xsan1 startdate=05/21/2005 starttime=01:00`
 - `define schedule standard proxy2 description="NODE-B proxy schedule" action=incremental options="-asnode=NODE-Z" objects=/Volumes/Xsan2 startdate=05/21/2005 starttime=01:00`
 - `define schedule standard proxy3 description="NODE-C proxy schedule" action=incremental options="-asnode=NODE-Z" objects=/Volumes/Xsan3 startdate=05/21/2005 starttime=01:00`

Note: Place the `asnodename` option in the schedule definition only. Do not place it in the client options file, on the command line, or in any other location.

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

Note:

- Each schedule can be started from a different workstation or LPAR.
- After running the schedules, any proxied client can query and restore all of the backed up data.
- A proxy operation uses the settings for the target node (such as **maxnummp** and **deduplication**) and schedules that are defined on the Tivoli Storage Manager server. The Tivoli Storage Manager server node settings and schedules for the agent node are ignored.

Related reference:

"Asnodename" on page 319

"Session settings and schedules for a proxy operation" on page 320

 **DEFINE SCHEDULE** command

Examples of how to schedule a backup of an IBM PowerHA SystemMirror cluster

This section shows lists some examples of how to back up an IBM PowerHA SystemMirror cluster.

About this task

Perform the following steps to enable scheduling of multiple nodes:

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

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

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

In the following examples, IBM PowerHA SystemMirror is configured for two AIX hosts, **host_a** and **host_b**. Along with their own local data, the hosts are sharing disk storage which has two file spaces: **/disk1** and **/disk2**.

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

- The administrator defines 3 nodes on the Tivoli Storage Manager server: **host_a**, **host_b**, **cluster_group**, using the following commands: (1) **REGISTER NODE host_a mysecretpa5s**, (2) **REGISTER NODE host_b mysecretpa5s**, (3) **REGISTER NODE cluster_group mysecretpa5s**.
- The administrator defines a **dsm.opt** file on **host_a** and **host_b** (note that the opt files are different on each host), using the following commands: (1) **NODENAME host_a** (option can be left as default), (2) **DOMAIN /home /usr ... etc..**
- The administrator defines a **dsm.opt** file located somewhere on one of the cluster disk groups, for example, **/disk1/tsm/dsm.opt**, using the following commands: (1) **NODENAME cluster_group**, (2) **DOMAIN /disk1 /disk2**, (3) **CLUSTERNODE YES**.
- The administrator defines a schedule on the Tivoli Storage Manager server, using the following command: **DEFINE SCHEDULE STANDARD CLUSTER_BACKUP**.
- The administrator defines associations for each of the 3 nodes, using the following command: **DEFINE ASSOC STANDARD CLUSTER_BACKUP host_a,host_b,cluster_group**. At any one time, there are three instances of the

backup-archive client schedule running (with the scheduler for cluster_group being part of the cluster resources that failover whenever the cluster group disk resources failover. Thus, it would be running on either host_a or host_b but not both simultaneously).

- All three node names contain data on the Tivoli Storage Manager server.

The ASNODE example shows a generic solution which could be applied to UNIX cluster solutions to which we do not have support, for example: Veritas Cluster Server for Solaris.

- The administrator defines 3 nodes on the Tivoli Storage Manager server host_a, host_b, cluster_group:

```
REGISTER NODE host_a mysecretpa5s
REGISTER NODE host_b mysecretpa5s
REGISTER NODE cluster_group mysecretpa5s
```
- The administrator defines a proxy node relationship between host_a and host_b to hacmp_cluster

```
GRANT PROXYNODE TARGET=cluster_group AGENT=host_a,host_b
```
- The administrator defines a dsm.opt file on host_a and host_b to handle the local file systems:

```
NODENAME host_a (option can be left as default)
DOMAIN /home /usr ... etc.

NODENAME host_b (option can be left as default)
DOMAIN /home /usr ... etc.
```
- The administrator defines a dsm.opt file on the cluster resource to handle the backup of the clustered resources, e.g. /disk1/tsm/dsmcluster.opt (the nodename is the default nodename, which is either host_a or host_b, depending on which workstation contains the cluster group at any given time):

```
DOMAIN /disk1 /disk2
ASNODE cluster_group
```
- The administrator defines a schedule on the Tivoli Storage Manager server:

```
DEFINE SCHEDULE STANDARD CLUSTER_BACKUP
```
- The administrator defines associations for each one of the 3 nodes.

```
DEFINE ASSOC STANDARD CLUSTER_BACKUP host_a,host_b,cluster_group
```
- At any one time, there are three instances of the backup-archive client schedule running with the scheduler for node hacmp_cluster running on either host_a or host_b but not both (it is included in the cluster resources that would failover). This scheduler would point to the dsmcluster.opt that is defined on each host. The three instances would be started as:

```
[host_a] dsmc sched
[host_b] dsmc sched
[cluster_group] dsmc sched -optfile=/disk/tsm/dsmcluster.opt
```
- All three node names contain data on the Tivoli Storage Manager server.

For more information about the server scheduler commands, see the server documentation.

Scheduling a backup of a GPFS file system

Use the scheduler and proxy relationships to back up a GPFS file system.

About this task

Assume that three nodes in a GPFS cluster participate in the backup operation. Nodes node_1, node_2, and node_3 are used for authentication only. The objects

are backed up to file spaces that belong to node `node_gpfs`.

Procedure

1. Define four nodes on the Tivoli Storage Manager server.

```
REGISTER NODE node_1 mysecretpa5s  
REGISTER NODE node_2 mysecretpa5s  
REGISTER NODE node_3 mysecretpa5s  
REGISTER NODE node_gpfs mysecretpa5s
```
2. Define a proxy relationship between the nodes.

```
GRANT PROXYNODE TARGET=node_gpfs AGENT=node_1, node_2, node_3
```
3. Define a schedule.

```
DEFINE SCHEDULE STANDARD GPFS_SCHEDULE ACTION=incremental  
OBJECTS="/gpfs"  
DEFINE ASSOCIATION STANDARD GPFS_SCHEDULE node_gpfs
```
4. Choose one of the GPFS systems to run the schedule. Specify the **nodename** and **asnodename** options in the `dsm.sys` options file on all systems in the GPFS cluster. The value for the **asnodename** option must be the same on all systems.

Definitions in the `dsm.sys` options file on node 1:

```
nodename node_1  
asnodename node_gpfs
```

Definitions in the `dsm.sys` options file on node 2:

```
nodename node_2  
asnodename node_gpfs
```

Definitions in the `dsm.sys` options file on node 3:

```
nodename node_3  
asnodename node_gpfs
```

5. Start the scheduler on the system that is chosen to run the schedule.

```
DSMC SCHED
```

Related information:

-  `mmbackup` command: Tivoli Storage Manager requirements
-  Guidance for integrating IBM Spectrum Scale AFM with IBM Spectrum Protect
-  Using IBM Spectrum Protect include and exclude options with IBM Spectrum Scale `mmbackup` command

Associate a local snapshot with a server file space (UNIX and Linux)

Use the `snapshotroot` option with the **incremental** and **selective** commands in conjunction with an independent software vendor application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server.

The `snapshotroot` option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot.

Related reference:

“`Snapshotroot`” on page 528

Image backup

From your local workstation, you can back up a logical volume as a single object (image backup) on your system.

The traditional static image backup prevents write access to the volume by other system applications during the operation.

You must be a root user to perform this task, and image backup does not apply to Mac OS X.

An image backup provides the following benefits:

- Backs up file systems that contain a large number of files faster than a full file system incremental backup.
- Improves the speed with which the client restores file systems that contain 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 must recall that information.
- Restores a corrupted file system or raw logical volume. Data is restored to the same state it was when the last logical volume backup was performed.

The traditional static image backup prevents write access to the volume by other system applications during the operation. Use the `dynamicimage` option to back up the volume as is, without remounting it read-only. Corruption of the backup can occur if applications continue to write to the volume while the backup is running. Writing to a volume while an image backup is running can result in inconsistent data and data loss after a restore operation is run. The `dynamicimage` option overrides the copy serialization value in the management class to perform an image backup. After restoring an image backup taken with the `dynamicimage` option, always run the `chkdsk` utility.

To restore an image backup of a volume, the backup-archive client must be able to obtain an exclusive lock on the volume that is being restored.

Restriction: Do not use dynamic image backups for file systems, because the file system might provide inconsistent data even when there is no write activity. Also, dynamic image backup might result in a fuzzy image, which might not be valid or complete when restored.

If the backup-archive client fails to mount the file system after it restores an image, run **fsck**. However, running **fsck** can affect the integrity of large amounts of data. Do not use dynamic image backup for AIX JFS2 file systems. The client does not allow dynamic image backup for AIX JFS2 file systems. If you specify `dynamicimage=yes` for a JFS2 file system, the client performs a snapshot-based image backup. If the snapshot cannot be created for some reason, the client instead performs a static image backup.

Attention: To prevent data loss, avoid using the `dynamicimage` option, and ensure that there is no write activity on the volume while the backup is in progress.

For AIX JFS2 file systems, the amount of data that is backed up to the Tivoli Storage Manager server during static or snapshot image backup is reduced by backing up only those blocks used by the file system or smaller than the

imagegapsize option. This method of backing up your data improves the performance of image backup. For more information, see “Imagegapsize” on page 417.

For AIX clients only: By default, the client performs an online snapshot image backup of JFS2 file systems, during which the volume is available to other system applications.

For Linux clients only: By default, the client performs a snapshot image backup of file systems that exist on a logical volume that is created by the Linux Logical Volume Manager. The volume is available to other system applications while the snapshot image backup is performed.

For Linux clients: Image backup of DASD devices with raw-track access mode on Linux on z Systems™ is not supported.

Attention: File systems that are managed by Tivoli Storage Manager for Space Management are not enabled for image backup.

Related tasks:

“Snapshot-based file backup and archive and snapshot-based image backup” on page 188

Performing prerequisite tasks before creating an image backup

This topic lists some items to consider before you perform an image backup.

About this task

The following items are the image backup considerations.

- Ensure that no other application is using the volume when you run a static image backup. To ensure a consistent image during backup processing, if a file space is detected on the volume the client unmounts and remounts the volume as read only, so that no other applications can write to it. If the volume is in use when the client attempts to unmount, the backup fails. If the client cannot unmount and remount the volume as read only because it is in use, and snapshot image backup is not available, you can use the `dynamicimage` option to force the client to perform an image backup without unmounting and remounting the volume in read-only mode. Set the `dynamicimage` option in an `include.image` statement or from the command line. The backup can be corrupted if applications write to the volume while the backup is in progress. This can be corrected by running `fsck` after a restore to fix any corrupted blocks.

If no file system is detected on the volume being backed up, ensure that all applications writing to the volumes are quiesced. The backup-archive client uses the file system table and mount table to detect the supported file systems.

Do not include system files in an image backup because file systems being actively used cannot be unmounted.

For AIX and Linux only: If you perform an image backup of a mounted file system which is mounted to another mount point and specified in the file system table, then after completing the image backup, all mount options for this file system, except read or write state, is lost.

Important: If a mounted file system has nested mount points, unmount them before attempting a backup. Otherwise, the client is unable to unmount the volume. The file system is rendered *busy* if it contains any mounts.

- Use the `include.image` option to assign a management class to the volume image. If you do not assign a management class, the default management class is used for the image.
- You can exclude a volume from image backup using the `exclude.image` option.
- You must use the mount point for the file system volume on which you want to perform an image backup. The client will not back up a file system volume without the use of a mount point. Back up file systems using the mounted name. For example, if `/dev/lv01` is formatted as a file system mounted on `/home`, enter this command to perform an image backup of this volume:

```
dsmc backup image /home
```

Back up raw volumes using the device name. For example, if `/dev/lv02` is a raw volume, enter this command to perform an image backup of this volume:

```
dsmc backup image /dev/lv02
```

If you back up a raw volume which is formatted as a file system, ensure that the file system is not mounted and does not have an entry in `/etc/filesystems`.

Related concepts:

Chapter 9, “Storage management policies,” on page 275

Related reference:

“Exclude options” on page 389

“Include options” on page 421

Volume device type support for an image backup

This topic lists several devices that are supported by the **backup image** command.

The following table lists the devices supported by the **backup image** command. A raw device might be a disk slice, a partition, or a logical volume.

Table 36. Volume device-type support for an image backup

Logical volume manager	Raw device types	Sample device name	Backup image command support
AIX Logical Volume Mgr	Logical Volumes	<code>/dev/lv00</code>	AIX
Sun Solstice DiskSuite Volume Manager	Metadevices	<code>/dev/md/dsk/dl</code>	Solaris
Solaris Volume Manager	Metadevices	<code>/dev/md/dsk/dl</code>	Solaris
Veritas Volume Mgr	Logical Volumes	<code>/dev/vx/dsk/rootdg/vol01</code> - AIX <code>/dev/vg00/lvol01</code> - Solaris	Solaris AIX
Raw Disk	Partitions	<code>/dev/hda1</code> , <code>/dev/sda3</code>	Linux x86_64, Linux on POWER
Raw Disk	Disk devices	<code>/dev/sda</code> <code>/dev/mapper/mpathX</code>	Linux x86_64
Raw Disk	Partitions	<code>/dev/dasdx</code>	Linux on System z
Linux Logical Volume Mgr	Logical Volumes	<code>/dev/myvolgroup/myvolume</code>	All Linux
Raw Disk	Disk Slices	<code>/dev/dsk/c0tld0s0</code>	Solaris
HP-UX Logical Volume Mgr	Logical Volumes	<code>/dev/dsk/c0t0d1</code> <code>/dev/vg00/lvol01</code>	HP-UX, HP-UX Itanium 2

For raw devices, the backup-archive client backs up the volume on an as-is basis. That is, no snapshot is taken, and applications can continue to write to the volume while it is being backed up. Tivoli Storage Manager cannot guarantee the consistency of the data when backing up at the physical disk level; corruption can occur if the data on the volume is changing while the backup is in progress.

The client must support the raw device type on the specific platform in order to perform an image backup of a raw device. If you want to perform an image backup for a file system mounted on a raw device, the raw device must be supported. Remember to specify raw devices by their block device name.

For the Linux clients, image backup is only supported on partitions with id 0x83 or logical volumes created with the Linux Logical Volume Manager. Backing up other partitions, such as extended partitions that contain mounted file systems or database data, might produce inconsistent backup data if the data changes during the image backup operation.

On HP-UX, image backup for raw volumes is only supported for devices that are created using HP-UX LVM logical volumes, such as `/dev/dsk/c0t0d1`. HP-UX LVM logical volume devices usually take the form `/dev/vgXY/lvolAB`.

For AIX and Solaris: You can perform image operations on volumes created using Veritas Volume Manager. The client initially supports static (default) and dynamic image type for backup.

For Solaris 10 clients, only use image backup for file systems that are assigned from the global zone to the non-global zone by exporting the device, specifying add device and set match. Do not use image backup for other file systems in the non-global zones because the non-global zone does not have the authority to mount or unmount the file system. Also, for Solaris 10 clients, do not use the overlap device of the root disk (`c0t0d0s2`) for raw device backup. Avoid using this feature on disks or slices that are used as swapping devices.

Meta devices created by the Veritas Volume Manager must be listed, including the disk group in `/etc/vfstab`, to be recognized by the backup-archive client for an image backup of file systems. The file systems should be unmounted. Raw devices should not be listed in `/etc/vfstab`. For example, the following is the correct meta device name to be used in the `/etc/vfstab` file:

```
/dev/vx/dsk/<disk group>/<meta device name>
```

Specifying `/dev/vx/dsk/` would not be recognized correctly, and you would receive an error (ANS1134E).

Disk slices containing cylinder 0 should not be backed up or restored. In this case the VTOC is overwritten. If you need to back up the first disk slice, exclude cylinder 0 by starting the disk slice from cylinder 1 (use the format utility). The backup-archive client does not check whether cylinder 0 is contained in the device that is overwritten during a restore.

Utilizing image backups to perform file system incremental backups

This topic lists the methods and steps to use image backups to perform efficient incremental backups of your file system.

These backup methods allow you to perform a point-in-time restore of your file systems and improve backup and restore performance. You can perform the backup only on formatted volumes; not on raw logical volumes.

You can use one of the following methods to perform image backups of volumes with mounted file systems.

Method 1: Using image backups with file system incremental backups

This topic lists the steps to perform image backups with file system incremental backup.

About this task

Procedure

1. Perform a full incremental backup of the file system. This establishes a baseline for future incremental backups.
2. Perform an image backup of the same file system to make image restores possible.
3. Perform incremental backups of the file system periodically to ensure that the server records additions and deletions accurately.
4. Perform an image backup periodically to ensure faster restore.
5. Restore your data by performing an incremental restore. Ensure that you select the **Image plus incremental directories and files** and **Delete inactive files from local** options in the Restore Options window before beginning the restore. During the restore, the client does the following:

Results

- Restores the most recent image on the server.
- Deletes all of the files restored in the previous step which are inactive on the server. These are files which existed at the time of the image backup, but were subsequently deleted and recorded by a later incremental backup.
- Restores new and changed files from the incremental backups.

Note: If an incremental backup is performed several times after backing up an image, make sure that the backup copy group of the Tivoli Storage Manager server has enough versions for existing and deleted files on the server so that the subsequent restore image with incremental and deletefiles options can delete files correctly.

Related tasks:

“Backing up data using the Java GUI” on page 169

“Performing an image backup using the GUI” on page 187

“Restoring an image using the GUI” on page 222

Method 2: Using image backups with incremental-by-date image backups

This topic lists the steps to perform image backups with incremental-by-date image backup.

Procedure

1. Perform an image backup of the file system.

2. Perform an incremental-by-date image backup of the file system. This sends only those files that were added or changed since the last image backup to the server.
3. Periodically, perform full image backups.
4. Restore your volume by performing an incremental restore. Ensure that you select the **Image plus incremental directories and files** option in the Restore Options window before beginning the restore. This first restores the most recent image and then restores all of the incremental backups performed since that date.

Results

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

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

This improves restore time because fewer changes are applied from incremental backups.

The following restrictions apply when using method 2:

- The file system can have no previous full incremental backups.
- Incremental-by-date image backup does not inactivate files on the server; therefore, when you restore an image with the incremental option, files deleted after the original image backup is present after the restore.
- If this is the first image backup for the file system, a full image backup is performed.
- If file systems are running at or near capacity, an out-of-space condition could result during the restore.

Related tasks:

“Performing an image backup using the GUI” on page 187

“Restoring an image using the GUI” on page 222

Comparing methods 1 and 2

This topic shows a comparison of methods 1 and 2: (1) Using image backup with file system incremental or (2) Using image backup with incremental-by-date image backup.

To help you decide which method is appropriate for your environment, the following table is a comparison of methods 1 and 2.

Table 37. Comparing incremental image backup methods

Method 1: Using image backup with file system incremental	Method 2: Using image backup with incremental-by-date image backup
Files are expired on the server when they are deleted from the file system. On restore, you have the option to delete files which are expired on server from image.	Files are not expired on server. After the image incremental restore completes, all of the files that are deleted on the file system after the image backup are present after the restore. If file systems are running at or near capacity, an out-of-space condition could result.

Table 37. Comparing incremental image backup methods (continued)

Method 1: Using image backup with file system incremental	Method 2: Using image backup with incremental-by-date image backup
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 image backup feature is configured, you can create an image backup where the real volume is available to other system applications.

About this task

A consistent image of the volume is maintained during the image backup.

When you perform an image backup using the backup-archive client GUI image backup option, the backup operation is run according to the setting of the `snapshotproviderimage` option. The `snapshotproviderimage` option defaults to an AIX JFS2 snapshot for AIX and a Linux LVM snapshot for Linux. You can override the default by using the Preferences editor Snapshot tab and the Image Snapshot Preferences.

For HP-UX and Solaris clients, selecting the image backup option performs a static image backup by default. For static image backup, the client unmounts and remounts the volume as read-only, so that no other applications can access it. You can override the default value by using the `include.image` option and selecting `dynamicimage yes`. For dynamic image backup, the client performs the image backup without making the file system read-only during the backup.

To create an image backup of your file system or raw logical volume, perform the following steps:

Procedure

1. Click on the **Backup** button in the Tivoli Storage Manager main window. The Backup window appears.
2. Expand the directory tree and select the objects you want to back up. To back up a raw logical volume, locate and expand the RAW directory tree object.
3. Click **Backup**. The Backup Task List window displays the backup processing status. The Backup Report window displays a detailed status report.

Results

- To perform a static image backup, select **Image Backup** from the drop-down list.
- For AIX and Linux clients *only*: To perform a snapshot image backup, use the `snapshotproviderimage` option.
- To perform an incremental-by-date image backup, select **Incremental image (date only)** from the drop-down list.

The following are some items to consider when you perform an snapshot-based image backup:

- To modify specific backup options, click the **Options** button. The options you select are effective only during the current session.
- To modify specific backup options, click the **Options** button. The options you select are effective only during the current session.

Linux only: The Tivoli Storage Manager Version 5.4 (and newer) client will not recognize any LVM1 volumes for image operations. However, it allows prior image backups of LVM1 volumes to be restored on LVM2 volumes. Table 38 shows the combinations involving the old and new client levels handling LVM1 and LVM2 volumes for different image operations.

Table 38. LVM1 and LVM2 image operation comparisons

Tivoli Storage Manager client version	LVM1 Backup and Restore	LVM2 Backup and Restore	Mixed Volumes	
			Backup: LVM1, Restore: LVM2	Backup: LVM2, Restore: LVM1
V5.3 and prior	YES	Only static image for file system	NO	NO - raw volumes are not supported
V5.4 and beyond	NO Error msg ANS1090E displayed	YES	YES LVM1 vol must have been backed up using prior client	NO Restore to LVM1 vol fails

Related reference:

“Snapshotproviderimage” on page 527

Performing an image backup using the command line

Use the **backup image** and **restore image** commands to perform image backup and restore operations on a single volume.

Use the mode option with the **backup image** command to perform an incremental-by-date image backup that backs up only new and changed files after the last full image backup. However, this only backs up files with a changed date, not files with changed permissions.

Related reference:

“Backup Image” on page 620

“Mode” on page 454

“Restore Image” on page 707

Snapshot-based file backup and archive and snapshot-based image backup

For backup-archive clients running on AIX 5.3 or later JFS2 file systems as root user, snapshot-based image backup is created using snapshots by default.

About this task

Optionally, you can enable snapshot-based file level backup and archive operations by specifying the snapshotproviderfs option. If for some reason a snapshot cannot be taken, the client attempts to perform a static image backup or regular file backup.

If you want to specify snapshot-based file backup and archive, set the option `snapshotproviderfs` to `JFS2`. This is applicable to all JFS2 file systems for that client.

Important: Use snapshot-based file backup and archive and snapshot-based image backup for all of your AIX JFS2 file systems.

For example, to turn *on* snapshot-based file backup and archive for all JFS2 file systems on the client, specify the following in the server stanza in the `dsm.sys` file:

```
snapshotproviderfs JFS2
```

To explicitly turn *off* snapshot-based file backup and archive for all JFS2 file systems on the client, specify the following in the server stanza in the `dsm.sys` file:

```
snapshotproviderfs NONE
```

To turn *on* snapshot-based file backup and archive for only one specific JFS2 file system on the client, specify the following in the server stanza in the `dsm.sys` file:

```
snapshotproviderfs NONE

include.fs /kalafs1 snapshotproviderfs=JFS2
```

To turn *off* snapshot-based file backup and archive for only one specific JFS2 file system on the client, specify the following in the server stanza in the `dsm.sys` file:

```
snapshotproviderfs JFS2

include.fs /kalafs2 snapshotproviderfs=NONE
```

To turn *on* snapshot-based file backup and archive for only one specific operation on the client, specify the following on the command line:

```
dsmc incr -snapshotproviderfs=JFS2 /kalafs1
```

To turn *off* snapshot-based file backup and archive for only one specific operation on the client, specify the following in the server stanza in the `dsm.sys` file:

```
snapshotproviderfs JFS2
```

Then perform the backup command. For example:

```
dsmc incr -snapshotproviderfs=NONE /kalafs2
```

The `snapshotproviderfs` option does not apply to **incremental** with the `snapdiff` and `diffsnapshot` options.

Related reference:

“`Snapshotproviderfs`” on page 526

Protecting Btrfs file systems

Btrfs file systems can be included as file specifications for backup and restore commands, archive and retrieve commands, and on **backup image** and **restore image** commands. You can also specify Btrfs subvolumes as file specification to the backup and restore, and archive and retrieve functions. You cannot use the backup-archive client image backup or image restore commands on a Btrfs subvolume.

Btrfs file systems are supported on SLES 11 SP2, or later, on IBMSystem x, System p, and System z.

If you want to create a static image backup of the entire Btrfs file system, you must unmount all the subvolumes so the backup-archive client can unmount or mount the Btrfs file system during the backup process. You can avoid the mounting and unmounting requirements if you perform a snapshot-based image backup of the Btrfs file system instead of a static image backup.

Image backup and image restore functionality is not available for Btrfs subvolumes. If you try to back up a subvolume by using the **image backup**, the following message is displayed:

```
ANS1162E Filesystem could not be mounted
```

You can mount a Btrfs subvolume by using either the subvolume name or the subvolume ID.

On Btrfs file systems, journal backup can be performed both at the file system and the subvolume level. If you perform journal-based backups on a Btrfs file system, the journal that is created is for the entire file system; there is not a separate journal for each subvolume.

Restriction: On Linux systems, some file systems such as ext2, ext3, ext4, btrfs, and xfs use a universally unique identifier (UUID) to identify themselves to the operating system. If you create an image backup of such a volume and you restore it to a different location, you might have two volumes with the same UUID. If you use UUID to define your file systems in `/etc/fstab`, be aware that the backup-archive client might be unable to correctly mount the restored file system because the UUIDs conflict. To avoid this situation, restore the image to its original location. If you must restore it to a different location, change the UUID of either the original or restored volume before you mount the restored file system. Refer to the Linux documentation for instructions on how to change a UUID. You might also need to manually edit the `/etc/fstab` file so the original volume, the restored volume, or both volumes can be mounted.

Backing up and restoring Btrfs file systems

You can back up or restore, or archive and retrieve, Btrfs file systems by using the backup-archive client **incremental**, **selective**, **restore**, **archive**, and **retrieve** commands.

About this task

If you used a version of the backup-archive client that is older than Version 7.1 to back up a Btrfs file system, the file system type was listed as Unknown, in the Tivoli Storage Manager server GUI and command output. The Unknown file system type is displayed because before Tivoli Storage Manager 7.1, Btrfs file systems were not formally supported. If you use a backup-archive V7.1 client (or newer) to back up that same Btrfs file system, all files that have Access Control Lists (ACLs) and extended attributes (XATTRs) are backed up again, even if their content has not changed since the last backup that was created by the older version of the client. Also, after a Btrfs file system is backed up by the V7.1 (or newer) client, the file system type is correctly shown as Btrfs in the Tivoli Storage Manager server GUI and command output.

Even with a V7.1 or newer client, copying a file on a Btrfs file system might cause the file to be included in the next backup operation. For example, if you copy a file by using the **cp** command with the **-p** or **-preserve** options (preserve mode, ownership, and time stamps), and if the file's attributes are changed, the access

ACL extended attribute (system.posix_acl_access) is changed. Because an extended attribute is changed, the client backs up the entire file, rather than just updating the attributes for the file.

Procedure

1. Mount the file system that you want to protect or recover. For example, use the following syntax to mount a file system: `mount /dev/sdb1 on /btreefs1 type btrfs (rw)`
2. Protect or recover the file system by performing one of the following operations:

Operation	Command
Back up the file system	<code>dsmc incr /btreefs1</code>
Restore the file system	<code>dsmc restore /btreefs1/ -subdir=yes -replace=yes</code>
Archive the file system	<code>dsmc archive /btreefs1/ -subdir=yes</code>
Retrieve the file system	<code>dsmc retrieve /btreefs1/ -subdir=yes -replace=yes</code>
Back up a file system snapshot	Create the file system snapshot. Use the btrfs subvolume snapshot command. The snapshot directory that is specified in this example is the <code>btreefs1_snap</code> directory on the file system named <code>/btreefs1</code>. <pre>btrfs subvolume snapshot /btreefs1/ /btreefs1/btreefs1_snap</pre> Issue the backup-archive client incremental command. Specify the snapshotroot option and the location of the Btrfs snapshot. <pre>\$DSM_DIR/dsmc incr /btreefs1 -snapshotroot=/btreefs1/btreefs1_snap</pre>
Perform an image backup	All subvolumes must be unmounted before you create an image backup. <pre>dsmc backup image /btreefs1 -snapshotproviderimage=none</pre> To avoid having to unmount the subvolumes, create a snapshot-based image backup. <pre>dsmc backup image /btreefs1</pre>
Restore an image backup	All subvolumes must be unmounted before you restore an image backup. <pre>dsmc restore image /btreefs1</pre>

Backing up and restoring Btrfs subvolumes

You can back up or restore, or archive and retrieve, Btrfs subvolumes by using the backup-archive client **incremental**, **selective**, **restore**, **archive**, and **retrieve** commands.

Procedure

1. List the subvolumes and determine their IDs.

```

btrfs subvolume list /btreefs1
ID 256 top level 5 path @
ID 262 top level 5 path @/btreefs1_sub1

```

2. Make the directory to use as the mount point for the subvolume.

```
mkdir /btreefs1_sub1
```
3. Mount the subvolume. For example, to mount the subvolume on device sdb1 at /btreefs1_sub1, use the following syntax: `mount -t btrfs -o subvolid=262 /dev/sdb1 /btreefs1_sub1`
 Protect or recover the subvolume by using one or more of the following operations:

Operation	Command
Back up a subvolume	Both incremental and selective backups are supported. <pre>dsmc incr /btreefs1_sub1</pre> <pre>dsmc sel /btreefs1_sub1/ -subdir=yes</pre>
Restore a subvolume	<pre>dsmc restore /btreefs1_sub1/</pre> <pre>-subdir=yes -replace=yes</pre>
Archive a subvolume	<pre>dsmc archive /btreefs1_sub1/</pre> <pre>-subdir=yes</pre>
Retrieve a subvolume	<pre>dsmc retrieve /btreefs1_sub1/</pre> <pre>-subdir=yes -replace=yes</pre>
Back up a Btrfs subvolume snapshot	Create the subvolume snapshot. Use the btrfs subvolume snapshot command. The snapshot directory that is specified in this example is the /btreefs1/btreefs1_sub1_snap directory, for the subvolume named btreefs1_sub1. <pre>btrfs subvolume snapshot</pre> <pre>/btreefs1/btreefs1_sub1</pre> <pre>/btreefs1/btreefs1_sub1_snap</pre> Issue the backup-archive client incremental command. Specify the snapshot root option and the location of the Btrfs snapshot. <pre>dsmc incr /btreefs1_sub1</pre> <pre>-snapshotroot=/btreefs1</pre> <pre>/btreefs1_sub1_snap</pre>

Back up NAS file systems using Network Data Management Protocol

Windows, AIX, and Solaris backup-archive clients can use Network Data Management Protocol (NDMP) to efficiently back up and restore network attached storage (NAS) file system images. The file system images can be backed up to, or be restored from, automated tape drives or libraries that are locally attached to Network Appliance or EMC Celerra NAS file servers, or to or from tape drives or libraries that are locally attached to a Tivoli Storage Manager server.

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

Incremental backup can also be used to back up NAS file system snapshots. See the **incremental** command and `snapshotroot`, `snappediff`, `createnewbase`, and `diffsnapshot` options for more information.

After configuring NDMP support, the server connects to the NAS device and uses NDMP to initiate, control, and monitor each backup and restore operation. The NAS device performs outboard data transfer to and from the NAS file system to a locally attached library.

Filer to server data transfer is available for NAS devices that support NDMP Version 4.

The benefits of performing backups using NDMP include the following:

- LAN-free data transfer.
- High performance and scalable backups and restores.
- Backup to local tape devices without network traffic.

The following support is provided:

- Full file system image backup of all files within a NAS file system.
- Differential file system image backup of all files that have changed since the last full image backup.
- Parallel backup and restore operations when processing multiple NAS file systems.
- Choice of interfaces to initiate, monitor, or cancel backup and restore operations:
 - Web client
 - Backup-archive client command interface
 - Administrative client command line interface (backup and restore operations can be scheduled using the administrative command scheduler)
 - Administrative web client

The following functions are *not* supported:

- Archive and retrieve
- Client scheduling. Use server commands to schedule a NAS backup.
- Detection of damaged files.
- Data-transfer operations for NAS data stored by Tivoli Storage Manager:
 - Migration
 - Reclamation
 - Export
 - Backup set generation

Related concepts:

“NDMP support requirements (Extended Edition only)” on page 9

Related reference:

“Diffsnapshot” on page 354

“Incremental” on page 654

“Snapdiff” on page 519

“Snapshotroot” on page 528

Backing up NAS file systems with the web client GUI using NDMP protocol

For both the web client GUI and the client command line interface, you must specify `passwordaccess=generate` (which is a current web client restriction for the client node) and `set authentication=on` must be specified at the server.

You are always prompted for a user ID and password. To display NAS nodes and perform NAS functions, you must enter an authorized administrative user ID and password. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web.

You can use the `toc` option with the `include.fs.nas` option in the client options file to specify whether the client saves Table of Contents (TOC) information for each file system backup. If you save TOC information, you can use the Windows web client to examine the entire file system tree and select files and directories to restore. Creation of a TOC requires that you define the `TOCDESTINATION` attribute in the backup copy group for the management class to which this backup image is bound. Note that TOC creation requires additional processing, network resources, storage pool space, and possibly a mount point during the backup operation.

The web client interface is available only for connections to Tivoli Storage Manager V7.1.7 and earlier servers, or V8.1.1 and earlier V8 servers.

To back up NAS file systems using the web client GUI:

1. Click **Backup** from the main window. The Backup window is displayed.
2. Expand the directory tree if necessary.

Note:

- a. The root node called **Nodes** is not selectable. This node only appears if a NAS plug-in is present on the client workstation.
 - b. NAS nodes display on the same level as the client workstation node. Only nodes for which the administrator has authority appear.
 - c. You can expand NAS nodes to reveal file spaces, but no further expansion is available (no file names).
3. Click the selection boxes next to the nodes or file systems you want to back up.
 4. Click the type of backup you want to perform in the backup type pull-down menu. The NAS backup type list is active only when you first select NAS backup objects. **Full backup** backs up the entire file system. **Differential** backs up the changes since the most recent full backup.
 5. Click **Backup**. The NAS Backup Task List window displays the backup processing status and progress bar. The number next to the progress bar indicates the number of bytes backed up so far. After the backup completes, the NAS Backup Report window displays processing details, including the actual size of the backup, including the total bytes backed up.

Note: If it is necessary to close the web browser session, current NAS operations continue after disconnect. You can use the **Dismiss** button on the NAS Backup Task List window to quit monitoring processing without ending the current operation.

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

Here are some items to consider when you back up NAS file systems using the web client GUI:

- Workstation and remote (NAS) backups are mutually exclusive in a Backup window. After selecting an item for backup, the next item you select must be of the same type (either NAS or non NAS).
- Details will not appear in the right-frame of the Backup window for NAS nodes or file systems. To view information about objects in a NAS node, highlight the object and select **View > File Details** from the menu.
- To delete NAS file spaces, select **Utilities > Delete Filespaces**.
- Backup options do not apply to NAS file spaces and are ignored during a NAS backup operation.

Related concepts:

“Web client configuration overview” on page 59

“Restore NAS file systems” on page 236

Related reference:

“Toc” on page 552

Back up NAS file systems using the command line

You can use the command line to back up NAS file system images.

Table 39 lists the commands and options that you can use to back up NAS file system images from the command line.

Table 39. NAS options and commands

Option or command	Definition	Page
<code>domain.nas</code>	Use the <code>domain.nas</code> option to specify the volumes to include in your default domain for NAS backups.	“Domain.nas” on page 366
<code>exclude.fs.nas</code>	Use the <code>exclude.fs.nas</code> option to exclude file systems on the NAS file server from an image backup when used with the backup nas command. This option is for AIX and Solaris clients <i>only</i> .	“Exclude options” on page 389
<code>include.fs.nas</code>	Use the <code>include.fs.nas</code> option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether Table of Contents (TOC) information is saved during a NAS file system image backup, using the <i>toc</i> option with the <code>include.fs.nas</code> option in your client options file.. This option is for AIX and Solaris clients <i>only</i> .	“Include options” on page 421
query node	Use the query node command to display all the nodes for which a particular administrative user ID has authority to perform operations. The administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using.	“Query Node” on page 681

Table 39. NAS options and commands (continued)

Option or command	Definition	Page
backup nas	Use the backup nas command to create an image backup of one or more file systems that belong to a Network Attached Storage (NAS) file server.	"Backup NAS" on page 626
toc	Use the toc option with the backup nas command or the include.fs.nas option to specify whether Table of Contents (TOC) information is saved for each file system backup.	"Toc" on page 552
monitor process	Use the monitor process command to display current backup and restore processes for all NAS nodes for which an administrative user has authority. The administrative user can then select one process to monitor.	"Monitor Process" on page 662
cancel process	Use the cancel process command to display current backup and restore processes for all NAS nodes for which an administrative user has authority. From the display, the administrative user can select one process to cancel.	"Cancel Process" on page 641
query backup	Use the query backup command with the class option to display information about file system images backed up for a NAS file server.	"Query Backup" on page 668
query filesystem	Use the query filesystem command with the class option to display a list of file spaces belonging to a NAS node.	"Query Filespace" on page 674
delete filesystem	Use the delete filesystem command with the class option to display a list of file spaces belonging to a NAS node so that you can choose one to delete.	"Delete Filespace" on page 648

- NAS nodes represent a new node type. The NAS node name uniquely identifies a NAS file server and its data to Tivoli Storage Manager. You can prefix the NAS node name to the file specification to specify the file server to which the include statement applies. If you do not specify a NAS node name, the file system you specify applies to all NAS file servers.
- Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.

Note: When you initiate a NAS backup operation by using the client command line interface, client GUI, or web client the server starts a process to initiate, control, and monitor the operation. It might take several moments before you notice progress at the client command line interface because the server must perform a mount operation, and other necessary tasks, before data movement occurs.

Related reference:

"Toc" on page 552

Backup network file systems

You can configure the backup-archive client to protect files that are accessed with either Network File System (NFS) or Common Internet File System (CIFS) protocols.

Backup performance is better when you install the backup-archive client where the file system physically exists. But sometimes it is necessary to access file systems by using NFS or CIFS to back up or recover data on remote shared drives. The backup-archive client on AIX, HP-UX, Linux, Mac OS X, and Solaris operating systems can back up, archive, restore, and retrieve file data on an NFS or CIFS-mounted shared drive. The operations are valid on all versions of the NFS protocol, including NFS version 2, NFS version 3, and NFS version 4.

The backup-archive can back up and restore access control lists when it is configured to use NFS version 4.

The following restrictions apply when the backup-archive client protects data on network file system volumes:

- Backup-archive clients cannot complete image backups of network file system volumes.
- Backup-archive clients on AIX cannot complete snapshot-based file backups or archive files on network file system volumes.
- Backup-archive clients cannot complete journal-based backups of network file system volumes.
- Backup-archive clients might not be able to back up NetApp volume snapshots if they are accessed by using the NFS protocol. If the NetApp filer provides different device identifiers for its volume snapshots, these snapshots might be excluded from backups. The behavior depends on the OS version, the NetApp filer version, and the settings.

Back up NFS file systems with the global namespace feature

NFS V4 clients can back up NFS file systems that are mounted by using the global namespace feature, which is called a *referral*. All file systems in the global namespace are backed up under a single file space.

The following examples show the file systems in the global namespace that are backed up under a single file space:

```
server 'publications' has /doc file system
server 'projects' has /projects file system
server 'data' has /data file system
```

The server account1 is the main NFS server that exports all these file systems by using a referral, and it is the server that all of the clients recognize. The /etc/exports directory on account1 looks like the following examples:

```
/doc -vers=4,refer=/doc@publications
/projects -vers=4,refer=/projects@projects
/data -vers=4,refer=/data@data
```

The client payroll mounts directories from the account1 server and can access all three file systems:

```
payroll:/#mount -o vers=4 account1:/ /mnt
payroll:/#ls /mnt
doc/ projects/ data/
```

The client payroll can back up the /mnt file as one NFS file system, which backs up all other file systems.

Important: Using the `virtualmountpoint` option can improve system performance when you back up NFSV4 file systems by using the global namespace. Add the following entries in a stanza in `dsm.sys` to back up each mounted directory as a separate file space:

```
VIRTUALMOUNTPOINT /doc
VIRTUALMOUNTPOINT /projects
VIRTUALMOUNTPOINT /data
```

Back up AIX workload partition file systems

Using the backup-archive client on AIX, you can back up and restore local partition file data within the global partition by using the local partition name space available within the global partition.

Each workload partition (WPAR) has its own security domain, so only the global root user is guaranteed to have access to all of the data.

The WPARs are partitions that are created entirely in software within a single AIX system image, with the following attributes:

- Usually the WPAR appears to be a complete stand-alone AIX system
- There is no hardware assist or configuration

Workload partitions provide a secure and isolated environment for enterprise applications in terms of process, signal, and file system space. Software running within the context of a workload partition appears to have its own separate instance of AIX.

The following example shows a WPAR configuration from within the global WPAR:

Global partition:

System name: shimla
File system: /home /opt

WPAR #1 configuration:

Name: wpar1
File system: /home; name in global WPAR: /wpars/wpar1/home

WPAR #2 configuration:

Name: wpar2
File system: /data; name in global WPAR: /wpars/wpar2/data

There are two ways to back up WPAR data, as follows:

- Back up all WPAR file systems as the file spaces within the global partition. The file space name must be used to identify the WPAR to which it belongs. All of the data is managed on one node by using one schedule. Using the example configuration, here is a sample `dsm.sys` file with one server stanza for all file systems, both global and local:

```
Servname shimla
TCPport      1500
TCPserveraddress server.example.com
```

```

nodename      shimla
PasswordAccess generate
Domain        /wpars/wpar1/home /wpars/wpar2/data /home /opt

```

- Back up each WPAR file system under a different node name. This method provides file space name segregation for each WPAR. Each WPAR must have a separate node name and a scheduler that is running within the global partition. Also, three scheduler services must be set up, each using a different `dsm.opt` file corresponding to the server stanza name. This method allows each WPAR backup operation to be managed independently of the others. Using the example configuration, here is a sample `dsm.sys` file with three server stanzas: one for `wpar1`, one for `wpar2`, and one for global partition `shimla`:

```

SErvername shimla_wpar1
TCPPort      1500
TCPSeveraddress server.example.com
nodename      wpar1
PasswordAccess generate
Domain        /wpars/wpar1/home

```

```

SErvername shimla_wpar2
TCPPort      1500
TCPSeveraddress server.example.com
nodename      wpar2
PasswordAccess generate
Domain        /wpars/wpar2/data

```

```

SErvername shimla
TCPPort      1500
TCPSeveraddress server.example.com
nodename      shimla
PasswordAccess generate
Domain        /home /opt

```

Backing up Solaris Zettabyte file systems

On Solaris SPARC and Solaris x86 systems, you can backup Zettabyte file systems (ZFS), by using ZFS snapshots. The advantage of this approach, over an ordinary incremental or selective backup, is that the files and folders in a snapshot are always in a read-only state, so they cannot be changed during a backup.

About this task

You create a ZFS snapshot by using Oracle Solaris ZFS commands. For example:

```
zfs snapshot tank/myZFS@mySnapshot
```

In this example, the ZFS pool name is called `tank` and the ZFS file system name is `myZFS`. Files that belong to this ZFS snapshot are in the subdirectory named `tank/myZFS/.zfs/snapshot/mySnapshot/`.

Procedure

Use either of these two methods to backup a ZFS snapshot.

- Backup each file of the snapshot by using the `snapshotroot` option. For example:

```
dsmc inc -snapshotroot=/tank/myZFS/.zfs/snapshot/mySnapshot /tank/myZFS
```

This option allows the administrator to replace the current snapshot path with the ZFS file system path, so that the files and folders are backed up under the original file system.

- Backup the complete snapshot by using Oracle Solaris ZFS commands. For example:

```
zfs send tank/myZFS@mySnapshot > /tmpdir/mySnapshotFile
```

The advantage of backing up the complete snapshot is that the full file system can be restored, in a disaster recovery scenario.

Related concepts:

“Restoring Solaris Zettabyte (ZFS) file systems” on page 243

Related reference:

“Snapshotroot” on page 528

AIX JFS2 encrypted file system backup

Use AIX JFS2 Encrypted File System (EFS) to back up files either in clear text or raw format. With clear text format, the file is decrypted by EFS as it is read. With raw format, the data is not decrypted. The default is raw format, but when you set the `efsdecrypt` option to `yes`, you get clear text backups.

About this task

Important: Whenever you run a backup that includes any files encrypted on an EFS, you must ensure that you use the correct specification of the `efsdecrypt` option. If the `efsdecrypt` option value changes between two incremental backups, all encrypted files on EFS file systems are backed up again, even if they have not changed since the last backup. For example, if you are running an incremental backup of encrypted files that were previously backed up as raw, then ensure that `efsdecrypt` is specified as `no`. If you change `efsdecrypt` to `yes`, all of the files are backed up again in clear text even if they are unchanged, so ensure that you use this option carefully.

If you attempt to restore an encrypted file to either a work station that does not support EFS, or a file system where EFS is not active, an error message is written and the file is skipped.

Here are some reasons to back up EFS using clear text encryption:

- This type of decryption is useful if you want to use the Tivoli Storage Manager backup-archive client encryption or another type of hardware encryption (for tape systems, for example).
- You can use clear text for long term archival of data, because the data is stored independent of the platform or encryption scheme.

Here are some things to consider when backing up a file in clear text:

- The user who invoked the backup-archive client must be able to decrypt it
- The user can have read access to a file, but not have access to the key

In the following scenarios an error message is issued:

Procedure

1. The user is running in root guard mode, and EFS has the concept of two types of root. Root admin is the traditional mode. A root in guard mode will not have access to the unencrypted data, unless the user is the owner or a member of the file group.

2. The user is running with a non-root user ID and attempting an archive of a file to which they have read access, but the user is not the owner or member of the file group. EFS will not allow the data to be decrypted.

Results

Here are some considerations when backing up EFS raw data:

- The backup-archive client will not honor the client encryption setting, which prevents double encryption, but only at the client. The server has no knowledge that the data is encrypted so any encryption done by a tape drive, for example, still occurs.
- The client will not honor the compression setting, so the client will not even try to compress the data.
- The client does not automatically back up or restore the keystore files. When you are restoring encrypted files, you might also have to restore keystores in order to decrypt the data.

Tips:

1. To protect the keystore, make sure the contents of `/var/efs` are included in your periodic backups.
 2. For the keystore data, use Tivoli Storage Manager storage policy with an unlimited number of versions.
- Encrypted file system (EFS) files backed up in raw mode (default) cannot be restored by a backup-archive client prior to V5.5, or by a client on another UNIX platform.

Back up AIX JFS2 extended attributes

AIX Enhanced Journal File System (JFS2) provides backup processing for named extended attributes for all file systems that support named extended attributes.

These extended attributes are automatically backed up with each object that contains extended attributes data, and no additional action is required.

When the file system is defined with the v2 format, the only file system that supports named extended attributes is JFS2. You can use JFS2 for extended attributes for files and directories, but you cannot use JFS2 for extended attributes on symbolic links.

Backing up VMware virtual machines

You can use the backup-archive client to back up and restore a VMware virtual machine (VM). Full backups of the virtual machine operate at a disk image level. Incremental backups copy only the data that is changed since the previous full backup.

Table 40 lists the backup and restore capabilities for VMware virtual machines that the backup-archive client can implement on Linux platforms.

Table 40. Backup and restore capabilities for VMware virtual machines on Linux platforms

Capability	Comment
Full VM backup:	Creates an image backup of the virtual machine.

Table 40. Backup and restore capabilities for VMware virtual machines on Linux platforms (continued)

Capability	Comment
Incremental VM backup:	Requires the IBM Tivoli Storage Manager for Virtual Environments licensed product. Backs up all changes that occurred since the previous backup of the virtual machine, whether the backup was a full backup, or another incremental backup. For this backup mode, you must schedule full backups of the virtual machine to occur periodically so that the historical data can be expired. Historical data is expired from a full backup and all incremental backups that depend on the full backup. Restoring data from incremental backups is not convenient or efficient because the process must automatically complete the following tasks: • Restore the most recent full backup. • Restore each incremental backup as required to restore your content to a specific point in time. For more efficient backups, use the incremental-forever backup modes.
Full VM incremental-forever backup:	Requires the IBM Tivoli Storage Manager for Virtual Environments licensed product. A full VM backup is required before you can create incremental backups. If you schedule incremental-forever backups, this backup type is selected automatically for the first backup if a full backup was not already created. Data from incremental backups is combined with data from the full backup to create a synthetic full backup image. Subsequent full VM incremental-forever backups read all used blocks and copy those blocks to the Tivoli Storage Manager server. Each full VM incremental-forever backup reads and copies all of the used blocks, whether the blocks are changed or not since the previous backup. You can still schedule a full VM backup, although a full backup is no longer necessary. For example, you might run a full VM backup to create a backup to a different node name with different retention settings. You cannot use this backup mode to back up a VMware virtual machine if the client is configured to encrypt the backup data.
Incremental-forever-incremental VM backup:	Requires the IBM Tivoli Storage Manager for Virtual Environments licensed product. Requires you to create a full VM backup one time only. The full VM backup copies all of the used disk blocks owned by a virtual machine to the Tivoli Storage Manager server. After the initial full backup is complete, all subsequent backups of the virtual machine are incremental-forever-incremental backups. Each incremental-forever-incremental backup copies only the blocks that are changed since the previous backup, irrespective of the type of the previous backup. The server uses a grouping technology that associates the changed blocks from the most recent backup with data already stored on the server from previous backups. A new full backup is then effectively created each time changed blocks are copied to the server by an incremental-forever-incremental backup. The incremental-forever-incremental backup mode provides the following benefits: • Improves the efficiency of backing up virtual machines. • Simplifies data restore operations. • Optimizes data restore operations. During a restore operation, you can specify options for point-in-time and point-in-date to recover data. The data is restored from the original full backup and all of the changed blocks that are associated with the data. You cannot use this backup mode to back up a VMware virtual machine if the client is configured to encrypt the backup data.

Table 40. Backup and restore capabilities for VMware virtual machines on Linux platforms (continued)

Capability	Comment
Item recovery for files and folders from a full backup of the virtual machine:	Requires the IBM Tivoli Storage Manager for Virtual Environments licensed product. Provides the capability to recover files and folders from a full backup of a virtual machine. Item recovery is available only with the Tivoli Storage Manager recovery agent.
Full restore of the virtual machine:	Restores all of the file systems, virtual disks, and the virtual machine configuration.

Related concepts:

“Parallel backups of virtual machines” on page 205

Related tasks:

“Preparing the environment for full backups of VMware virtual machines”

“Creating full backups for VMware virtual machines” on page 204

Preparing the environment for full backups of VMware virtual machines

Complete the following steps to prepare the VMware environment for backing up full VMware virtual machines. The vStorage backup server can run either a Windows or Linux client.

Procedure

1. To configure the storage environment for backing up, complete the following steps:
 - a. Configure your storage environment so that the vStorage backup server can access the storage volumes that are in your ESX server farm.
 - b. If you are using network-attached storage (NAS) or direct-attach storage, ensure that the vStorage backup server is accessing the volumes with a network-based transport.
 - c. Optional: For data access, make the following settings:
 - Create storage area network (SAN) zones that your vStorage backup server can use to access the storage logical units (LUNs) that host your VMware datastores.
 - Configure your disk subsystem host mappings so that all ESX servers and the backup proxy can access the same disk volumes.
2. To configure the vStorage backup server, complete the following steps:
 - a. Set and export the **LD_LIBRARY_PATH** environment variable to point to the client installation directory. For example:


```
export LD_LIBRARY_PATH=/opt/tivoli/tsm/client/ba/bin
```
 - b. Add the client installation directory to the path of each account that uses backup-archive client commands, for example, **dsmc**, **dsmcad**, or **dsmj**.
3. To modify Tivoli Storage Manager, complete the following steps:
 - a. Access the administrative command line on the backup-archive client.
 - b. From the backup-archive client on the vStorage backup server, run the following command to register the node:


```
register node my_server_name my_password
```

Where *my_server_name* is the full computer name of the vStorage backup server and *my_password* is the password to access the server.

Related tasks:

“Creating full backups for VMware virtual machines”

Related reference:

“Backup VM” on page 634

“Query VM” on page 691

“Restore VM” on page 715

“Vmchost” on page 570

“Vmcpw” on page 571

“Vmcuser” on page 573

“Vmvstortransport” on page 597

Creating full backups for VMware virtual machines

A full backup of a VMware virtual machine is a backup of an entire virtual machine, including the virtual disks and the virtual machine configuration file. This type of backup is similar to an image backup. To create the full backup, you configure the backup-archive client on the vStorage backup server. The vStorage backup server must run a Windows client or a Linux client.

Procedure

1. To prepare the environment, complete the steps in the following topic:
“Preparing the environment for full backups of VMware virtual machines” on page 203
2. To configure the backup-archive client on the vStorage backup server, complete the following steps:
 - a. From the welcome page of the backup-archive client GUI, click **Edit > Client Preferences**.
 - b. Select the **VM Backup** tab.
 - c. Select **VMWare Full VM**.
 - d. In the **Domain Backup Types** list, select **Domain Full VM**.
 - e. In the **Host** field, enter either the host name of each ESX server or the host name of the Virtual Center. If you specify the Virtual Center, you can back up virtual machines from any of the VMware servers that are managed by the Virtual Center.
 - f. Enter the user ID and password information for the host that you specify in the **Host** field.
 - g. Optional: If you want to override the default management class for full virtual machine backups, specify the management class that you want to use.
 - h. In the **Datastore Location** field, enter the path to the directory where the files are stored.
 - i. Click **OK** to save your changes.
3. To create a backup of one of the virtual machines, complete the following steps:
 - a. At the command line of the vStorage backup server, run the following command:

```
dsmc backup vm my_vm_name -mode=full -vmbackuptype=fullvm
```

Where *my_vm_name* is the name of the virtual machine.

- b. Verify that the command is completed without errors. The following message indicates successful completion:

```
Backup VM command complete
Total number of virtual machines backed up successfully: 1
virtual machine vmname backed up to nodename NODE
Total number of virtual machines failed: 0
Total number of virtual machines processed: 1
```

4. To verify that you can restore the files for the virtual machine, complete the following steps:
 - a. At the command-line interface of the vStorage backup server, run the following command:
`dsmc restore vm my_vm_name`
 - b. If errors occur in the restore processing, view the client error log for more information.

Tip: The log file is saved to `/opt/ibm/Tivoli/TSM/baclient/dsmerror.log`

Related concepts:

“Parallel backups of virtual machines”

Related tasks:

“Preparing the environment for full backups of VMware virtual machines” on page 203

Related reference:

“Backup VM” on page 634

“Domain.vmfull” on page 369

“Query VM” on page 691

“Restore VM” on page 715

“Mode” on page 454

“Vmchost” on page 570

“Vmcpw” on page 571

“Vmcuser” on page 573

“Vmmc” on page 584

“Vmvstortransport” on page 597

Parallel backups of virtual machines

A single Tivoli Storage Manager data mover node can be used to concurrently back up multiple virtual machines.

When the backups are initiated, the client establishes parallel sessions to copy the data to the Tivoli Storage Manager server.

For backing up VMware virtual machines, client options are provided so you can optimize the backups so they do not adversely affect the ESX servers that are hosting the virtual machines. The options are described in detail in “Client options reference” on page 316. A short summary is provided here to illustrate their use.

- The `vmmaxparallel` option limits how many virtual machines can be backed up at any one time.
- The `vmlimitperhost` option limits how many virtual machines can be backed up, at any one time, per ESX server.

- The `vm-limit-per-datastore` option limits how many virtual machines can be backed up, at any one time, per datastore.

`vm-max-parallel`

The `vm-max-parallel` option specifies how many virtual machines can be backed up at any one time. The optimal value for `vm-max-parallel` is not obvious; it depends on the processing power of the vStorage server that the data mover node runs on, and the performance of I/O between the client and the Tivoli Storage Manager server. For example, if you are moving data to the server over a busy LAN, you might need to limit the number of virtual machines in each parallel backup session. Similarly, if the vStorage server processing capabilities are limited, for any reason, this is also a reason to restrict the value for `vm-max-parallel`. The default for this option is 1. Using the default value mimics the same behavior of older clients, where only one virtual machine is backed up during a session.

`vm-limit-per-host`

The `vm-limit-per-host` option specifies how many virtual machines can be backed up from a single ESX/ESXi host, at the same time. Like the value for `vm-max-parallel`, you might have to experiment with this setting to find the optimum value. On ESX/ESXi servers that are heavily used, you might need to restrict the value for `vm-limit-per-host` so you do not adversely affect the vSphere server performance. On servers that are not as heavily used, you can include more virtual machines.

`vm-limit-per-datastore`

The `vm-limit-per-datastore` option limits the maximum number of virtual machines that can be backed up, at any one time, from a datastore. In a multiple datastore VMware environment, you can use this option to reduce the burden that is placed on any one datastore during a parallel backup operation.

These options work in concert with each other. Examples are provided to illustrate how these options determine how many virtual machines are backed up from any host or datastore.

Parallel backup examples

In the following figures, the circled virtual machines are the virtual machines that are selected for backup processing, which is based on the option settings in `domain.vmfull`.

Example 1: Each VM is stored in a single datastore

Figure 1 on page 207 shows that each of the circled virtual machines has its data saved in a unique datastore. Assume that the parallel backup options are set to the following values:

- `vm-max-parallel 3`
- `vm-limit-per-host 1`
- `vm-limit-per-datastore 1`




Figure 1. Virtual machines using unique datastores.

In Host A, only virtual machines 1 and 5 match the selection criteria on a `domain.vmlimitperhost` statement. In Host B, only virtual machine 2 matches the selection criteria. In this configuration, each virtual machine has a separate datastore, so the `vmlimitperdatastore` setting is valid. But, since `vmlimitperhost` is set to one, only one virtual machine (vm1 or vm5) from Host A and one virtual machine (vm2) from Host B are included when the Backup VM operation is run; only two virtual machines are included in this backup session.

Example 2: Same as example 1, but with a different setting for `vmlimitperhost`

Figure 2 on page 208 shows that each of the circled virtual machines has its data saved in a unique datastore. In this configuration, the `vmlimitperhost` is increased to two to illustrate how the option increase changes the Backup VM operation. Assume that the parallel backup options are now set to the following values:

- `vmmxparallel 3`

- `vmlimitperhost` 2 (an increase of 1)
- `vmlimitperdatastore` 1




Figure 2. Virtual machines using unique datastores, with different option value for `vmlimitperhost`.

The same virtual machines match the `domain.vmfull` criteria as they did in the previous example. However, with the increase in the `vmlimitperhost` setting, now a total of three virtual machines are included in a Backup VM operation (vm1 and vm5 from Host A, and vm2 from Host B).

Example 3: Some VMs share datastores

Figure 3 on page 209 shows that the VMDK and configuration files for virtual machine 5 in Host A is stored in two datastores. To include both vm1 and vm5 in Host A in the parallel backup operation, the value of `vmlimitperdatastore` must be increased to at least two. If `vmlimitperdatastore` is not increased to two, or higher, the backup of the second virtual machine (vm5), in Host A, cannot be started until

the first virtual machine (vm1) backup is completed because the two VMs share data in datastore1.




Figure 3. Virtual machines sharing a datastore.

Back up and archive Tivoli Storage Manager FastBack data

Use Tivoli Storage Manager FastBack to back up and archive the latest snapshots for short-term retention.

Use the **archive fastback** and **backup fastback** commands to archive and back up volumes that are specified by the `fbpolicyname`, `fbclientname` and `fbvolumename` options for short-term retention.

Related concepts:

“Installation requirements for backing up and archiving Tivoli Storage Manager FastBack client data” on page 9

“Configuring the client to back up and archive Tivoli Storage Manager FastBack data” on page 94

Related reference:

“Fbclientname” on page 396

“Fbpolicyname” on page 398

“Fbvolumename” on page 401

Display backup processing status

During a backup, by default the backup-archive client displays the status of each file it attempts to back up.

The client reports the size, path, file name, total number of bytes transferred, and whether the backup attempt was successful for the file. These are also recorded in the `dsmsched.log` file for scheduled commands.

The web client and backup-archive client Java GUI provide a **Task List** window that displays information about files during processing. When a task completes, a **Backup Report** window displays processing details. Click the **Help** button in the **Backup Report** window for context help.

On the backup-archive command line, the name of each file is displayed after it is sent to the server. The progress indicator shows overall progress.

Table 41 lists some informational messages and meanings.

Table 41. Client command line informational messages

Informational message	Meaning
Directory-->	Indicates the directory that you back up.
Normal File-->	Any file that is not a directory, symbolic link, or special file.
Special File-->	Special files define devices for the system or temporary files that are created by processes. There are three basic types of special files: FIFO (first-in, first-out), block, and character. FIFO files are also called pipes. Pipes are created by one process to temporarily allow communication with another process. These files cease to exist when the first process finishes. Block and character files define devices. The client processes only device and named pipe special files. Socket special files are not processed.
Symbolic Link-->	Indicates that the client backs up a symbolic link.
Updating-->	Indicates that only the file meta data is sent, not the data itself.
Expiring-->	Indicates an object (file or directory) on the server that no longer exists on the client is expired and made inactive on the server.
Total number of objects inspected:	As indicated. When using journal-based backup, the number of objects that are inspected might be less than the number of objects that are backed up. When you use the snapshot difference incremental backup, the number of objects that are inspected is zero. The number is zero because the client performs an incremental backup of the files that NetApp reported as changed. The client does not scan the volume looking for files that have changed.
Total number of objects backed up:	As indicated.

Table 41. Client command line informational messages (continued)

Informational message	Meaning
Total number of objects encrypted:	This is a count of the objects that were encrypted during backup or archive processing.
Data encryption type:	Specifies the encryption algorithm type (e.g 256-bit AES), if one or more objects are encrypted during backup or archive processing.
Total number of objects updated:	These are files whose attributes, such as file owner or file permissions, have changed.
Total number of objects rebound:	See “Bind management classes to files” on page 283 for more information.
Total number of objects deleted:	This is a count of the objects that are deleted from the client workstation after being successfully archived on the server. The count is zero for all backup commands.
Total number of objects expired:	See the section about full and partial incremental backup for more information.
Total number of objects failed:	Objects can fail for several reasons. Check the dsmerror.log for details.
Total snapshot difference objects:	For snapshot difference incremental backups, this represents the total number of objects backed up and the total number of objects expired.
Total objects deduplicated:	Specifies the number of files that are deduplicated.
Total number of bytes inspected:	Specifies the sum of the sizes of the files that are selected for the operation. For example, the total number of bytes that are inspected for this command is the number of bytes that are used on the volume /Volumes/BUILD: dsmc INCREMENTAL /Volumes/BUILD/* -SU=Yes
Total bytes before deduplication:	Specifies the number of bytes to send to the Tivoli Storage Manager server if the client does not eliminate redundant data. Compare this amount with Total bytes after deduplication. Includes metadata size and might be greater than bytes inspected.
Total bytes after deduplication:	Specifies the number of bytes that are sent to the Tivoli Storage Manager server after deduplication of the files on the client computer. Includes metadata size and might be greater than bytes processed.
Total number of bytes processed:	Specifies the sum of the sizes of the files that are processed for the operation.
Data transfer time:	The total time to transfer data across the network. Transfer statistics might not match the file statistics if the operation was retried due to a communications failure or session loss. The transfer statistics display the bytes attempted to be transferred across all command attempts.
Network data transfer rate:	The average rate at which the network transfers data between the client and the server. This is calculated by dividing the total number of bytes transferred by the time to transfer the data over the network. The time it takes to process objects is not included in the network transfer rate. Therefore, the network transfer rate is higher than the aggregate transfer rate.

Table 41. Client command line informational messages (continued)

Informational message	Meaning
Aggregate data transfer rate:	The average rate at which Tivoli Storage Manager and the network transfer data between the client and the server. This is calculated by dividing the total number of bytes transferred by the time that elapses from the beginning to the end of the process. Both Tivoli Storage Manager processing and network time are included in the aggregate transfer rate. Therefore, the aggregate transfer rate is lower than the network transfer rate. Note: On occasion, the aggregate data transfer rate might be higher than the network data transfer rate. This is because the backup-archive client can have multiple simultaneous sessions with the backup server. If you set the <code>resourceutilization</code> option, the client attempts to improve performance and load balancing by using multiple sessions when it backs up a file space or other set of files. When multiple sessions are open during backup, the data transfer time represents the sum of the times reported by all sessions. In this case, aggregate data transfer time is incorrectly reported as higher. However, when running with a single session, the aggregate data transfer rate should always be reported as lower than the network data transfer rate.
Objects compressed by:	Specifies the percentage of data sent over the network divided by the original size of the file on disk. For example, if the net data-bytes are 10K and the file is 100K, then Objects compressed by: $((1 - (10240/102400)) \times 100 = 90\%$.
Total number of objects grew:	The total number of files that grew larger as a result of compression.
Deduplication reduction:	Specifies the size of the duplicate extents that were found, divided by the initial file or data size. For example, if the initial object size is 100 MB, after deduplication it is 25 MB. The reduction would be: $(1 - 25/100) \times 100 = 75\%$.
Total data reduction ratio:	Adds incremental and compression effects. For example, if the bytes inspected are 100 MB and the bytes sent are 10 MB, the reduction would be: $(1 - 10/100) \times 100 = 90\%$
Elapsed processing time:	The active processing time to complete a command. This is calculated by subtracting the starting time of a command process from the ending time of the completed command process.
Total number of bytes transferred:	As indicated.
LanFree bytes transferred:	The total number of data bytes transferred during a lan-free operation. If the <code>enablelanfree</code> option is set to <i>no</i> , this line will not appear.
Total number of bytes inspected:	A sum of sizes of files selected for the operation.
Total number of retries:	The total number of retries during a backup operation. Depending on the settings for the serialization attribute and the changingretries option, a file that is opened by another process might not be backed up on the first backup try. The backup-archive client might try to back up a file several times during a backup operation. This message indicates the total retries for all files that are included in the backup operation.

Backup (UNIX and Linux): Additional considerations

There are some special situations that you need to consider before you back up your data.

Stored files

When you back up and archive files, Tivoli Storage Manager stores the backups and archives in a file space in storage that has the same name as the file system or virtual mount point from which the files originated.

For example, if you have a file system named `/home`, and you back up a file named `doc1` in the `/home/monnett` directory, Tivoli Storage Manager stores the file in a file space named `/home`. If you later define `/home/monnett` as a virtual mount point, any files you back up from the `/home/monnett` directory, such as `doc2`, are stored in a file space named `/home/monnett`. If you enter this command:

```
dsmc query backup "/home/monnett/*"
```

Tivoli Storage Manager looks for files in the `/home/monnett` file space. It always looks for a file in the file space with the longest name that matches the file specification you include in a command. It locates the file named `doc2` that was backed up after the virtual mount point was defined. However, it does not locate the file named `doc1` because that file was backed up before the virtual mount point was defined and the backup was stored in the `/home` file space.

To list or restore the `doc1` file using a command, you must explicitly specify the file space name by enclosing it in braces. For example:

```
dsmc query backup "{/home}/monnett/*"  
dsmc restore {/home}/monnett/doc1
```

If you subsequently remove the `/home/monnett` virtual mount point, and you then back up additional files in the `/home/monnett` directory, the backups are once again stored in the `/home` file space. For example, if you now back up a file named `doc3` in the `/home/monnett` directory, it is stored in the `/home` file space. It is not stored in the existing `/home/monnett` file space.

However, because the `/home/monnett` file space already exists, when you try to query or restore the `doc3` file, Tivoli Storage Manager looks for the file in the `/home/monnett` file space unless you specify the correct file space name. For example:

```
dsmc query backup "{/home}/monnett/*"  
dsmc restore {/home}/monnett/doc2
```

Note: You must explicitly specify the file space name only when there can be more than one resolution to the file specification.

For example, if the following file spaces exist in storage:

```
/home  
/home/monnett  
/home/monnett/project1  
/home/monnett/project1/planning
```

then enter:

```
dsmc query backup "/home/monnett/project1/planning/*"
```

Tivoli Storage Manager looks for files only in the `/home/monnett/project1/planning` file space, even if one or more of the other file spaces contains a path with the same name. But, when you enter one of the following:

```
dsmc query backup "{/home}/monnett/project1/planning/*"  
dsmc query backup "{/home/monnett}/project1/planning/*"  
dsmc query backup "{/home/monnett/project1}/planning/*"
```

Tivoli Storage Manager looks for files only in the /home file space, the /home/monnett file space, or the /home/monnett/project1 file space, depending on which form you use.

Special file systems

Special file systems contain dynamic information generated by the operating system; they contain no data or files. The backup-archive client ignores special file systems and their contents.

Special file systems include the following:

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

NFS or virtual mount points

When files are backed up and archived from a file system or virtual mount point, the client does not follow the nested NFS or virtual mount points (if any are defined on a file system). The nested NFS or virtual mount points will not be backed up or archived.

Management classes

Tivoli Storage Manager uses management classes to determine how to manage your backups on the server.

Every time you back up a file, the file is assigned a management class. The management class used is either a default selected for you, or one assigned to the file with an include option in the include-exclude options list. The selected management class must contain a backup copy group in order for the file to be backed up.

Select **Utilities** → **View Policy Information** from the Java or web client GUI to view the backup policies defined by the Tivoli Storage Manager server for your client node.

Related concepts:

Chapter 9, “Storage management policies,” on page 275

Back up symbolic links

The backup-archive client backs up symbolic links differently than it does regular files and directories.

The way that the client backs up symbolic links depends on options settings, whether the target directory is accessible, and the way you specify objects.

A *UNIX symbolic link* is a file that contains a pointer to another file or directory. The object the symbolic link points to is called the target object.

A symbolic link can be backed up as path information to a target directory, or it can be backed up as a directory. If the symbolic link is backed up as a directory, the files and folders in the target directory can also be backed up.

Note: Symbolic link processing as described here does not apply to Mac OS X. Symbolic links are always backed up as files and are never followed.

Related reference:

“Archsymlinkasfile” on page 319

“Followsymbolic” on page 407

“Virtualmountpoint” on page 566

Examples: Incremental or selective backup of symbolic links

How the client backs up a symbolic link depends on whether the target of the symbolic link is a file or a directory, and how you specify the symbolic link on the incremental or selective backup command.

If a symbolic link points to a file, the client only backs up the path information. The client does not back up a file that is the target of a symbolic link.

If a symbolic link points to a directory, the backup depends on how the directory is specified on the command.

If a directory is specified with a trailing slash on a selective or incremental backup command, the client saves the symbolic link as a directory, and backs up the contents of the target directory.

If the symbolic link is entered without a trailing slash, or if a symbolic link is not explicitly stated in a backup file specification, the client backs up only the path information to the target directory. The contents of the target directory are not backed up.

In the following examples, assume that `symdir` is a symbolic link to target directory `/fs1/guest/`. `/fs1/guest/` contains these objects:

`/fs1/guest/file` (a file)

`/fs1/guest/dir1` (a directory)

`/fs1/guest/dir1/file1` (a file)

Example 1

```
dsmc incr /home/gillis/symdir/
```

In this example, the client backs up the symbolic link as a directory, and backs up the contents of the target directory `/fs1/guest/`. If you specify the `subdir=yes` option, the client backs up subdirectories of `/fs1/guest/`.

Example 2

```
dsmc incr /home/gillis/symdir/dir1
```

Example 3

```
dsmc incr /home/gillis/symdir/dir1/
```

In examples 2 and 3, the client backs up the symbolic link as a directory, and backs up the `/dir1/` subdirectory of the target directory. the trailing slash is relevant only for the symbolic link; it is not relevant for subdirectories of the symbolic link. If you specify the `subdir=yes` option, the client backs up subdirectories of `/fs1/guest/dir1`. Backup copies that are stored on the Tivoli Storage Manager server have a path like `/home/gillis/symdir/dir1/file1`.

Example 4

```
dsmc incr /home/gillis/symdir
```

In example 4, because there is no trailing slash after the symbolic link, the client backs up only the path to the target directory. The client does not back up the symbolic link as a directory, and does not back up files nor folders in the target directory.

Example 5

```
dsmc incr /home/gillis/
```

In example 5, because the symbolic link is not explicitly stated in the backup file specification, the client backs up only the path to the target directory. The client does not back up the symbolic link as a directory, and does not back up files nor folders in the target directory.

Restriction: If you back up a symbolic link as a directory, a future incremental backup that does not back up that symbolic link as a directory expires that symbolic link as a directory, and expires the files and directories in that directory.

For example, assume that you first back up the symbolic link `symdir` as a directory, and back up the contents of the target directory. The command in example 1 does this. The client creates backup copies with a high-level path `/home/gillis/symdir/`. In this example, the client creates backup copies with these paths:

```
/home/gillis/symdir/  
/home/gillis/symdir/file  
/home/gillis/symdir/dir1  
/home/gillis/symdir/dir1/file1
```

The contents of `/home/gillis` are backed up using the following command:

```
dsmc inc /home/gillis/ -subdir=yes
```

This command processes the value `symdir` as a symbolic link and does not process any objects that the symbolic link points to. Hence, the client expires backup copies in the `/home/gillis/symdir/` directory that were created in example 1.

Incremental backup of a domain only

The client backs up a symbolic link during an incremental backup of the domain, if the symbolic link is defined as a virtual mount point and the `followsymbolic` option is set to `yes`.

The client backs up a symbolic link and the target directory when all of the following conditions are true:

- The client performs an incremental backup of the domain.
- The symbolic link is defined as a virtual mount point using the `virtualmountpoint` option.
- `followsymbolic=yes`

The `virtualmountpoint` and `followsymbolic` options add the symbolic link to the domain. The **Incremental** command backs up the domain, which includes the symbolic link target.

Related reference:

“Followsymbolic” on page 407

“Virtualmountpoint” on page 566

Hard links

When you back up files that are hard-linked, the backup-archive client backs up each instance of the linked file.

For example, if you back up two files that are hard-linked, the client backs up the file data twice.

When you restore hard-linked files, the client attempts to reestablish the links. For example, if you had a hard-linked pair of files, and only one of the hard-linked files is on your workstation, when you restore both files, they are hard-linked. The files are also hard-linked even if neither of the files exists at the time of restore, if both of the files are restored together in a single command. The one exception to this procedure occurs if you back up two files that are hard-linked and then break the connection between them on your workstation. If you restore the two files from the server using the standard (or classic) restore process, the client respects the current file system and does not re-establish the hard link.

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

Sparse files

Sparse files do not have disk space allocated for every block in the whole address space, leading to holes within the file. Holes are detected by their content, which is always zeros, and these zeros take up space.

The default is to restore the sparse file without the holes, which would leave more free disk space. The Tivoli Storage Manager client detects sparse files during a backup operation and marks them as sparse on the Tivoli Storage Manager server.

Note: Sparse files do not apply to Mac OS X.

The backup-archive client backs up a sparse file as a regular file if client compression is off.

Related reference:

“Compression” on page 339

“Makesparsefile” on page 447

NFS hard and soft mounts

When the backup-archive client connects to an NFS file system, you can use either a hard mount or a soft mount.

The client uses the **nfstimeout** option value to determine how long to wait for an NFS system call to respond before timing out; this setting applies to hard and soft mounts. The default is 0 seconds. This means that the client uses the default behavior of NFS system calls.

Be aware of the consequences of hard and soft mounts if the mount becomes stale (for example, if the server for the file system is not available).

Hard mount

If the NFS file system is hard mounted, the NFS daemons try repeatedly to contact the server. The NFS daemon retries will not time out, they affect system performance, and you cannot interrupt them, but control returns to the client when the **nfstimeout** value is reached.

Soft mount

If the NFS file system is soft mounted, NFS tries repeatedly to contact the server until either:

- A connection is established
- The NFS retry threshold is met
- The **nfstimeout** value is reached

When one of these events occurs, control returns to the calling program.

Note: On UNIX and Linux systems, the **nfstimeout** option can fail if the NFS mount is hard. If a hang occurs, deactivate the **nfstimeout** option and mount the NFS file system soft mounted, as follows:

```
mount -o soft,timeo=5,retry=5 machine:/filesystem /mountpoint
```

The parameters are defined as follows:

soft Generates a soft mount of the NFS file system. If an error occurs, the **stat()** function returns with an error. If the option **hard** is used, **stat()** does not return until the file system is available.

timeo=n
Sets the timeout period for a soft mount error to *n* tenths of a second.

retry=n
Sets the number of times to try the mount, where *n* is an integer; the default is 10000.

Deleted file systems

When a file system or drive has been deleted, or it is no longer backed up by the backup-archive client, the existing backup versions for each file are managed according to the following policy attributes: Number of days to keep inactive backup versions, and number of days to keep the last backup version (if there is no active version)

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

If you do not need to keep any of the backup versions, use the **delete backup** command to delete all backup versions in the file space. Your Tivoli Storage Manager server administrator must give you the authority to use this command. Use the **query session** command to determine whether you have "delete backup" authority. Alternatively, you can ask your Tivoli Storage Manager server administrator to delete the file space for you.

Related concepts:

Chapter 9, "Storage management policies," on page 275

Opened files

The backup-archive client looks for files that have changed between the start and the completion of the backup of the file.

Some files on your system might be in use, or open, when you try to back them up. Because an open file can change, a backup action might not reflect the correct contents of the file at a given time.

Consider whether the file is important, and whether you can build the file again. If the file is not important, you might not want to back it up. Or, if the file is important, a root user on your workstation can ensure the file is closed before backup.

If your backups run on a schedule, a root user can use the **preschedulecmd** option to enter a command to close the file. For example, if the open file is a database, use

the **quiesce** command of the database to shut down the database. A root user can use the `postschedulecmd` option to restart the application that uses the file after the backup completes. If you are not using a schedule for the backup, ensure that you close the application that uses the file before you start the backup.

The client can back up the file even if it is open and gets changed during the backup. This is only useful if the file is usable even if it changes during backup. To back up these files, assign the files a management class with the serialization *dynamic* or *shared dynamic*.

Related concepts:

“Display information about management classes and copy groups” on page 277

“Select a management class for files” on page 280

Wildcard characters

You can use the operating system wildcard characters in file specifications with the backup-archive client. These characters let you select groups of files that have similar names.

In a command, wildcard characters can only be used in the file name or extension. They cannot be used to specify destination files, file systems, or directories. When using wildcard characters in non-loop mode, as in `dsmc sel "/home/ledger.*"`, enclose the parameter containing the asterisk in quotation marks to ensure the system does not interpret the wildcard character and produce unexpected results. Wildcard character information is covered in the following table.

Important: Use an asterisk (*) instead of a question mark (?) as a wildcard character when trying to match a pattern on a multibyte code page, to avoid unexpected results.

This table shows some wildcard patterns and how to specify them.

* (Asterisk)	Zero or more characters that match all files:
*.cpp	With a cpp extension
hm*.*	Starting with hm, regardless of extension, but must have the '.' character
hm*	Starting with hm, whether an extension exists or not
h.*	With an h somewhere in the file name, regardless of extension, but must have .
? (Question mark)	One character that matches all files with:
?..cpp	The extension cpp with one, and only one, character in the file name
hm?.cpp	Three-character names beginning with hm and that have the cpp extension
* ? (Asterisk and question mark)	Asterisk and question mark combinations matching:
??hm.*	All four-character file names ending in hm., no matter what extension they have

In a path name for a file specification, you cannot specify a directory whose name contains an asterisk (*) or a question mark (?). The client recognizes those characters only as wildcard characters.

Chapter 5. Restoring your data

Use Tivoli Storage Manager to restore backup versions of specific files, a group of files with similar names, or entire directories.

You can restore these backup versions if the original files are lost or damaged. Select the files that you want to restore by using a file specification (file path, name, and extension), a directory list, or a subdirectory path to a directory and its subdirectories.

All client backup and restore procedures that are referenced by this topic also apply to the web client. However, the web client does not provide a Preferences Editor for setting client options.

Attention: Do not restore operating system files, like base system directories, kernel modules, or patches, to their original location while the file system is running. The operating system might hang or crash.

The following are the primary restore tasks:

- “Restoring an image”
- “Restoring data using the GUI” on page 239
- “Command line restore examples” on page 240
- “Restore data from a backup set” on page 224
- “Restoring data to a point in time” on page 233
- “Restore NAS file systems” on page 236
- “Authorizing another user to restore or retrieve your files” on page 244
- “Restoring or retrieving files from another client node” on page 245
- “Restore or retrieve files to another workstation” on page 246
- “Restoring a disk in case of disk loss” on page 246
- “Deleting file spaces” on page 247

Refer to *IBM Tivoli Storage Manager for Space Management for UNIX and Linux* for details about restoring migrated files and the `restoremigstate` option.

Related tasks:

“Starting a web client session” on page 140

Restoring an image

There are some items to consider before you begin restoring images on your system.

Before you restore an image (offline or online), you must have administrative authority on the system.

Here is a list of items to consider before you restore an image:

- Restoring the image of a volume restores the data to the same state that it was in when you performed your last image backup. Be absolutely sure that you need to restore an image, because it replaces your entire current file system or raw volume with the image on the server.

- Ensure that the volume to which you are restoring the image is at least as large as the image that is being restored.
- On Linux systems, some file systems such as ext2, ext3, ext4, btrfs, and xfs use a universally unique identifier (UUID) to identify themselves to the operating system. If you create an image backup of such a volume and you restore it to a different location, you might have two volumes with the same UUID. If you use UUID to define your file systems in `/etc/fstab`, be aware that the backup-archive client might be unable to correctly mount the restored file system because the UUIDs conflict. To avoid this situation, restore the image to its original location. If you must restore it to a different location, change the UUID of either the original or restored volume before you mount the restored file system. Refer to the Linux documentation for instructions on how to change a UUID. You might also need to manually edit the `/etc/fstab` file so the original volume, the restored volume, or both volumes can be mounted.
- The file system or volume you are restoring to must be the same type as the original.
- Ensure that the target volume of the restore is not in use. The client locks the volume before starting the restore. The client unlocks the volume after the restore completes. If the volume is in use when the client attempts to lock the file system, the restore fails.
- You cannot restore an image to where the Tivoli Storage Manager client program is installed.
- If you have run progressive incremental backups *and* image backups of your file system, you can perform an incremental image restore of the file system. The process restores individual files after the complete image is restored. The individual files restored are those backed up after the original image. Optionally, if files were deleted after the original backup, the incremental restore can delete those files from the base image.

Deletion of files is performed correctly if the backup copy group of the Tivoli Storage Manager server has enough versions for existing and deleted files. Incremental backups and restores can be performed only on mounted file systems, not on raw logical volumes.

- If for some reason a restored image is corrupted, you can use the `fsck` tool to attempt to repair the image.

You can use the `verifyimage` option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, the client issues a warning message on the console and in the error log.

If bad sectors are present on the target volume, you can use the `imagetofile` option with the **restore image** command to specify that you want to restore the source image to a file. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume.

Related reference:

"Imagetofile" on page 418

"Verifyimage" on page 565

Restoring an image using the GUI

You can use the GUI to restore an image of your file system or raw logical volume.

About this task

Follow these steps to restore an image of your file system or raw logical volume:

Procedure

1. Click **Restore** from the main window. The Restore window appears.
2. Expand the directory tree.
3. Locate the object in the tree named **Image** and expand it. Click the selection box next to the image you want to restore. You can obtain detailed information about the object by highlighting the object and selecting **View → File Details...** from the main window or click the **View File details** button.
4. **(Optional)** To perform an incremental image restore, click the **Options** button to open the Restore Options window and select the **Image plus incremental directories and files** option. If you want to delete inactive files from your local file system, select the **Delete inactive files from local** check box. Click the **OK** button.
5. Click **Restore**. The Restore Destination window appears. The image can be restored to the volume with the mount point from which it was originally backed up. Alternatively, a different volume can be chosen for the restore location.
6. Click the **Restore** button to begin the restore. The **Task List** window appears showing the progress of the restore. The Restore Report window displays a detailed status report.

Results

The following are some items to consider when you perform an image restore using the GUI:

- You can select **View → File Details** from the main window or click the **View File details** button to display the following statistics about file system images backed up by the client:
 - Image Size - This is the volume size which was backed up.
 - Stored Size - This is the actual image size stored on the server. The stored image on the Tivoli Storage Manager server is the same size as the volume capacity.
 - File system type
 - Backup date and time
 - Management class assigned to image backup
 - Whether the image backup is an active or inactive copy
- To modify specific restore options, click the **Options** button. Any options you change are effective during the current session *only*.
- In the Restore Options window, you can choose to restore the image only or the image and incremental directories files. If you choose **Image Only**, you restore the image from your last image backup only. This is the default.

If you ran incremental-by-date image backup on a volume or image backups on a volume with incrementals, you can choose the **Image plus incremental directories and files** option. If you choose **Image plus incremental directories and files**, you can also select **Delete inactive files from local** to delete the inactive files that are restored to your local file system. If incremental-by-date image backup was the only type of incremental backup you performed on the file system, deletion of files will not occur.

Important: Be absolutely sure that you need to perform an incremental restore because it replaces your entire file system with the image from the server and then restore the files that you backed up using the incremental image backup operation.

Restoring an image using the command line

Use the **restore image** command to restore an image using the Tivoli Storage Manager command line client.

Related reference:

“Imagetofile” on page 418

“Verifyimage” on page 565

Restore data from a backup set

Your Tivoli Storage Manager administrator can generate a backup set, which is a collection of your files that reside on the server, onto portable media created on a device using a format that is compatible with the client device.

You can restore data from a backup set from the Tivoli Storage Manager server, or when the backup set is locally available as a file or on a tape device.

You can restore backup sets from the following locations:

- From the Tivoli Storage Manager server
- From portable media on a device attached to your client workstation
- From a backup set file on your client workstation

Backup sets can provide you with instant archive and rapid recovery capability as described in the following list.

Instant archive

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

Rapid recovery with local backup sets

Typically, restores are performed from normal file backups that are stored on the Tivoli Storage Manager server outside of backup sets. This restore approach gives you the ability to restore the most recent backup version of every file. It is possible that a backup set does not contain the most recent backup version of your files.

In some cases restoring data from a backup set can be a better option than restoring data from normal backup files on the Tivoli Storage Manager server. Restoring from a backup set can be a better option for the following reasons:

- A backup set restore can provide for a faster recovery because all of the required files for restore are contained together within a smaller number of storage volumes.
- A backup set provides a point-in-time collection of files. You can restore to a point in time rather than restoring what is currently available from a normal file-level restore from the server.

Restoring a backup set from the Tivoli Storage Manager server provides a larger set of restore options than restoring from a local backup set. However, restoring from a local backup set can be preferable in some cases:

- It is possible that you need to restore your data when a network connection to the Tivoli Storage Manager server is not available. This is possible in a disaster recovery situation.

- The local restore may be faster than restoring over a network connection to your Tivoli Storage Manager server.

A backup set can be restored from the Tivoli Storage Manager server while the backup set volumes are available to the server, or they can be moved to the client system for a local backup set restore. A backup set can be generated with or without a table of contents (TOC), and can contain file data or image data.

Your ability to restore data from backup sets is restricted by the location of the backup set and the type of data in the backup set. The command-line client can restore some data that the GUI cannot restore, but the GUI can allow you to browse and choose which objects to restore. Generally, backup sets from the server with a TOC allow more options when restoring. However, local backup sets provide options that are sometimes preferable to restoring from the Tivoli Storage Manager server.

The restrictions for restoring data from backup sets using the GUI are summarized in the following table. Each interior cell represents one combination of data type and backup set location. For each situation, the cell indicates if you can use the GUI to restore only the entire backup set, to select objects within the backup set, or if you cannot use the GUI to restore the backup set.

Table 42. Backup set GUI restore restrictions

Data type in the backup set	Backup set location		
	Local (location=file or location=tape)	Tivoli Storage Manager Server (TOC available)	Tivoli Storage Manager Server (TOC not available)
file	Restore entire backup set only.	Restore entire backup set, or selected objects in the backup set.	Restore entire backup set only.
image	Cannot be restored.	Restore entire backup set, or selected objects in the backup set.	Cannot be restored.
system state	Restore entire backup set only.	Restore entire backup set, or selected objects in the backup set.	Restore entire backup set only.

The restrictions for restoring data from backup sets using the command-line client are summarized in the following table. Each interior cell represents one combination of data type and backup set location. For each situation, the cell lists the restore commands you can use. Except as noted, you can restore specific objects within a backup set, as well as the entire backup set.

Table 43. Backup set command-line restore restrictions

Data type in the backup set	Backup set location		
	Local (location=file or location=tape)	Tivoli Storage Manager Server (TOC available)	Tivoli Storage Manager Server (TOC not available)
file	Commands: restore restore backupset	Commands: restore restore backupset	Commands: restore backupset
image	Cannot be restored	Command: restore image	Cannot be restored

Table 43. Backup set command-line restore restrictions (continued)

Data type in the backup set	Backup set location		
	Local (location=file or location=tape)	Tivoli Storage Manager Server (TOC available)	Tivoli Storage Manager Server (TOC not available)
system state	Command: restore backupset	Commands: restore backupset restore systemstate	Command: restore backupset

Restriction: When restoring system state data using the **restore backupset** command, you cannot specify individual objects. You can only restore the entire system state.

Related reference:

“Localbackupset” on page 446

“Query Backupset” on page 670

“Query Image” on page 678

“Restore” on page 693

“Restore Backupset” on page 698

“Restore Image” on page 707

Restore backup sets: considerations and restrictions

This topic lists some considerations and restrictions that you must be aware of when restoring backup sets.

Backup set restore considerations

Consider the following when restoring backup sets:

- If the object you want to restore was generated from a client node whose name is different from your current node, specify the original node name with the **filespace** parameter on any of the restore commands.
- If you are unable to restore a backup set from portable media, check with your Tivoli Storage Manager administrator to ensure that the portable media was created on a device using a compatible format.
- If you use the **restore backupset** command on the initial command line with the parameter **-location=tape** or **-location=file**, the client does not attempt to contact the Tivoli Storage Manager server.
- When restoring a group from a backup set:
 - The entire group, or all groups, in the virtual file space are restored. You cannot restore a single group by specifying the group name, if there are several groups in the same virtual file space. You cannot restore a part of a group by specifying a file path.
 - Specify a group by using the following values:
 - Specify the virtual file space name with the **filespace** parameter.
 - Use the **subdir** option to include subdirectories.
- Limited support is provided for restoring backup sets from tape devices attached to the client system. A native device driver provided by the device manufacturer must always be used. The device driver provided by Tivoli to be used with the Tivoli Storage Manager server cannot be used on the client system for restoring local backup sets.

- If a backup set contains files from several owners, the backup set itself is owned by the root user ID, and non-root user IDs cannot see the backup set. In this case, non-root user IDs can restore their files by obtaining the backup set name from the Tivoli Storage Manager administrator. Non-root users can restore only their own files.
- To enable the client GUI to restore a backup set from a local device, without requiring a server connection, use the `localbackupset` option.

Backup set restore restrictions

Be aware of the following restrictions when restoring backup sets:

- A backup set data that was backed up with the API cannot be restored or used.
- You cannot restore image data from a backup set using the **restore backupset** command. You can restore image data from a backup set only with the **restore image** command.
- You cannot restore image data from a local backup set (`location=tape` or `location=file`). You can restore image data from a backup set only from the Tivoli Storage Manager server.

Related reference:

“**Localbackupset**” on page 446

“**Restore**” on page 693

“**Restore Image**” on page 707

“**Restore Backupset**” on page 698

Backup set restore

Tivoli Storage Manager considers a backup set as one object containing the whole file structure. You can restore the entire backup set or, in some cases, you can select portions. The backup set media is self-describing and contains all the information required to perform a successful restore.

If you are connected to a Tivoli Storage Manager Version 5.4 or later server, your Tivoli Storage Manager administrator can create backup sets that are stacked. Stacked backup sets can contain data from multiple client nodes, and they can contain different types of data for a particular client node. The types of data can be file data or image data.

Restriction: Image data and application data restore processing is only available when restoring from the Tivoli Storage Manager server. You cannot restore image data and application data from a client local backup set restore.

When a backup set is stacked, you can only restore data for your own node. Data for all other nodes is skipped. When restoring data from a stacked backup set on a local device, you can only restore file level data for your own client node. It is important that the `nodename` option is set to match the node name used to generate the backup set for one of the nodes in the stack.

Important: Due to the portability of local backup sets, you must take additional steps to secure your local backup sets on portable media. The backup set media should be physically secured because the backup set can be restored locally without authenticating with the server. Each user has access to all of the data on the stacked backup set, which means that the user has access to data that they do

not own, by changing the node name or viewing the backup set in its raw format. Encryption or physical protection of the media are the only methods to ensure that the data is protected.

If you restore backup set data from the server, individual files, directories or entire backup set data can be restored in a single operation from the GUI or the command line. When you restore backup set data locally, the GUI can only display and restore an entire backup set. The command line can be used to restore individual files or directories stored in a backup set locally.

Restoring backup sets using the GUI

The client GUI can restore data from a backup set from the server, from a local file, or from a local tape device. You can use the GUI to restore individual files from a backup set from the Tivoli Storage Manager server with a TOC, but not from a local backup set nor from a backup set from the server without a TOC.

About this task

Important: Before you begin a restore operation, be aware that backup sets can contain data for multiple file spaces. If you specify a destination other than the original location, data from *all* file spaces are restored to the location you specify.

To restore a backup set from the GUI, perform the following steps:

1. Click **Restore** from the GUI main window. The Restore window appears.
2. Locate the **Backup Sets** directory tree object and expand it by clicking the plus sign (+) beside it.
 - To restore the backup set from a local device, expand the **Local** object and the Specify backup set location window is displayed. On the window, select **File name:** or **Tape name:** from the list and enter the tape or file name location. You can also click the **Browse** button to open a file selection window and select a backup set.
 - To restore data from backup set from the server, first expand the **Server** object and then either **Filelevel** or **Image**, depending on the type of restore requested.
3. Click the selection box next to the backup set or directory or file within the backup set that you want to restore.

You can select files from within a backup set if that backup set is from the server and has a table of contents.
4. Click **Restore**. The Restore Destination window appears. Enter the appropriate information.
5. Click **Restore**. The Task List window displays the restore processing status.

Note:

- If the object you want to restore is part of a backup set generated on a node, and the node name is changed on the server, any backup set objects that were generated prior to the name change will not match the new node name. Ensure that the node name is the same as the node for which the backup set was generated.
- The client can be used to restore a backup set on an attached device with or without a server connection. If the server connection fails, a prompt appears to continue for purposes of local backup set restore. Also, the `localbackupset` option can be used to tell the client not to attempt the connection to the server.

- Certain local devices such as tape devices (tape devices do not apply to Mac OS X) require device drivers to be set up prior to performing a restore. See the device manual for assistance with this task. You also need to know the device address in order to perform the restore.
- The following features of a backup set restore from the server are not available when restoring locally:
 1. Image restore.
 2. The GUI display and restore of individual files and directories. The command line can be used to restore an individual directory or file from a local backup set.

Backup set restores using the client command-line interface

The client command line interface can restore data from a backup set from the server, from a local file, or from a local tape device. You can use the client command line interface to restore individual files from local backup sets and from backup sets without a TOC.

To restore a backup set from the client command line interface, use the **query backupset** command to display what backup set data is available, then use restore commands to restore the data.

You can use the following commands to restore data from backup sets:

- **restore**
- **restore backupset**
- **restore image**

Use the appropriate command for the location of the backup set and the data in the backup set. For more information, see Table 43 on page 225.

Related reference:

“Query Backupset” on page 670

“Query Image” on page 678

“Restore” on page 693

“Restore Backupset” on page 698

“Restore Image” on page 707

Restoring or retrieving data during a failover

When the client fails over to the secondary server, you can restore or retrieve replicated data from the secondary server.

Before you begin

Before you begin to restore or retrieve data during a failover:

- Ensure that the client is configured for automated client failover.
- Ensure that you are connected to a Tivoli Storage Manager server that replicates client nodes. For more information about failover requirements, see “Requirements for automated client failover” on page 88.

Restriction: In failover mode, you cannot back up or archive data to the secondary server.

Procedure

To restore or retrieve data during a failover, complete the following steps:

1. Verify the replication status of the client data on the secondary server. The replication status indicates whether the most recent backup was replicated to the secondary server.
2. Restore or retrieve your data as you would normally do from the client GUI or from the command-line interface.

Tip: Restartable restore operations function as expected when you are connected to the secondary server. However, restore operations that are interrupted when the primary server goes down cannot be restarted after the client fails over. You must run the whole restore operation again after the client fails over to the secondary server.

Results

If the replicated data on the secondary server is not current, you are prompted to continue or to stop the restore or retrieve operation.

For example, to restore the `build.sh` directory at the command-line interface, you issue the following command:

```
dsmc res /build.sh
```

The following output is displayed:

```
IBM Tivoli Storage Manager
Command Line Backup-Archive Client Interface
  Client Version 7, Release 1, Level 0.0
  Client date/time: 05/21/2013 14:33:54
(c) Copyright by IBM Corporation and other(s) 1990, 2013. All Rights Reserved.

Node Name: MY_NODE_NAME
ANS2106I Connection to primary TSM server 192.0.2.1 failed

ANS2107I Attempting to connect to secondary server TARGET at
192.0.2.9 : 1501

Node Name: MY_NODE_NAME
Session established with server TARGET: Windows
  Server Version 7, Release 1, Level 0.0
  Server date/time: 05/21/2013 14:33:54  Last access: 05/21/2013 14:13:32

  Session established in failover mode to secondary server
ANS2108I Connected to secondary server TARGET.
Restore function invoked.

ANS2120W The last store operation date reported by the server TARGET of
05/16/2013 22:38:23 does not match the last store operation date of
05/21/2013 21:32:20 stored by the client.
Continue (Yes (Y)/No (N))
```

If you respond with N, the following message is displayed:

```
ANS1074W The operation was stopped by the user.
```

If you respond with Y, restore processing continues as normal, but the data that you restore might not be the most current.

Related concepts:

“Automated client failover configuration and use” on page 87

Related tasks:

“Determining the status of replicated client data” on page 92

Restore an image to file

When you back up an image, Tivoli Storage Manager backs up the first sector of the volume, but when the data is restored, it skips the first sector to preserve the original logical volume control block of the destination volume.

When you restore an image to file, entire volume contents, including the first sector, are restored to the file.

AIX LVM volumes from original volume groups contain the Logical Volume Control Block (LVCB) on the first sector (512 bytes) of the volume. The LVCB contains volume specific meta-data that should be preserved by applications using the volume.

When you copy the file, containing the image, onto an LVM volume from the original volume group, you need to skip the LVCB from both the file and destination volume. The following **dd** command can be used for this purpose.

```
dd if=<filename> of=/dev/<vol> bs=512 skip=1 seek=1
```

The **dd** command sets the block size to 512 bytes, which makes copying very slow. It is better to use **bs=1m** or similar. Here is an alternative way to copy image data:

1. Save the original first sector to a file:

```
dd if=/dev/<vol> of=firstblk.tmp bs=512 count=1
```

2. Copy the restored image:

```
dd if=<filename> of=/dev/<vol> bs=1m
```

3. Restore the original first sector:

```
dd if=firstblk.tmp of=/dev/<vol> bs=512 count=1
```

With the introduction of big and scalable volume group formats on AIX, it is possible that the first sector of the logical volume cannot contain LVCB and is available for the data. If you use big or scalable volume groups on your system, and need to restore the whole volume including the first sector, restore the volume to file and then copy it to a destination volume. The following **dd** command can be used for this purpose.

```
dd if=<filename> of=/dev/<vol> bs=1m
```

Related concepts:

“Restoring an image using the command line” on page 224

Related tasks:

“Restoring an image using the GUI” on page 222

Manage GPFS file system data with storage pools

With Global Parallel File Systems (GPFS) technology, you can manage your data using storage pools. A storage pool is a collection of disks or RAID configurations with similar properties that are managed together as a group.

The group under which the storage pools are managed together is the file system. The automated placement and management of files on the storage pool level is done by policies. A policy is a set of rules that describes the life cycle of user data, based on the attributes of the file.

When a file is created, the placement policy determines the initial location of the data of the file and assigns the file to a storage pool. All data written to that file is placed in the assigned storage pool. The management policy determines file management operation, such as migration and deletion. The files within a GPFS file system are distributed over different storage pools, depending on the enabled placement and migration policies.

During restore, the files are placed on the correct storage pool. The Tivoli Storage Manager server is not aware of pool-to-pool migrations, so the files are placed on the storage pool from where the backup has taken place. The policy engine replaces the files based on migration policies.

If a storage pool ID is stored in the extended attributes of the file, and that storage pool is available, the file is always placed in that storage pool. If the storage pool is not available, the file is placed according to the placement policy. If the placement policy does not match the file, the file is placed in the system pool.

GPFS handles the placement of files after a restore as follows:

- The file is placed in the pool that can be selected by matching the saved file attributes to a RESTORE rule
- The file is placed in the pool that it was in when it was backed up
- The file is placed based on the current placement policy
- The file is placed in the system storage pool

The GPFS RESTORE rule allows you to match files against their saved attributes rather than the current file attributes. If the file attributes do not match, GPFS tries to restore the file in the sequence described above.

For more information about the GPFS RESTORE rule, read the GPFS documentation about policies and rules.

The following restrictions apply:

- The restore of stub files does not work with multiple storage pools, or with files that have ACLs
- Unlink of filesets are not allowed
- The ctime option of GPFS should be set to no (default), to prevent unwanted Backup-Archive backups of files after GPFS file migration from pool to pool

For information about using storage pools, see the Tivoli Storage Manager server documentation.

Related reference:

 [Data storage in storage pools](#)

Related information:

 [GPFS product information](#)

 [mmbackup command: Tivoli Storage Manager requirements](#)

 [Using IBM Spectrum Protect include and exclude options with IBM Spectrum Scale mmbackup command](#)

Restoring data to a point in time

Use a *point-in-time* restore to restore files to the state that existed at a specific date and time.

About this task

A point-in-time restore can eliminate the effect of data corruption by restoring data from a time prior to known corruption, or recover a basic configuration to a prior condition.

You can perform a point-in-time restore of a file space, directory, or file.

You can also perform a point-in-time restore of image backups.

Perform incremental backups to support a point-in-time restore. During an incremental backup, the backup-archive 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 before the oldest version maintained by the Tivoli Storage Manager server, the object is not restored to your system. Files that were deleted from your workstation before the point-in-time specified are not restored.

Note:

1. Your administrator must define copy group settings that maintain enough inactive versions of a file to guarantee that you can restore that file to a specific date and time. If enough versions are not maintained, the client might not be able to restore all objects to the point-in-time you specify.
2. If you delete a file or directory, the next time you run an incremental backup, the active backup version becomes inactive and the oldest versions that exceed the number specified by the *versions data deleted* attribute of the management class are deleted.

When you perform a point-in-time restore, consider the following information:

- The client 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, the client cannot restore that object.
- Point-in-time restore restores files that were deleted from the client workstation after the point-in-time date but not files that were deleted before this date.
- The client cannot restore a file that was 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.

Procedure

To perform a point-in-time restore by using the client GUI, complete 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 that 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.

Results

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 by using the `pitdate` and `pittime` options with the **query backup** and **restore** commands. For example, when you use the `pitdate` and `pittime` options with the **query backup** command, you establish the point-in-time for which file information is returned. When you use `pitdate` and `pittime` with the **restore** command, the date and time values you specify establish the point-in-time for which files are returned. If you specify `pitdate` without a `pittime` value, `pittime` defaults to 23:59:59. If you specify `pittime` without a `pitdate` value, it is ignored.

Related concepts:

Chapter 9, "Storage management policies," on page 275

Related reference:

"**Backup Image**" on page 620

Restore AIX encrypted files

When files are backed up in raw format from an AIX JFS2 Encrypted File System (EFS), you can only restore them to the same or another JFS2 EFS. They cannot be restored to any different file system, or on a different platform.

When EFS files are backed up in clear text, then you can restore them anywhere. If you restore them to a JFS2 EFS, they are automatically re-encrypted only if the directory to which they are restored has the AIX "EFS inheritance" option set.

After restoring a file that was backed up in raw format, you might find that the file cannot be decrypted. The encryption key originally used for the file might no longer be available in the keystore of the user. In this case, you must restore the keystore used at the time of backup. Restore it to a location different from `/var/efs`, and use the AIX command `"efskeymgr -f "` to add its keys to the kernel key ring.

For information on backing up EFS data, refer to "AIX JFS2 encrypted file system backup" on page 200.

Restore AIX workload partition file systems

All the files that are created by the local workload partition (WPAR), and backed up by the backup-archive client that is installed at the global WPAR, can be restored by the client installed at the global WPAR.

Here are some global partition and WPAR configuration examples:

Global partition:

```
system name: shimla
file system: /home /opt
```

WPAR #1 configuration:

```
name: wpar1
file system: /home; name in global WPAR: /wpars/wpar1/home
```

WPAR #2 configuration:

```
name: wpar2
file system: /data; name in global WPAR: /wpars/wpar2/data
```

There are two ways to restore WPAR data, depending on the method used to back up the WPAR data files:

- Restore all WPAR file systems as the file spaces within the global partition. The file space name must be used to identify the WPAR to which it belongs. All of the data is managed on one node using one schedule. Using the example configuration mentioned previously, here is a sample `dsm.sys` file with one server stanza for all file systems, both global and local:

```
Servname shimla
TCPPort 1500
TCPServeraddress server.example.com
nodename shimla
PasswordAccess generate
Domain /wpars/wpar1/home /wpars/wpar2/data /home /opt
```

Use the following command to restore each file space:

```
dsmc restore /wpars/wpar1/home/*
dsmc restore /wpars/wpar2/data/*
dsmc restore /home/*
dsmc restore /opt/
```

- Restore each WPAR file system from a different node name, if it is backed up under a different node name. Each WPAR must have a separate node name and a scheduler running within the global partition. Also, three scheduler services must be set up, each using a different `dsm.opt` file corresponding to the server stanza name. This method allows each WPAR restore operation to be managed independent of the others. Using the example configuration mentioned previously, here is a sample `dsm.sys` file with three server stanzas: one for `wpar1`, one for `wpar2`, and one for global partition `shimla`:

```
Servname shimla_wpar1
TCPPort 1500
TCPServeraddress server.example.com
nodename wpar1
```

```

PasswordAccess generate
Domain      /wpars/wpar1/home

Servername  shimla_wpar2
TCPPort     1500
TCPServeraddress server.example.com
nodename    wpar2
PasswordAccess generate
Domain      /wpars/wpar2/data

Servername  shimla
TCPPort     1500
TCPServeraddress server.example.com
nodename    shimla
PasswordAccess generate
Domain      /home /opt

```

Table 44. Sample WPAR restore commands with dsm.opt file

In dsm.opt file	Sample restore command
servername shimla_wpar1	dsmc restore /wpars/wpar1/home/*
servername shimla_wpar2	dsmc restore /wpars/wpar2/data/*
servername shimla	dsmc restore /home/* dsmc restore /opt/*

Restore NAS file systems

You restore NAS file system images using the web client or command line interface. The web client interface is available only for connections to Tivoli Storage Manager V7.1.7 and earlier servers, or V8.1.1 and earlier V8 servers.

You can restore full or differential NAS file system images that were backed up previously. If you restore a differential image, Tivoli Storage Manager automatically restores the full backup image first, followed by the differential image. It is not necessary for a client node to mount a NAS file system to perform backup or restore operations on that file system.

Related concepts:

“Web client configuration overview” on page 59

Restoring NAS file systems using the web client

This section lists the steps to follow to restore NAS file systems using the web client GUI.

Before you begin

The web client interface is available only for connections to Tivoli Storage Manager V7.1.7 and earlier servers, or V8.1.1 and earlier V8 servers.

Procedure

1. Click the **Restore** button from the main window. The Restore window appears.
2. Expand the directory tree if necessary. To expand a node in the tree, click the plus sign (+) next to an object in the tree. Nodes shown are those that have been backed up and to which your administrator has authority. The root node called **Nodes** is not selectable. This node only appears if a NAS plug-in is

present on the client workstation. NAS nodes display on the same level as the node of the client workstation. Only nodes to which the administrator has authority appear.

3. Expand the NAS node to reveal the Image object.
4. Expand the Image object to display volumes that you can restore. You cannot expand Volume objects.
5. Click the selection boxes next to the volumes under the Image object that you want to restore. If you want to restore a NAS image that was backed up on a particular date, click the **Point In Time** button. After you select a date, the last object that was backed up on or prior to that date appears, including any inactive objects. If you want to display all images (including active images and inactive images), before you select them, select **View → Display active/inactive files** from the menu bar.
6. Click **Restore**. The Restore Destination window appears. Enter the information in the Restore Destination window. If you choose to restore to a different destination, you can only restore one volume at a time to a different destination. You can restore NAS file system images to any volume on the NAS file server from which they were backed up. You cannot restore images to another NAS file server.
7. Click **Restore**. The NAS Restore **Task List** window displays the restore processing status and progress bar. If there is a number next to the progress bar, it indicates the size of the restore, if known. After the restore completes, the NAS Restore Report window displays processing details. If you must close the web browser session, current NAS operations continue after you disconnect. You can use the **Dismiss** button on the NAS Restore **Task List** window to quit monitoring processes without ending the current operation.
8. (Optional) To monitor processing of an operation, select the **Actions → TSM Activities** from the main window.

Results

Considerations:

- Workstation and remote (NAS) backups are mutually exclusive in a Restore window. After selecting an item for restore, the next item you select must be of the same type (either NAS or non NAS).
- Details will not appear in the right-frame of the Restore window for NAS nodes or images. To view information about a NAS image, highlight the NAS image and select **View → File Details** from the menu.
- To delete NAS file spaces, select **Utilities → Delete Filespaces**. You can delete both workstation and remote objects.

Options and commands to restore NAS file systems from the command line

This topic lists some examples of options and commands you can use to restore NAS file system images from the command line.

Table 45. NAS options and commands

Option or command	Definition	Page
query node	Displays all the nodes for which a particular administrative user ID has authority to perform operations. The administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web client.	"Query Node" on page 681
query backup	Use the query backup command with the class option to display information about file system images backed up for a NAS file server.	"Query Backup" on page 668
query filesystem	Use the query filesystem command with the class option to display a list of file spaces belonging to a NAS node.	"Query Filespace" on page 674
restore nas	Restores the image of a file system belonging to a Network Attached Storage (NAS) file server.	"Restore NAS" on page 710
monitor process	Displays current backup and restore processes for all NAS nodes for which an administrative user has authority. The administrative user can then select one process to monitor.	"Monitor Process" on page 662
cancel process	Displays current backup and restore processes for all NAS nodes for which an administrative user has authority. From the display, the administrative user can select one process to cancel.	"Cancel Process" on page 641
delete filesystem	Use the delete filesystem with the class option to display a list of file spaces belonging to a NAS node so that you can choose one to delete.	"Delete Filespace" on page 648

Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol1/vol10.

Note: When you initiate a NAS restore operation using the command line client or the web client, the server starts a process to initiate, control, and monitor the operation. It might take several moments before you notice progress at the client command line interface because the server must perform a mount and other necessary tasks before data movement occurs. The Tivoli Storage Manager command line client might display an Interrupted ... message when the mount occurs. You can ignore this message.

Restore active or inactive backups

Your administrator determines how many backup versions Tivoli Storage Manager maintains for each file on your workstation.

Having multiple versions of a file permits you to restore older versions if the most recent backup is damaged. The most recent backup version is the *active* version. Any other backup version is an *inactive* version.

Every time Tivoli Storage Manager backs up your files, it marks the new backup version as the active backup, and the last active backup becomes an inactive backup. When the maximum number of inactive versions is reached, Tivoli Storage Manager deletes the oldest inactive version.

To restore a backup version that is inactive, you must display both active and inactive versions by clicking on the **View** menu → **Display active/inactive files** item. To display only the active versions (the default), click on the **View** menu → **Display active files only** item. If you try to restore more than one version at a time, only the active version is restored.

On the Tivoli Storage Manager command line, use the `inactive` option to display both active and inactive objects.

Related reference:

“Inactive” on page 419

Restoring data using the GUI

This section lists the steps to follow to restore backup versions of individual files or subdirectories.

Procedure

1. Click **Restore** from the main window. The Restore window appears.
2. Expand the directory tree. Select the selection boxes next to the files or directories you want to restore. To search or filter files, click the **Find** icon on the tool bar.
3. Enter your search criteria in the Find Files (Restore) window.
4. Click the **Search** button. The Matching Files (Restore) window appears.
5. Click the selection boxes next to the files you want to restore and close the Matching Files (Restore) window.
6. Enter your filter criteria in the Find Files (Restore) window.
7. Click the **Filter** button. The Restore window displays the filtered files.
8. Click the selection boxes next to the filtered files or directories you want to restore.
9. To modify specific restore options, click the **Options** button. Any options you change are effective during the current session *only*.
10. Click **Restore**. The Restore Destination window appears. Enter the information in the Restore Destination window.
11. Click **Restore**. The Restore **Task List** window displays the restore processing status.

Results

Note: On Mac OS X, consider the following items when restoring data using the GUI:

1. When **TSM Tools for Administrators** is used to start the client, the client is running with a UID of zero. This means that if you create a folder to restore your files to, that folder is owned by root. To access the files you must change the permissions of the folder. You can change the folder owner from a terminal window using the `sudo chown` command. See your operating system documentation for more information on how to accomplish this.

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

Command line restore examples

This topic lists some examples of **restore** commands to use for specific tasks.

The following table shows examples of how to use the **restore** command to restore objects from Tivoli Storage Manager server storage.

Table 46. Command-line restore examples

Task	Command	Considerations
Restore the most recent backup version of the /Users/monnett/Documents/h1.doc file, even if the backup is inactive.	<code>dsmc restore /Users/monnett/Documents/h1.doc -latest</code>	If the file you are restoring no longer resides on your workstation, and you have run an incremental backup since deleting the file, there is no active backup of the file on the server. In this case, use the latest option to restore the most recent backup version. Tivoli Storage Manager restores the latest backup version, whether it is active or inactive. See "Latest" on page 445 for more information.
Display a list of active and inactive backup versions of files from which you can select versions to restore.	<code>dsmc restore "/Users/monnett/Documents/*"-pick -inactive</code>	If you try to restore both an active and inactive version of a file at the same time, only the active version is restored. See "Pick" on page 472 and "Inactive" on page 419 for more information.
Restore the /Users/monnett/Documents/h1.doc file to its original directory.	<code>dsmc restore /Users/monnett/Documents/h1.doc</code>	If you do not specify a destination, the files are restored to their original location.
Restore the /Users/monnett/Documents/h1.doc file under a new name and directory.	<code>dsmc restore /Users/monnett/Documents/h1.doc /Users/gordon/Documents/h2.doc</code>	None
Restore the files in the /Users directory and all of its subdirectories.	<code>dsmc restore /Users/ -subdir=yes</code>	When restoring a specific path and file, Tivoli Storage Manager recursively restores <i>all</i> subdirectories under that path, and any instances of the specified file that exist under <i>any</i> of those subdirectories. See "Subdir" on page 539 for more information about the <i>subdir</i> option.
Restore all files in the /Users/gordon/Documents directory to their state as of 1:00 PM on August 17, 2003.	<code>dsmc restore -pitd=8/17/2003 -pitt=13:00:00 /Users/gordon/Documents/</code>	See "Pitdate" on page 473 and "Pittime" on page 474 for more information about the pitdate and pittime options.

Table 46. Command-line restore examples (continued)

Task	Command	Considerations
Restore all files from the /Users/mike/Documents directory that end with .bak to the /Users/mike/projectn/ directory.	<code>dsmc restore "/Users/mike/Documents/*.bak" /Users/mike/projectn/</code>	If the destination is a directory, specify the delimiter (/) as the last character of the destination. If you omit the delimiter and your specified source is a directory or a file spec with a wildcard, you receive an error. If the projectn directory does not exist, it is created.
Restore files specified in the restorelist.txt file to a different location.	<code>dsmc restore -filelist=/Users/user2/Documents/restorelist.txt /Users/NewRestoreLocation/</code>	See “Filelist” on page 402 for more information about restoring a list of files.

Related reference:

“Restore” on page 693

Examples: Command line restores for large amounts of data

If you need to restore a large number of files, you can get faster performance by using the **restore** command instead of the GUI. In addition, you can improve performance by entering multiple **restore** commands at one time.

For example, to restore all the files in your /home file system, enter:

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

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

For example, you could enter these commands:

```
dsmc restore /home/monnett/ -subdir=yes -replace=all -tapeprompt=no
dsmc restore /home/gillis/ -subdir=yes -replace=all -tapeprompt=no
dsmc restore /home/stewart/ -subdir=yes -replace=all -tapeprompt=no
```

You can also use the quiet option with the **restore** commands to save processing time. However, you will not receive informational messages for individual files.

Note: If you already have the appropriate values set for the `subdir`, `replace`, `tapeprompt`, and `quiet` options in your client user-options file, you do not need to include those options in the commands.

When you enter multiple commands to restore your files, you must specify a unique part of the file space in each **restore** command. Be sure you do not use any overlapping file specifications in the commands.

To display a list of the directories in a file space, use the **query backup** command. For example:

```
dsmc query backup -dironly -subdir=no /Users/
```

As a general rule, you can enter from two to four **restore** commands at one time. The maximum number you can run at one time without degrading performance depends on factors such as how much memory you have and network utilization.

The speed at which you can restore the files also depends on how many tape drives are available on the server, and whether your administrator is using collocation to keep file spaces assigned to as few volumes as possible.

For example, if /Users/user1 and /Users/user2 are on the same tape, the restore for /Users/user2 must wait until the restore for /Users/user1 is complete. However, if /Users/user3 is on a different tape, and there are at least two tape drives available, the restore for /Users/user3 can begin at the same time as the restore for /Users/user1.

Set the system ulimit values to unlimited (-1) if you are restoring very large (2 GB) files with HSM or the backup-archive client. The client can restore these large files with enough system resources. If the ulimits are set to lower values, there might be restore failures.

Standard query restore, no-query restore, and restartable restore

This topic describes the standard (or classic) restore method, the no-query restore method, and the restartable restore method.

Standard query restore process

The standard query restore process is also known as classic restore. This topic explains how standard query restore works.

Here is how standard query restore works:

- The client queries the server for a list of files backed up for the client file space you want to restore.
- The server sends a list of backed up files that match the restore criteria. If you want to restore both active and inactive files, the server sends information about all backed up files to the client.
- The list of files returned from the server is sorted in client memory to determine the file restore order and to minimize tape mounts required to perform the restore.
- The client tells the server to restore file data and directory objects.
- The directories and files you want to restore are sent from the server to the client.

No-query restore process

In the no-query restore process, a single restore request is sent to the server instead of querying the server for each object to be restored.

1. The client tells the server that a no-query restore is going to be completed and provides the server with details about file spaces, directories, and files.
2. The server uses a separate table to track entries which guide the restore.
3. The data to be restored is sent to the client. File and directory objects that are stored on disk are sent immediately since sorting for such data is not required before the object is restored.
4. You can use multiple sessions to restore the data. If the data is on multiple tapes, there are multiple mount points available at the server. The combination of using the **resourceutilization** option and **MAXNUMMP** allows multiple sessions.

When you enter an unrestricted wildcard source file specification on the **restore** command and do not specify any of the options: **inactive**, **latest**, **pick**, **fromdate**, **todate**, the client uses a *no-query restore* method for restoring files and directories from the server. This method is called no-query restore because instead of querying the server for each object to be restored, a single restore request is sent to the server. In this case, the server returns the files and directories to the client without

further action by the client. The client merely accepts the data that comes from the server and restores it to the destination named on the **restore** command.

Using the Tivoli Storage Manager GUI client, an example of an unrestricted wildcard command would be to select a folder from the restore tree window. An example of a restricted wildcard command would be to select individual files from a folder.

Using the command-line client, an example of an unrestricted wildcard command would be:

```
"/Users/user1/Documents/2004/*"  
/home/mydocs/2004/*
```

An example of a restricted wildcard file specification would be:

```
/Users/user1/Documents/2004/sales.*  
/home/mydocs/2004/sales.*
```

Restartable restore process

If the restore process stops because of a power outage or network failure, the server records the point at which this occurred.

This record is known to the client as a *restartable restore*. It is possible to have more than one restartable restore session. Use the **query restore** command or choose **restartable restores** from the Actions menu to find out if your client has any restartable restore sessions in the server database.

You must complete a restartable restore before attempting further backups of the file system. If you attempt to repeat the restore that was interrupted or try to back up the destination file space, the attempt fails because you did not complete the original restore. You can restart the restore at the point of interruption by entering the **restart restore** command, or you can delete the restartable restore using the **cancel restore** command.

From the Tivoli Storage Manager GUI **Restartable restores** dialog box you can select the interrupted restore and delete it, or you can choose to restart the restore. If you restart the interrupted restore, it restarts with the first transaction, which might consist of one or more files, not completely restored when the interruption occurred. Because of this, you might receive some replace prompts for files from the interrupted transaction which were already restored.

To perform restartable restores using the GUI, follow these steps:

1. Select **Actions** → **Restartable restores** from the main panel.
2. Select the restartable restore session you want to complete.
3. Click the **Restart** button at the bottom of the panel.

Related reference:

"Resourceutilization" on page 499

"Restore" on page 693

Restoring Solaris Zettabyte (ZFS) file systems

Zettabyte File Systems (ZFS) use storage pools to manage physical storage.

How you restore a ZFS file system depends on how it was backed up.

- If you backed up all files and folders as separate objects, you can restore them by performing a file-level restore. For example:

```
dsmc restore /tank/myZFS/ -subdir=yes -replace=all
```

Do not perform a file-level restore operation in a disaster recovery scenario. Even though you successfully restore all system files and folders from a Tivoli Storage Manager client-created backup, the restored system might be unstable or fail.

- If you backed up an entire ZFS snapshot as a single file, you need to restore the snapshot file from the server into a temporary location. For example:

```
dsmc restore /tmpdir/mySnapshotfile
```

You can then restore the file system from the snapshot file by using the Oracle Solaris ZFS commands. For example:

```
zfs receive tank/myZFS@mySnapshot < /tmpdir/mySnapshotFile
```

The advantage of restoring ZFS from a snapshot file is that the full file system can be restored, in a disaster recovery scenario.

For detailed information about restoring data on ZFS file systems, see the product documentation that is available from Oracle. If you are restoring a ZFS root pool, see the topics that describe how to re-create your root pool and recover root pool snapshots.

Related tasks:

“Backing up Solaris Zettabyte file systems” on page 199

Additional restore tasks

This section discusses some advanced considerations for restoring data.

Authorizing another user to restore or retrieve your files

You can authorize another user on the same workstation or a different workstation to restore backup versions or retrieve archive copies of your files.

About this task

This permits you to share files with other people or with other workstations that you use with a different node name. To authorize a user on another workstation to restore or retrieve your files, the other workstation must be running one of the UNIX clients and must be registered with your server.

Note: Mac OS X can *only* restore Mac OS X nodes.

To authorize another user to restore or retrieve your files:

Procedure

1. Click **Utilities** → **Node Access List** from the main window. The Node Access List window appears.
2. Click the **Add** button. The Add Access Rule window appears.
3. In the Add Access Rule window, select an item in the Permit Access to field to specify the type of data that the other user can access. You can select either Backed up Objects or Archived Objects.
4. In the Grant Access to Node field, type the node name of the host workstation of the user that can access your data.

5. In the User field, type the name of the user on a node who can access your data.
6. In the Filespace and Directory field, select the file space and the directory that the user can access. You can select one file space and one directory at a time. If you want to give the user access to another file space or directory, you must create another access rule.
7. If you want to limit the user to specific files in the directory, type the name or pattern of the files on the server that the other user can access in the Filename field. You can make only one entry in the Filename field. It can either be a single file name or a pattern which matches one or more files. You can use a wildcard character as part of the pattern. Your entry must match files that have been stored on the server.
8. For the Java GUI: If you want to give access to all files that match the file name specification within the selected directory including its subdirectories, click **Include subdirectories**.
9. Click the **OK** button to save the access rule and close the Add Access Rule window.
10. The access rule that you created is displayed in the list box in the Node Access List window. When you have finished working with the Node Access List window, click the **OK** button. If you do not want to save your changes, click **Cancel** or close the window.

Results

In the client command line interface, use the **set access** command to authorize another node to restore or retrieve your files. You can also use the **query access** command to see your current list, and **delete access** to delete nodes from the list.

Related reference:

“Delete Access” on page 642

“Query Access” on page 665

“Set Access” on page 728

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.

About this task

You can display file spaces of another user on the server, restore the backup versions of another user, or retrieve the archive copies of another user to your local file system:

Procedure

1. Click **Utilities** from the main window.
2. Click **Access Another Node**. The Access Another Node window appears.
3. Type the node name of the host workstation of the user in the Node name field. Type the user name in the User name field.
4. Click the **Set** button.

Results

If you are using commands, use the `fromnode` and `fromowner` options to indicate the node name and the name of the user who owns the files.

For example, to restore files to one of your own file systems that were backed up from a workstation named `Node1` and owned by a user named `Ann`, enter:

```
dsmc restore -fromn=node1 -fromo=ann "/home/proj/*" /home/gillis/
```

Use the **query filespace** command to get a list of file spaces. For example, to get a list of file spaces owned by `Ann` on `Node1`, enter:

```
dsmc query filespace -fromn=node1 -fromo=ann
```

Related reference:

“**Fromnode**” on page 410

“**Query Filespace**” on page 674

“**Restore**” on page 693

“**Retrieve**” on page 720

Restore or retrieve files to another workstation

From a different workstation, you can restore or retrieve files you have already backed up from your own workstation. You must know the Tivoli Storage Manager password assigned to your node.

To restore or retrieve files to another workstation, use the `virtualnodename` option to specify the node name of the workstation from which you backed up the files. The `virtualnodename` option cannot be set to the hostname of the workstation. You can use the `virtualnodename` option when you start Tivoli Storage Manager or you can add the `virtualnodename` option to your client user options file `dsm.opt`. Use the `virtualnodename` option on the **dsmj** command if you are borrowing the workstation of another user and you do not want to update their client user-options file.

Tivoli Storage Manager prompts you for the password for your original node. After you enter the correct password, all file systems from your original workstation appear in the Restore or Retrieve window. You can restore or retrieve files as if you were working on your own workstation.

Important: When you use this method to access files, you have access to all files backed up and archived from your workstation. You are considered a virtual root user.

You can use the `virtualnodename` option in a command. For example, to restore your *projx* files, enter:

```
dsmc restore -virtualnodename=nodeone "/home/monnett/projx/*"
```

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

The considerations for retrieving files are the same as restoring files.

Restoring a disk in case of disk loss

You can only recover your files if you can run the client. If the disk that contains the client is lost (from theft or hardware failure, for example), you must reinstall

the client before you can recover your files. If you also lose the disk that contains the operating system and communication software, you must recover them before you can connect to the Tivoli Storage Manager server.

About this task

To protect yourself against these kinds of losses, you need to put together a set of installation media that you can use to restore your system to a state that lets you contact the server and begin recovering data. The installation media should contain:

Procedure

1. A startable operating system that lets you perform basic functions.
2. A correctly configured communication program that lets you establish communications with the server.
3. A client with appropriate customized options files. You can use the client command line interface to complete this task.

Results

The communication package you use determines what files you need. Consult your operating system and communication software manuals to set up your installation media.

If you also have the Tivoli Storage Manager for Space Management installed on your workstation, your installation media should include the HSM command line client.

Note: Your administrator can schedule restore operations, which can be very useful when you need to restore a large number of files.

Related concepts:

 Backup and restore on space managed file systems

Deleting file spaces

If your Tivoli Storage Manager administrator gives you authority, you can delete entire file spaces from the server.

About this task

When you delete a file space, you delete all the files and images, both backup versions and archive copies, that are contained within the file space. For example, if you delete the file space for your `/home/monnet` file system, you are deleting every backup for every file in that file system and every file you archived from that file system. **Carefully consider whether you want to delete a file space.** You must be an authorized user to perform this task.

You can delete individual backup versions by using the **delete backup** command.

You can delete file spaces using the backup-archive client GUI or client command line interface. To delete NAS file spaces, use the web client or client command line interface.

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

Procedure

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

Results

You can also delete a file space using the **delete filesystem** command. Use the *class* option with the **delete filesystem** command to delete NAS file spaces.

Related reference:

"Class" on page 333

"Delete Backup" on page 645

"Delete Filespace" on page 648

Enable SELinux to restore files on the Red Hat Enterprise Linux 5 client

If you are a non-root user, and you are trying to restore files on the Red Hat Enterprise Linux 5 client, you must first enable SELinux.

If you do not enable SELinux, you will have problems if you restore files that have modified extended attributes.

Chapter 6. Archive and retrieve your data (UNIX and Linux)

You can archive infrequently used files to the Tivoli Storage Manager server and retrieve them when necessary. Archiving and retrieving files is similar to backing up and restoring files. Many of the windows and concepts are similar.

All the primary archive and retrieve procedures in this topic also apply to the web client, except for the Preferences editor procedures.

You can complete the following primary archive and retrieve tasks:

- “Archiving data with the GUI”
- “Archive data examples by using the command line” on page 250
- “Deleting archive data” on page 253
- “Retrieving data with the GUI” on page 256
- “Retrieve data examples by using the command line” on page 256

Related concepts:

Chapter 4, “Backing up your data,” on page 147

Related tasks:

“Starting a web client session” on page 140

Archive files

To archive files, you must specifically select the files to archive. You can select the files by using a file specification or by selecting them from a directory tree.

Your administrator might set up schedules to archive certain files on your workstation automatically. The following sections cover how to archive files without using a schedule.

Related tasks:

“Setting the client scheduler process to run as a background task and start automatically at startup” on page 264

Archiving data with the GUI

You can archive a file or a group of files by using file names. You can select files that match your search criteria by using a directory tree.

Procedure

Archive files with the following procedure.

1. Click **Archive** from the main window.
2. In the Archive window, 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 toolbar.
3. Enter your search criteria in the Find Files window.
4. Click **Search**.
5. In the Matching Files window, click the selection boxes next to the files you want to archive and close the Matching Files window.
6. Enter your filter criteria in the Find Files window.

7. Click **Filter**. The Archive window displays the filtered files.
8. Click the selection boxes next to the filtered files or directories that you want to archive.
9. Enter the description, accept the default description, or select an existing description for your archive package in the **Description** box. The maximum length of a description is 254 characters. When an existing archive description is used, the files or directories that are selected are added to the archive package. All archived packages with the same description are grouped for retrieves, queries, and deletions.
10. To modify specific archive options, click **Options**. Any options that you change are effective during the current session only.
11. Click **Archive**. The archive Task List window displays the archive processing status.

Archive data examples by 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. Examples of archiving data by using the command line are shown.

You can archive a single file, a group of files, or all the files in a directory or subdirectory. After you archive a file, you can choose to delete the original file from your workstation.

The following table shows examples of using the **archive** command to archive objects.

Table 47. Command line archive examples

Task	Command	Considerations
Archive all files in the /home/proj1 directory with a file extension of .txt.	<code>dsmc archive "/home/proj1/*.txt"</code>	Use wildcards to archive more than one file at a time.
Archive all files in the /home/jones/proj/ directory and delete the files on your workstation.	<code>dsmc archive /home/jones/proj/ -deletefiles</code>	Retrieve the archived files to your workstation whenever you need them again. For more information about the <code>deletefiles</code> option, see "Deletefiles" on page 351.
Archive the /home/jones/h1.doc and /home/jones/test.doc files.	<code>dsmc archive /home/jones/h1.doc /home/jones/test.doc</code>	If you specify the <code>removeoperandlimit</code> option with the archive command, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits. With this option, you can specify more than 20 files on a single command. For more information about this option, see "Removeoperandlimit" on page 488.
Archive a list of files in the /home/avi/filelist.txt file.	<code>dsmc archive -filelist=/home/avi/filelist.txt</code>	Use the <code>filelist</code> option to process a list of files. For more information, see "Filelist" on page 402.

Table 47. Command line archive examples (continued)

Task	Command	Considerations
Archive the /home/jones/ch1.doc file and assign a description to the archive.	<code>dsmc archive /home/jones/ch1.doc -description="Chapter 1, first version"</code>	If you do not specify a description with the archive command, the default is Archive Date:x, where x is the current system date. For more information about the description option, see “Description” on page 351.
Archive all of the files in the /home/jones/proj/ directory and its subdirectories.	<code>dsmc archive /home/jones/proj/ -subdir=yes</code>	For more information about the <code>subdir</code> option, see “Subdir” on page 539.
Use the <code>v2archive</code> option with the archive command to archive only files in the /home/re1x/dir1 directory, but not the re1x or dir1 directories.	<code>dsmc archive "/home/re1x/dir1/" -v2archive</code>	The backup-archive client archives only files in the /home/re1x/dir1 directory. Directories that exist in the path are not processed. For more information about the <code>v2archive</code> option, see “V2archive” on page 560.
Use the <code>archmc</code> option with the archive command to specify the available management class for your policy domain to which you want to bind your archived files.	<code>dsmc archive -archmc=ret2yrs /home/plan/proj1/budget.jan</code>	For more information about the <code>archmc</code> option, see “Archmc” on page 318. For more information about management classes, see Chapter 9, “Storage management policies,” on page 275.
Assume that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1. You archive the /usr/dir1/sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name /usr.	<code>dsmc archive /usr/dir1/sub1/ -subdir=yes -snapshotroot=/snapshot/day1</code>	The client considers the <code>snapshotroot</code> value as a file space name. For more information, see “Snapshotroot” on page 528.

Related reference:

“Archive” on page 611

Associate a local snapshot with a server file space

To associate data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server, use the `snapshotroot` option.

By using the `snapshotroot` option with the **archive** command with a vendor-acquired application that provides a snapshot of a logical volume, you can associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server.

You cannot use the `snapshotroot` option to take a volume snapshot, but you can use the option to manage data that is created by a volume snapshot.

Related reference:

“Snapshotroot” on page 528

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.

Before you begin

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.

Consider the following features of a proxied session:

- A proxy operation uses the settings for the target node (such as **maxnummp** and **deduplication**) and schedules that are defined on the Tivoli Storage Manager server. The Tivoli Storage Manager server node settings and schedules for the agent node are ignored.
- You cannot perform a system state or system services backup or restore.
- You cannot access another node (either from the GUI drop down or use of the `fromnode` option).
- You cannot perform a NAS backup or restore.

About this task

Consolidating archived files to a common target node name on the server is useful when the workstation responsible for performing the archive can change over time, such as with a Xsan or cluster. The `asnodename` option also allows data to be restored from a different system than the one which performed the backup. Use the `asnodename` option with the appropriate command to back up, archive, restore, and retrieve data under the target node name on the Tivoli Storage Manager server. This support is only available with Tivoli Storage Manager V5.3 and higher.

Tivoli Storage Manager FastBack clients are also backed up using client node proxy.

Configuring your environment for proxied operations is a multiple step procedure that involves setting options and commands on the backup-archive client and on the server.

Procedure

Perform steps 1 through 5 to install the client and grant proxy authority to the nodes that can perform archive procedures on behalf of another node.

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

Step 6 sets ensures that archived files are encrypted on the server.

6. Set the `encryptkey` option in the options file.

Specify `encryptkey=save` in the options file to save the encryption key in the Tivoli Storage Manager password file. Back up at least one file with `asnode=ProxyNodeName` to create a local encryption key on each agent node in the multiple node environment.

Specify `encryptkey=prompt` in the options file if you want the node users to manage the encryption key. Ensure that users of the agent nodes in the multiple node environment are using the same encryption key.

Repeat this step if you change the encryption key. Use the same encryption key for all files that are backed up, in the shared environment.

Perform steps 7 to step 10 to enable multinode operation, from the GUI.

7. 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.
8. Select **Edit > Preferences** to open the preferences window.
9. Select the **General** tab and fill in the **As Node Name** field with the name of the proxy authorized target node.
10. Click **Apply** and then **OK** to close the preferences window.

Perform step 11 to verify that your client node is now accessing the server as the target node.

11. Open the tree window and verify that the target node name specified by the **As Node Name** field is displayed. Alternatively, you can verify that the target node name shows in the **Accessing As Node** field in the **Connection Information** window.
12. Optional: To return to single node operation, delete the **As Node Name** from the **Accessing As Node** field in the **General > Preferences** tab.

Related reference:

“Asnodename” on page 319

“Session settings and schedules for a proxy operation” on page 320

Deleting archive data

You can delete individual archive objects from the Tivoli Storage Manager server, without having to delete the entire file space to which they belong.

Before you begin

Your Tivoli Storage Manager administrator must grant you the authority to delete archived objects. To determine whether you have this authority, select **File > Connection Information** from the backup-archive client GUI or from the main menu in the web client. Your archive delete authority status is listed in the **Delete Archive Files** field. If this field shows **No**, you cannot delete archived objects unless your administrator grants you the authority to delete them.

Procedure

To delete an archived object from the server, perform the following steps in the web client or GUI. As an alternative to using the web client or GUI, you can also delete archived objects from the command line by using the **delete archive** command.

1. Select **Delete Archive Data** from the **Utilities** menu.

2. In the Archive Delete window, 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. Select the archived objects that you want to delete.
4. Click **Delete**. The client prompts you for confirmation before it starts to delete the selected objects. The Archive Delete Task List window shows the progress of the delete operation.

Related reference:

"Delete Archive" on page 643

Advanced archive tasks

Access permissions, symbolic links, and hard links are advanced functions to consider when you archive data.

Access permissions

When you archive a file, the client saves standard UNIX access permissions that are assigned to the file.

Depending on your operating system, it also saves extended permissions. For example, for files on an AIX workstation, the client saves access control lists.

If you are a user, and you archive a file to which you have read access, you own the archived copy of the file. You are the only user who can retrieve the archived file unless you grant access to another user.

Archive and retrieve symbolic links

The backup-archive client archives and retrieves symbolic links differently than it does regular files and directories.

The way that the client archives and retrieves symbolic links depends on options settings, whether the target directory is accessible, and the way you specify objects.

A *UNIX symbolic link* is a file that contains a pointer to another file or directory. The object the symbolic link points to is called the *target object*.

A symbolic link can be backed up as path information to a target directory, or it can be backed up as a directory. If the symbolic link is backed up as a directory, the files and folders in the target directory can also be backed up.

What you restore depends on how the symbolic link was backed up, the scope of the restore, the setting of the `followsymbolic` option, and whether the target directory is accessible at the time of restore.

For more information on how symbolic links are handled during archive, see the `archsymbasfile` option.

Note: Symbolic link processing as described here does not apply to Mac OS X. Symbolic links are always archived as files and are never followed.

The following table shows symbolic link archive and retrieve functions and the action taken:

Table 48. Symbolic link management table for archive and retrieve

Function	Action taken
Archive of a file link.	Archives the file to which the symbolic link points.
Archive of a directory link.	Archives the directory and its contents.
Archive of a file with subdir=yes.	Archives the file, directory path and all like-named files in the subtree.
Archive of a directory with subdir=yes.	Archives the directory, its contents, and contents of subdirectories.
Archive of a symbolic link that points to a file or directory that does not exist.	Archives the symbolic link.
Retrieve a symbolic link that points to file; the file and link exist.	Replaces the file if replace=y is set.
Retrieve a symbolic link that points to file; the symbolic link no longer exists.	Retrieves the file replacing the file name with the symbolic link name and places it in the directory where the symbolic link resided.
Retrieve a symbolic link that points to a directory; the symbolic link and directory no longer exist.	A directory is created in the directory where the symbolic link resides, and all files and subdirectories are restored to that directory. The symbolic link name is used as the new directory name.
Retrieve a symbolic link that points to a directory; the symbolic link and directory still exist.	The directory is not retrieved as long as the symbolic link exists.

Related reference:

“Archsymlinkasfile” on page 319

Hard links

When you archive files that are hard-linked, the backup-archive client archives each instance of the linked file.

For example, if you archive two files that are hard-linked, the client archives the file data twice.

When you retrieve hard-linked files, the client reestablishes the links. For example, if you had a hard-linked pair of files, and only one of the hard-linked files is on your workstation, when you retrieve both files, they are hard-linked. The only exception to this procedure occurs if you archive two files that are hard-linked and then break the connection between them on your workstation. If you retrieve the two files from the server, the client respects the current file system and does not retrieve the hard link.

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

Retrieve archives

Retrieve a file when you want to return an archive copy from the server to your workstation.

Many of the advanced considerations for retrieving files are the same as for restoring files.

Important: When you retrieve a file without any specifications, and more than one version of the archive copy exists on the server, the client retrieves all of the copies. After the first copy is retrieved, the second copy is retrieved. If there is an existing copy on your client workstation, you are prompted to replace, skip, or cancel.

Related concepts:

“Restore or retrieve files to another workstation” on page 246

Related tasks:

“Authorizing another user to restore or retrieve your files” on page 244

“Restoring or retrieving files from another client node” on page 245

Retrieving data with the GUI

You can retrieve an archived file with the GUI.

Procedure

1. Click **Retrieve** from the client Java GUI main window. The Retrieve window displays.
2. Expand the directory tree by clicking the plus sign (+) or the folder icon next to an object that you want to expand. To search or filter files, click the **Search** icon from the toolbar.
3. Enter your search criteria in the Find Files window.
4. Click **Search**. The Matching Files window displays.
5. Click the selection boxes next to the files that you want to retrieve and close the Matching Files window.
6. Enter your filter criteria in the Find Files window.
7. Click **Filter**. The Retrieve window displays the filtered files.
8. Click the selection boxes of the filtered files or directories that you want to retrieve.
9. To modify specific retrieve options, click **Options**. Any options that you change are effective during the current session only.
10. Click **Retrieve**. The Retrieve Destination window displays. Enter the appropriate information in the Retrieve Destination window.
11. Click **Retrieve**. The Task List window displays the retrieve processing status.

Retrieve data examples by using the command line

You can retrieve a single file, a group of files, or all the files in a directory or subdirectory.

When you retrieve a file, a copy of that file is sent from the Tivoli Storage Manager server. The archived file remains in storage.

Use the **retrieve** command to retrieve files from storage to your workstation. The following table shows examples of using the **retrieve** command.

Table 49. Command line examples of retrieving archives

Task	Command	Considerations
Retrieve the /home/jones/h1.doc file to its original directory.	dsmc retrieve /home/jones/h1.doc	If you do not specify a destination, the files are retrieved to their original location.
Retrieve the /home/jones/h1.doc file with a new name and directory.	dsmc retrieve /home/jones/h1.doc /home/smith/h2.doc	None.
Retrieve all files from the /home/jones directory that end with the characters .bak to the /home/smith directory.	dsmc retrieve "/home/jones/*.bak" /home/smith/	None.
Retrieve the /home/jones/ch1.doc file and assign a description.	dsmc retrieve /home/jones/ch1.doc -description="Chapter 1, first version"	If you do not specify a description with the retrieve command, the default is Retrieve Date:x, where x is the current system date.
Use the pick option to display a list of archives from which you can select files to retrieve.	dsmc retrieve "/home/jones/*" -pick	None.
Retrieve a list of files that are specified in the retrievelist.txt file to their original directory.	dsmc retrieve -filelist=/home/dir2/retrievelist.txt	None.

Related reference:

“Retrieve” on page 720

“Description” on page 351

“Filelist” on page 402

“Pick” on page 472

Archive management classes

The backup-archive client checks the include options in your include-exclude options list to determine which management class to assign to your archived files.

If you do not assign a management class to a file with the include option, the client assigns the default management class to the file. The client can archive only a file if the selected management class contains an archive copy group.

You can override the default management class by using the **archmc** option, or by clicking **Options** in the Archive window in the GUI, clicking **Override include/exclude list**, and then selecting the management class.

You can also add include-exclude statements in the backup-archive client Java GUI or web client directory tree. Then, you can use the **Utilities Preview Include-Exclude** function to preview the include-exclude list before you send data to the server.

Related concepts:

“Assign a management class to files” on page 281

“Display information about management classes and copy groups” on page 277

Related reference:

“Preview Archive” on page 663

“Preview Backup” on page 664

Chapter 7. Tivoli Storage Manager scheduler overview

The Tivoli Storage Manager central scheduler allows client operations to occur automatically at specified times.

To understand scheduling with Tivoli Storage Manager, several terms need to be defined:

schedule definition

A schedule definition on the Tivoli Storage Manager server specifies critical properties of an automated activity, including the type of action, the time the action should take place, and how frequently the action takes place. Numerous other properties can be set for a schedule. For information about the **DEFINE SCHEDULE**, see the Tivoli Storage Manager server documentation.

schedule association

A schedule association is 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 that are defined to a certain policy domain to be associated with schedules defined in that domain.

scheduled event

A scheduled event is a specific occurrence of when a schedule is run for a node. The following conditions must be met before automatic scheduled events take place for a client:

- A schedule definition must exist for a specific policy domain.
- A schedule association must exist for the required node, which belongs to that policy domain.
- The client scheduler process must be running on the client system.

When creating a schedule definition on the Tivoli Storage Manager server, schedule actions that you can take include incremental, selective, archive, restore, retrieve, image backup (does not apply to Mac OS X), image restore (does not apply to Mac OS X), command, and macro. The scheduled action that is most frequently used is incremental with the **objects** parameter left undefined. With this setting, the backup-archive client performs a domain incremental backup of all file systems defined by the client domain option. A schedule definition using the **command** action allows an operating system command or shell script to be executed. When automating tasks for Tivoli Storage Manager for Data Protection clients, you must use **command** action schedule definitions, which invoke the command-line utilities for those applications.

The schedule *startup window* indicates the acceptable time period for a scheduled event to start. The startup window is defined by these schedule definition parameters: **startdate**, **starttime**, **durunits**, and **duration**. The **startdate** and **starttime** options define the beginning of the startup window for the very first scheduled event. The beginning of the startup windows for subsequent scheduled events vary depending on the **period** and **perunit** values of the schedule definition. The **duration** and **durunits** parameters define the length of the startup window. The schedule action is required to start within the startup window. To illustrate, consider the results of the following schedule definition:

```
define schedule standard test1 action=incremental starttime=12:00:00 period=1
perunits=hour dur=30 duru=minutes
```

Event	Window start	Window end	Actual start (just an example, times vary)
1	12:00:00	12:30:00	12:05:33
2	13:00:00	13:30:00	13:15:02
3	14:00:00	14:30:00	14:02:00
and so on			

The variation in actual start times is a result of the randomization feature provided by the Tivoli Storage Manager central scheduler which helps to balance the load of scheduled sessions on the Tivoli Storage Manager server.

Examples: Blank spaces in file names in schedule definitions

When you define or update a schedule **objects** parameter or the schedule **options** parameter with file specifications that contain blank spaces, put quotation marks (") around each file specification that contains blanks, then add single quotes (') around the entire specification.

The following examples show how to delimit schedule **object** parameters when file specifications contain space characters:

```
objects='/home/proj1/Some file.doc'
objects='/home/proj1/Some file.doc' "/home/Another file.txt" /home/noblanks.txt'
objects='/home/My Directory With Blank Spaces/'
objects='/Users/user1/Documents/Some file.doc'
objects='/Users/user1/Documents/Some file.doc'
"/Users/user5/Documents/Another file.txt" /Users/user3/Documents/noblanks.txt'
objects='/Users/user1/My Directory With Blank Spaces/'
```

This syntax ensures that a file specification containing a space, such as /home/proj1/Some file.doc, is treated as a single file name, and not as two separate files (/home/proj1/Some, and file.doc).

The following examples show how to delimit schedule **options** parameters when file specifications contain space characters:

```
options='-preschedulecmd="/home/me/my files/bin/myscript"
-postschedulecmd="/home/me/my files/bin/mypostsript" -quiet'
options='-presched="/home/me/my files/bin/precmd" -postsched=finish'
```

You can also refer to the **objects** and **options** parameter information for the **DEFINE SCHEDULE** and **UPDATE SCHEDULE** commands. For descriptions of these commands and parameters, see the Tivoli Storage Manager server documentation..

Related concepts:

"Specifying input strings that contain blank spaces or quotation marks" on page 138

Preferential start times for certain nodes

Occasionally, you might want to ensure that a particular node begins its scheduled activity as close as possible to the defined start time of the schedule. The need for this typically arises when prompted mode scheduling is in use.

Depending on the number of client nodes associated with the schedule and where the node is in the prompting sequence, the node might be prompted significantly later than the start time for the schedule.

In this case, you can perform the following steps:

1. Copy the schedule to a new schedule with a different name (or define a new schedule with the preferred attributes).
2. Set the new schedule priority attribute so that it has a higher priority than the original schedule.
3. Delete the association for the node from the original schedule, then associate the node to the new schedule.

Now the Tivoli Storage Manager server processes the new schedule first.

Scheduler processing options

Scheduler processing options determine what operations are performed when a scheduler job is started.

You can define most of these scheduler processing options in the client options file. However, some of these options can be set on the Tivoli Storage Manager server, so they affect all clients.

The following table shows which options are defined by the client and server, and which options are overridden by the server. An X in a column indicates where the option can be specified.

Option	Client defined	Server defined	Server global override
manageservices	X		
maxcmdretries	X		SET MAXCMDRETRIES command
maxschedsessions		X	
postschedulecmd, postnschedulecmd	X		
preschedulecmd, prenschedulecmd	X		
queryschedperiod	X		SET QUERYSCHEDPERIOD command
randomize		X	
retryperiod	X		SET RETRYPERIOD command
schedcmddisabled	X		
schedlogname	X		
schedlogretention	X		
schedmode	X		SET SCHEDMODES command
sessioninitiation	X	X	UPDATE NODE command

Option	Client defined	Server defined	Server global override
tcpclientaddress	X	X (also defined on server when sessioninit=serveronly as part of the node definition)	
tcpclientport	X	X (also defined on server when sessioninit=serveronly as part of the node definition)	

Client defined options are defined in the `dsm.sys` or `dsm.opt` file, depending on the option and platform. The Tivoli Storage Manager server can also define some options in a client options set, or as part of the options parameter of the schedule definition. The Tivoli Storage Manager server can also set some options globally for all clients. By default, the client setting for these options is honored. If the global override on the Tivoli Storage Manager server is set, the client setting for the option is ignored. Defining client options as part of the schedule definition is useful if you want to use specific options for a scheduled action that differ from the option settings normally used by the client node, or are different for each schedule the node executes.

The `schedmode` option controls the communication interaction between the Tivoli Storage Manager client and server. There are two variations on the schedule mode: *client polling* and *server prompted*. These variations are explained in the Tivoli Storage Manager server documentation.

Evaluate schedule return codes in schedule scripts

You can use environment variables to determine the current Tivoli Storage Manager return code before you run a script by using either the `preschedulecmd` or `postschedulecmd` client options.

Tivoli Storage Manager provides the current value of the return code in the environment variable called `TSM_PRE_CMD_RC`. The `TSM_PRE_CMD_RC` variable is the current value of the Tivoli Storage Manager return code before you run a schedule script. The value of the `TSM_PRE_CMD_RC` variable is not necessarily the same as the return code issued by Tivoli Storage Manager following the execution of the schedule script. The `TSM_PRE_CMD_RC` variable can be used in schedule scripts to determine the current state of the schedule.

The `TSM_PRE_CMD_RC` variable is set on each of the following schedule options: `preschedule`, `prenschedule`, `postschedule`, and `postnschedule`. `TSM_PRE_CMD_RC` affects those schedules that have the `ACTION=COMMAND` option specified.

An example of the `TSM_PRE_CMD_RC` variable in use:

```
if [[ -n ${TSM_PRE_CMD_RC} ]] ; then
    if [[ ${TSM_PRE_CMD_RC} == 0 ]] ; then
        echo "The TSM_PRE_CMD_RC is 0"
    elif [[ ${TSM_PRE_CMD_RC} == 4 ]] ; then
        echo "The TSM_PRE_CMD_RC is 4"
```

```

elif [[ ${TSM_PRE_CMD_RC} == 8 ]] ; then
    echo "The TSM_PRE_CMD_RC is 8"

elif [[ ${TSM_PRE_CMD_RC} == 12 ]] ; then
    echo "The TSM_PRE_CMD_RC is 12"
else
    echo "The TSM_PRE_CMD_RC is an unexpected value: ${TSM_PRE_CMD_RC}"
fi

else
    echo "The TSM_PRE_CMD_RC is not set"
fi

```

Return codes from preschedulecmd and postschedulecmd scripts

The return codes that you might see when you use the preschedulecmd and postschedulecmd options are described.

- If the command specified by the preschedulecmd option ends with a nonzero return code, Tivoli Storage Manager assumes that the command failed. In this case, the scheduled event and any postschedulecmd or postnschedulecmd command cannot run. The administrative **query event** command with format=detailed option shows that the event failed with return code 12.
- If the command specified by the postschedulecmd option ends with a nonzero return code, Tivoli Storage Manager considers the command to be failed. The administrative **query event** command with format=detailed option shows that the event completed with return code 8. The exception is if the scheduled operation completed with a higher return code, in which case the higher return code takes precedence. Therefore, if the scheduled operation completes with return code 0 or 4 and the postschedulecmd command fails, the administrative **query event** command shows that the event completed with return code 8. If the scheduled operation completes with return code 12, that return code takes precedence, and **query event** shows that the event failed with return code 12.

When you interpret the return code from a command, Tivoli Storage Manager considers 0 to mean success, and anything else to mean failure. While this behavior is widely accepted in the industry, it is not 100% guaranteed. For example, the developer of the widget command might exit with return code 3, if widget ran successfully. Therefore, it is possible that the preschedulecmd or postschedulecmd command might end with a nonzero return code and still be successful. To prevent Tivoli Storage Manager from treating such commands as failed, you can wrap these commands in a script, and code the script so that it interprets the command return codes correctly. The script exits with return code 0 if the command was successful; otherwise it exits with a nonzero return code. The logic for a script running widget might look like this example:

```

run 'widget'
if lastcc == 3
    exit 0
else
    exit 1

```

Related reference:

“Postschedulecmd/Postnschedulecmd” on page 475

“Preschedulecmd/Prenschedulecmd” on page 478

Client-acceptor scheduler services versus the traditional scheduler services

You can configure the Tivoli Storage Manager client to manage the scheduler process using the Tivoli Storage Manager client acceptor daemon.

The client acceptor daemon provides a light-weight timer which automatically starts and stops the scheduler process as needed. Alternatively, the traditional method keeps the Tivoli Storage Manager scheduler process running continuously. Generally, using the client acceptor daemon to manage the scheduler is the preferred method.

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

Client acceptor daemon-managed services

- Defined using the `manageservices 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; if you delete an option from a client options set, you must restart the scheduler so the scheduler is made aware of the deletion.

Tip: 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 startup

You can configure the Tivoli Storage Manager client scheduler to run as a background system task that starts automatically when your system is started.

About this task

You can complete this task whether you use the client acceptor to manage the scheduler or whether you use the traditional method to start the scheduler client scheduler.

When you are running a client acceptor-managed schedule, set the client acceptor process to start automatically at startup time; not the scheduler process. For the traditional method, set the scheduler process to start automatically at startup time.

You can configure the client acceptor to run as a background system task that starts automatically when your system is started. To configure the client acceptor to manage scheduled backups, you use the `manageservices` option to specify whether the client acceptor manages only the scheduler, only the web client, or both the scheduler and web client. The method for setting up the client acceptor as a system task varies for each platform.


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 client command such as `dsmc query session`. For testing purposes, you can always start the scheduler in the foreground by running `dsmc sched` from a command prompt (without a `manageservices` stanza set).

To start the scheduler automatically at startup time, use either the client acceptor-managed method or the traditional method.

Client acceptor-managed method

1. In your `dsm.sys` file, set the `manageservices` option to **schedule** or **schedule webclient**.
2. Start the client acceptor.
 - a. On AIX, HP-UX, and Solaris clients, add the following entry into the system startup file (`/etc/inittab` for most platforms):


```
tsm::once:/usr/bin/dsmcad > /dev/null 2>&1 # TSM Client
Acceptor Daemon
```
 - b. On Linux clients, the installation program creates a startup script for the client acceptor (`dsmcad`) in `/etc/init.d`. The client acceptor (`dsmcad`) must be started before it can manage scheduler tasks, or manage the web client. As root, use the following command to start, stop, or restart the client acceptor, or check its status:



To enable the client acceptor to start automatically after a system restart, add the service as follows, at a shell prompt:

```
# chkconfig --add dsmcad
```

If the Linux operating system runs the `systemd` initialization service, complete the following steps to start the `dsmcad` and to run it at system start time:

- 1) Copy the provided `systemd` unit file `/opt/tivoli/tsm/client/ba/bin/dsmcad.service` to the `/etc/systemd/system/` directory.
- 2) Run the following command to refresh the `systemd` unit list:


```
systemctl daemon-reload
```
- 3) Run the following command to start the client acceptor at system start time:


```
systemctl enable dsmcad.service
```
- 4) Run the following command to start the client acceptor:


```
systemctl start dsmcad.service
```
- c. On Mac OS X, the client acceptor must be installed as a Startup Item. A system administrator must use the Tivoli Storage Manager

Tools for Administrators to install and start the client acceptor. To start, stop, or restart the client acceptor, use the following command:

```
▶ sudo /sbin/SystemStarter start dsmcad
                        stop
                        restart
```

3. In your `dsm.sys` file, set the `passwordaccess` option to **generate**.
4. Run a command like `dsmc query sess` to store the node password.

Traditional method:

1. Set the `managedservices` option.
 - On AIX, HP-UX, Linux, and Solaris clients, either remove the option entirely (it defaults to **webclient**) or set it to **webclient**.
 - On Mac OS X clients, set the `managedservices` option to either **webclient** or **none**. Do not set the option to `schedule`.
2. On AIX, HP-UX, Linux, and Solaris, add the following entry into the system startup file, for example, `/etc/inittab`, where it is supported:
`tsmsched::once:/usr/bin/dsmc sched > /dev/null 2>&1 # TSM scheduler`
3. In your `dsm.sys` file, set the `passwordaccess` option to **generate**.
4. Run a command like `dsmc query sess` to store the node password.
5. To start the client scheduler on your client node and connect to the server schedule, enter the following command:

```
dsmc schedule
```

If the current directory is not in your `PATH` environment variable, enter the following command:

```
./dsmc schedule
```

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

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

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

Related reference:

“Cadlistenonport” on page 332

Examples: Display information about scheduled work

Schedules can be classic or enhanced, depending on how the interval to the next execution is defined.

Classic schedules allow the period to be as small as an hour. Enhanced schedules allow actions to be executed on specific days.

To view schedules that are defined for your client node, enter:

```
dsmc query schedule
```

The backup-archive client displays detailed information about all scheduled work for your client node. Table 50 on page 267 displays sample classic **query schedule** output.

Table 50. Sample classic query schedule output

Schedule Name:	DAILY_INC
Description:	Daily System-wide backup
Schedule Style:	Classic
Action:	Incremental
Options:	QUIET
Objects:	
Priority:	1
Next Execution:	30 minutes
Duration:	4 Hours
Period:	1 Day
Day of Week:	Any
Month:	
Day of Month:	
Week of Month:	
Expire:	Never
Schedule Name:	WEEKLY_INC
Description:	Weekly backup for project files
Schedule Style:	Classic
Action:	Incremental
Options:	QUIET
Objects:	/proj
Priority:	1
Next Execution:	60 minutes
Duration:	8 Hours
Period:	7 Days
Day of Week:	Friday
Month:	
Day of Month:	
Week of Month:	
Expire:	Never

The schedule name, **WEEKLY_INC**, starts a weekly incremental backup in the /proj file system.

The schedule name, **DAILY_INC**, starts a daily incremental backup. The next incremental backup starts in 30 minutes. Because no objects are listed, the client runs the incremental backup on your default domain. The schedule has no expiration date.

To more accurately determine the status of scheduled events, the **query schedule** output for an enhanced schedule, on Tivoli Storage Manager Version 5.3 client and above, includes new fields. These fields are always displayed, even if it is a classic schedule or a Tivoli Storage Manager Version 5.3 client session with a pre-Version 5.3 server, but the new fields are blank. Note that for a down-level (prior to Tivoli Storage Manager Version 5.3) client, the server reports the period as indefinite and the day of week as an illegal day. Table 51 on page 268 displays sample enhanced **query schedule** output.

Table 51. Sample enhanced query schedule output

Schedule Name: QUARTERLY_FULL
Description: Quarterly full backup
Schedule Style: Enhanced
Action: Selective
Options: subdir=yes
Objects: /* /Volumes/fs2/*
Priority: 5
Next Execution: 1744 Hours and 26 Minutes
Duration: 1 Day
Period:
Day of Week: Friday
Month: March, June, September, December
Day of Month: Any
Week of Month: Last
Expire: Never

Display information about completed work

When you run the **schedule** command in the foreground, your screen displays output from the scheduled commands.

Output is also directed to the `dsmsched.log` file in the installation directory unless you change the directory and file name using the `schedlogname` option.

When you run the **schedule** command in the background, output from scheduled commands is directed to the `dsmsched.log` file in the current directory, or to the path and file name that you specified. The `dsmsched.log` cannot be a symbolic link.

Note: On Mac OS X, by default the log can be found in one of these locations:

```
~/Library/Logs/tivoli/tsm
/Library/Logs/tivoli/tsm
```

After scheduled work is performed, check the schedule log to verify that all work completed successfully.

When a scheduled command is processed the schedule log contains the following entry:

```
Scheduled event eventname completed successfully
```

If the scheduled event does not complete successfully, you receive a message similar to the following:

```
ANS1512E Scheduled event eventname failed. Return code = code.
```

The client indicates whether Tivoli Storage Manager successfully issued the scheduled command associated with the *eventname* (action=command). No attempt is made to determine the success or failure of the command. You can assess the status of the command by evaluating the return code from the scheduled command in the schedule log. The schedule log entry for the return code of the command is prefaced with the following text:

```
Finished command. Return code is:
```

The schedule log continues to grow unless you prune it using the `schedlogretention` option or specify a maximum size using the `schedlogmax` option.

Related concepts:

Specify scheduling options

You can modify scheduling options in the client options file or the graphical user interface (GUI).

However, if your administrator specifies a value for these options, that value overrides the value in your client.

Related concepts:

“Scheduling options” on page 302

Scheduler options for commands

The scheduler executes commands under a user ID of 0 (root); however, some commands might need to be executed under a user ID other than 0.

In this case, your Tivoli Storage Manager administrator can define schedules for commands that are executed under a user ID different from the scheduler user ID using the `schedcmduser` server option.

The `schedcmduser` option specifies the name of a valid user on the system where a scheduled command is executed. This option can only be defined by the Tivoli Storage Manager server administrator. If this option is specified, the command is executed with the authorization of the specified user. Otherwise, it is executed with the scheduler authorization.

►►—SCHEDCMDUser—*user_name*—◄◄

user_name

Specifies the name of a valid user on the system where a scheduled command is executed.

Note: The `schedcmduser` option does *not* affect the user ID used for the pre-schedule and post-schedule commands. Pre-schedule and post-schedule always run as root (user ID 0).

Enable or disable scheduled commands

You can use the `schedcmddisabled` option to disable the scheduling of commands by the server.

Commands are scheduled by using the `action=command` option on the DEFINE SCHEDULE server command.

The `schedcmddisabled` option does not disable the `preschedulecmd` and `postschedulecmd` commands. However, you can specify `preschedulecmd` or `postschedulecmd` with a blank or a null string to disable the scheduling of these commands.

You can use the `schedrestretrdisabled` option to prevent the Tivoli Storage Manager Server administrator from executing restore or retrieve schedule operations.

You can use the `srvprepostscheddisabled` option to prevent the Tivoli Storage Manager Server administrator from executing pre-schedule and post-schedule commands when performing scheduled operations.

You can use the `srvprepostsnapdisabled` option to prevent the Tivoli Storage Manager Server administrator from executing pre-snapshot and post-snapshot commands when performing scheduled image snapshot backup operations.

Related reference:

“`Schedcmddisabled`” on page 503

“`Schedrestretrdisabled`” on page 510

“`Srvprepostscheddisabled`” on page 531

“`Srvprepostsnapdisabled`” on page 532

Manage multiple schedule requirements on one system

In certain situations it is preferable to have more than one scheduled activity for each client system.

About this task

Normally, you can do this by associating a node with more than one schedule definition. This is the standard method of running multiple schedules on one system.

You must ensure that the schedule windows for each schedule do not overlap. A single client scheduler process is not capable of executing multiple scheduled actions simultaneously, so if there is overlap, the second schedule to start is missed if the first schedule does not complete before the end of the startup window of the second schedule.

Suppose that most of the file systems on your client system must be backed up daily, and that one file system containing critical data must be backed up hourly. In this case, you would need to define two schedules to handle this requirement. To avoid conflict between the hourly and daily backup schedule, the *starttime* of each schedule needs to be varied.

In certain cases, it is necessary to run more than one scheduler process on a system. Multiple processes require a separate options file for each process and must contain the following information:

- Define a unique node name for each process
- Specify unique schedule and error logs for each process
- When running in prompted mode, you must use the `tcpclientport` option to specify a unique port for each process.

The advantages of using multiple schedule processes:

- You can run more than one scheduled backup at the same time.
- You can specify different backup criteria for each schedule started, with the client option file or Tivoli Storage Manager server override options.

The disadvantages of using multiple schedule processes:

- A unique file space for each node name on the Tivoli Storage Manager server is created.

- When restoring the data, you must use the same node name associated with the backup.

Multiple schedule processes can run on UNIX and Linux platforms with either the client acceptor daemon-managed method, or the traditional method of running the scheduler. In either case, there are certain setup requirements:

- Each process must run using a different node name.
- You must create multiple stanzas in the dsm.sys file for each scheduler process. In each stanza, you must define a unique node name, along with unique values for the options `errorlogname` and `schedlogname`. You might also choose to define customized domain, include, and exclude statements for each stanza.
- In your dsm.sys file, set the `passwordaccess` option to `generate` in each stanza. The password must be generated for each node name that is running a scheduler process, by running a command such as `dsmc query sess`.
- If running with the `schedmode` option set to *prompt*, you should set a unique `tcpclientport` value for each stanza.

You must start each `dsmc sched` command or instance with the `-servername` option to reference its unique stanza name in dsm.sys. For `dsmcad`, it is necessary to define the environment variable `DSM_CONFIG` for each instance of `dsmcad` to reference its unique option file.

The following is an example configuration of two schedule processes managed by the client acceptor daemon in the dsm.sys file. Note that you must use full paths for the log file names to avoid the files being written in the root directory):

```
servername tsm1_sched1
nodename      aixsvt01_sched1
tcpserv       firebat
tcpclientport 1507
passwordaccess generate
domain        /svt1
schedmode     prompted
schedlogname  /tsm/dsmsched1.log
errorlogname  /tsm/dsmerror1.log
managedservices schedule

servername tsm1_sched2
nodename      aixsvt01_sched2
tcpserv       firebat
tcpclientport 1508
passwordaccess generate
domain        /svt1
schedmode     prompted
schedlogname  /tsm/dsmsched2.log
errorlogname  /tsm/dsmerror2.log
managedservices schedule
```

Contents of /test/dsm.opt1:

```
servername tsm1_sched1
```

Contents of /test/dsm.opt2:

```
servername tsm1_sched2
```

Open two shell command windows:

- In shell command window 1, enter:

```
export DSM_CONFIG=/test/dsm.opt1
sudo dsmcad
```
- In shell command window 2, enter:

```
export DSM_CONFIG=/test/dsm.opt2  
sudo dsmcad
```

Note: You should enter these commands into a shell script if you intend to have the dsmcad processes started directly from /etc/inittab so that the proper DSM_CONFIG variable can be set prior to launching dsmcad.

Chapter 8. Client return codes

The backup-archive command-line interface and the scheduler exit with return codes that accurately reflect the success or failure of the client operation.

Scripts, batch files, and other automation facilities can use the return code from the command-line interface. For operations that use the Tivoli Storage Manager scheduler, the return codes are shown in the output of the **QUERY EVENT** administrative command.

In general, the return code is related to the highest severity message during the client operation.

- If the highest severity message is informational (ANSnnnnI), then the return code is 0.
- If the highest severity message is a warning (ANSnnnnW), then the return code is 8.
- If the highest severity message is an error (ANSnnnnE or ANSnnnnS), then the return code is 12.

An exception to these rules is made when warning or error messages indicate that individual files could not be processed. For files that cannot be processed, the return code is 4. Examine the `dsmerror.log` file to determine the cause of errors that occur during client operations. Errors that occur during scheduled events are recorded in the `dsmsched.log` file.

Table 52 describes the return codes and their meanings.

Table 52. Client return codes and their meanings

Code	Explanation
0	All operations completed successfully.
4	The operation completed successfully, but some files were not processed. There were no other errors or warnings. This return code is common. Files are not processed for various reasons; the following reasons are the most common. • The file satisfies an entry in an exclude list. Excluded files generate log entries only during selective backups. • The file was in use by another application and could not be accessed by the client. • The file changed during the operation to an extent prohibited by the copy serialization attribute. See “Copy serialization attribute” on page 279.
8	The operation completed with at least one warning message. For scheduled events, the status is <code>Completed</code> . Review the <code>dsmerror.log</code> file (and <code>dsmsched.log</code> for scheduled events) to determine what warning messages were issued and to assess their impact on the operation.
12	The operation completed with at least one error message (except for error messages for skipped files). For scheduled events, the status is <code>Failed</code> . Review the <code>dsmerror.log</code> file (and <code>dsmsched.log</code> for scheduled events) to determine what error messages were issued and to assess their impact on the operation. Generally, this return code means that the error was severe enough to prevent the successful completion of the operation. For example, an error that prevents an entire file system or file specification from being processed yields return code 12.

Table 52. Client return codes and their meanings (continued)

Code	Explanation
<i>other</i>	For scheduled operations where the scheduled action is COMMAND, the return code is the return code from the command that was run. If the return code is 0, the status of the scheduled operation is Completed. If the return code is nonzero, then the status is Failed. Some commands might issue a nonzero return code to indicate success. For these commands, you can avoid a Failed status by wrapping the command in a script that starts the command, interprets the results, and exits. The script should produce return code 0 if the command was successful, or a nonzero return code if the command failed. Then, ask your Tivoli Storage Manager server administrator to modify the schedule definition to run your script instead of the command.

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

```
selective "/home/devel/*" -subdir=yes
incremental "/home/devel/TestDriver/*" -subdir=yes
archive "/home/plan/proj1/*" -subdir=yes
```

If the first command completes with return code 0, and the second command completes with return code 8, and the third command completed with return code 4, the return code for the macro is 8.

For more information about the **QUERY EVENT** command, see the Tivoli Storage Manager server documentation.

Related concepts:

“Scheduler options for commands” on page 269

Chapter 9. Storage management policies

Storage management policies are rules your administrator defines in order to manage your backups and archives on the server.

Your data is associated (or bound) to these policies; then when the data is backed up or archived, it is managed according to policy criteria. Policy criteria include a policy domain, a policy set, a management class, and a copy group.

Policies determine:

- Whether a file is eligible for backup or archive services.
- How many backup versions to keep.
- How long to keep inactive backup versions and archive copies.
- Where to place the copies in storage.
- For incremental backup, policies also determine:
 - How frequently a file can be backed up.
 - Whether a file must change before it is backed up again.

If you have the Tivoli HSM client installed, your administrator also defines rules that determine whether files are eligible for migration from your local file systems to storage.

This topic explains:

- Policy criteria (policy domains, policy sets, copy groups, and management classes).
- How to display policies.
- How your data is associated with policies.

Policy domains and policy sets

A *policy domain* is a group of clients with similar requirements for backing up and archiving data.

Policy domains contain one or more policy sets. An administrator uses policy domains to manage a group of client nodes in a logical way.

For example, a policy domain might include:

- A department, such as Accounting.
- A physical location, such as a particular building or floor.
- A local area network, such as all clients associated with a particular file server.

Tivoli Storage Manager includes a default policy domain named *Standard*. At first, your client node might be associated with the default policy domain. However, your administrator can define additional policy domains if there are groups of users with unique backup and archive requirements.

A *policy set* is a group of one or more management classes. Each policy domain can hold many policy sets. The administrator uses a policy set to implement different management classes based on business and user needs. Only one of these policy sets can be active at a time. This is called the *active policy set*. Each policy set contains a *default management class* and any number of additional management classes.

Management classes and copy groups

A *management class* is a collection of backup and archive copy groups that establishes and contains specific storage management requirements for backing up and archiving data.

An administrator can establish separate management classes to meet the backup and archive requirements for different kinds of data, such as:

- System data that is critical for the business.
- Application data that changes frequently.
- Report data that Management reviews monthly.
- Legal information that must be retained indefinitely, requiring a large amount of disk space.

Note: If you have the Tivoli Storage Manager for Space Management installed, it can also contain specific requirements for migrating files to storage.

Most of the work you do with storage management policies is with management classes. Each file and directory that you back up, and each file that you archive, is associated with (or *bound* to) a management class, as follows:

- If your data is not associated with a management class, Tivoli Storage Manager uses the default management class in the active policy set.
- When backing up directories, you can specify a management class with an *include* statement or the *di rmc* option. If you do not specify a management class, Tivoli Storage Manager uses the management class in the active policy set specifying the longest "Retain Only" retention period. If there are multiple management classes that meet this criteria, Tivoli Storage Manager uses the last one found, in alphabetical order.
- For archiving directories, you can specify a management class with an *include.archive* statement or the *archmc* option. If you do not specify a management class, the server assigns the default management class to the archived directory. If the default management class has no archive copy group, the server assigns the management class that currently has the archive copy group with the shortest retention time.

You can use *include* statements in your include-exclude list to associate files with management classes. In your client options file, you can associate directories with a management class, using the *di rmc* option.

Within a management class, the specific backup and archive requirements are in *copy groups*. Copy groups define the specific storage management attributes that describe how the server manages backed up or archived data. Copy groups include both *backup copy groups* and *archive copy groups*. A management class can have one backup copy group, one archive copy group, both, or neither.

A *backup copy group* contains attributes that are used during the backup process to determine:

- How many days must elapse before a file is backed up again.
- How a file is processed during a backup if it is in use.

It also contains attributes to manage the backup versions of your files on the server. These attributes control:

- On which media type the server stores backup versions of your files and directories.
- How many backup versions the server keeps of your files and directories.

- How long the server keeps backup versions of your files and directories.
- How long the server keeps inactive backup versions.
- How long the last remaining inactive version of a file is kept.

An *archive copy group* contains attributes that control:

- Whether a file is archived if it is in use
- On which media type the server stores archived copies of your files
- How long the server keeps archived copies of your files

Related concepts:

“Select a management class for files” on page 280

“Retention grace period” on page 284

Display information about management classes and copy groups

You can display policy information with the command-line interface or with a graphical user interface.

On a graphical user interface, click **View policy information** from the Utilities menu. The **Policy information** window displays the available management classes. On a command line, use the **query mgmtclass** command to view the available management classes. The **detail** option provides more information.

Table 53 shows the default values for the backup and archive copy groups in the standard management class.

Table 53. Default attribute values in the standard management class

Attribute	Backup default	Archive default
Copy group name	Standard	Standard
Copy type	Backup	Archive
Copy frequency	0 days	CMD (Command)
Versions data exists	Two versions	Does not apply
Versions data deleted	One version	Does not apply
Retain extra versions	30 days	Does not apply
Retain only version	60 days	Does not apply
Copy serialization	Shared static	Shared static
Copy mode	Modified	Absolute
Copy destination	Backuppool	Archivepool
Retain versions	Does not apply	365 days
Lan free	Destination	No
Deduplication enabled	No	No

Copy group name attribute

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

Copy type attribute

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

Copy frequency attribute

The *copy frequency* attribute is the minimum number of days that must elapse between successive incremental backups. Use this attribute during a full incremental backup.

Copy frequency works with the **mode** parameter. For example, if frequency=0 and mode=modified, a file or directory is backed up only if it changed since the last incremental backup. If frequency=0 and mode=absolute, an object is backed up every time you run an incremental backup against it. If frequency=0 and mode=absolute, changes and number of days since the last backup do not affect the current backup operation. The frequency attribute is not checked for selective backups.

For archive copy groups, copy frequency is always CMD (command). There is no restriction on how often you archive an object.

Copy frequency is ignored during a journal-based backup.

Versions data exists attribute

The *versions data exists* attribute specifies the maximum number of different backup versions retained for files and directories.

If you select a management class that permits more than one backup version, the most recent version is called the *active* version. All other versions are called *inactive* versions. If the maximum number of versions permitted is five, and you run a backup that creates a sixth version, the oldest version is deleted from server storage.

Versions data deleted attribute

The *versions data deleted* attribute specifies the maximum number of different backup versions retained for files and directories that you deleted.

This parameter is ignored until you delete the file or directory.

If you delete the file or directory, the next time you run an incremental backup, the active backup version is changed to inactive. The Tivoli Storage Manager server deletes the oldest versions in excess of the number specified by this parameter.

The expiration date for the remaining versions is based on the *retain extra versions* and *retain only version* parameters.

Retain extra versions attribute

The *retain extra versions* attribute specifies how many days all but the most recent backup version is retained.

The most recent version is the active version, and active versions are never erased. If *Nolimit* is specified, then extra versions are kept until the number of backup versions exceeds the *versions data exists* or *versions data deleted* parameter settings. In this case, the oldest extra version is deleted immediately.

Retain only version attribute

The *retain only version* attribute specifies the number of days the last remaining inactive version of a file or directory is retained.

If *Nolimit* is specified, the last version is retained indefinitely.

This parameter goes into effect during the next incremental backup after a file is deleted from the client system. Any subsequent updates to this parameter will not affect files that are already inactive. For example: If this parameter is set to 10 days when a file is inactivated during an incremental backup, the file is deleted from the server in 10 days.

Copy serialization attribute

The copy serialization attribute determines whether a file can be in use during a backup or archive, and what to do if it is.

The value for this attribute can be one of the following:

- **Static.** A file or directory must not be modified during a backup or archive. If the object is changed during a backup or archive attempt, it is not backed up or archived.
- **Shared static.** A file or directory must not be modified during backup or archive. The client attempts to perform a backup or archive as many as four additional times, depending on the value specified on the *changingretries* option in your options file. If the object is changed during every backup or archive attempt, it is not backed up or archived.
- **Dynamic.** A file or directory is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.
- **Shared dynamic.** A file or directory is backed up or archived regardless of whether it changes during a backup or archive. The client attempts to back up or archive as many as four additional times. The number of attempts depend on the value that was specified on the *changingretries* option in your options file, without the file changing during the attempt. The file is backed up or archived on the last try even if it has changed.

If you select a management class that permits a file to be backed up or archived while it is in use, the backup version or archived copy that is stored on the server might be a fuzzy copy. A *fuzzy copy* is a backup version or archived copy that does not accurately reflect what is currently in the file. It might contain some, but not all, of the changes. If that is not acceptable, select a management class that creates a backup version or archive copy only if the file does not change during a backup or archive. When you use static serialization, applications cannot open a file for write access while the file is being backed up.

If you restore or retrieve a file that contains a fuzzy copy, the file might not be usable. Do not use dynamic or shared dynamic serialization to back up files unless you are certain that a fuzzy copy that is restored is usable.

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

Related reference:

“*Snapshotproviderimage*” on page 527

Copy mode parameter

The copy **mode** parameter determines whether a file or directory is considered for incremental backup regardless of whether it changed or not since the last backup.

The client does not check the mode parameter when it runs selective backups.

The value for this parameter can be one of the following settings:

modified

The object is considered for incremental backup only if it has changed since the last backup. An object is considered changed if any of the following conditions are true:

- The date or time of the last modification is different.
- The size is different.
- If only the metadata changes (such as access permissions), the client might back up only the metadata.
- The owner is different.

absolute

The object 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 an object is archived regardless of whether it changed since the last archive request.

Related reference:

"Absolute" on page 316

Copy destination attribute

The *copy destination* attribute names the destination where backups or archives are stored.

The destination can be either a storage pool of disk devices or a storage pool of devices that support removable media, such as tape.

Retain versions attribute

The *retain versions* attribute specifies the number of days an archived file remains in storage.

When the specified number of days elapse for an archived copy of a file, it is deleted from server storage.

Deduplicate data attribute

The *deduplicate data* attribute specifies whether redundant data is transferred to the Tivoli Storage Manager server during backup and archive processing.

Related concepts:

"Client-side data deduplication" on page 79

Related reference:

"Deduplication" on page 349

"Enablededupcache" on page 378

"Exclude options" on page 389

Select a management class for files

If the default management class meets the backup and archive requirements for all the files on your workstation, it is not necessary to take any action to associate your files with that management class. This is done automatically when you back up or archive your files.

When selecting a different management class for your files, consider these questions:

- Does the management class contain a backup copy group?

If you attempt to back up a file associated with a management class that does not contain a backup copy group, the file is not backed up.

- Does the management class contain an archive copy group?

You cannot archive a file associated with a management class that does not contain an archive copy group.

- Does the backup copy group contain attributes that back up your files often enough?

Mode and frequency work together to control how often a file is backed up when you use incremental backup. These attributes are not checked for selective backup.

- What serialization method does the copy group use?

The serialization method determines how Tivoli Storage Manager functions when a file changes while it is being backed up.

- Does the backup copy group specify an adequate number of backup versions to keep, along with an adequate length of time to keep them?
- Does the archive copy group specify an adequate length of time to keep archived copies of files?

Related concepts:

“Copy serialization attribute” on page 279

Assign a management class to files

A management class defines when your files are included in a backup, how long they are kept on the server, and how many versions of the file the server should keep.

The server administrator selects a default management class. You can specify your own management class to override the default management class.

To assign a management class other than the default to directories, use the `dirmc` option in your options file.

You can assign a management class for a file or file group by using an `include` statement in your options file. You can also assign a management class by using an `include` statement in include-exclude file specified by the `inclxcl` option. Management class names are not case-sensitive.

Using the command-line client, to associate all files in the `costs` directory with the management class named `budget`, you would enter:

```
include /home/proj2/costs/* budget
```

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

```
include /* managall
```

The following examples show how to assign a management class to files:

```
exclude /*.sno
include /home/winter/*.* ice mcweekly
include /home/winter/december/*.* mcdaily
include /home/winter/january/*.* mcmonthly
include /home/winter/february/white.sno
```

Processing follows these steps:

1. The file `white.sno` in the february directory in the winter directory is backed up following bottom-up processing rules. Because you did not specify a management class on this statement, the file is assigned to the default management class.
2. Any file with an extension of `ice` in the january directory is assigned to the management class named `mcmonthly`.
3. Any file with an extension of `ice` in the december directory is assigned to the management class named `mcdaily`.
4. Any other files with an extension of `ice` in any directory under the winter directory are assigned to the management class named `mcweekly`.
5. Any file with an extension of `sno` in any directory is excluded from backup. The exception to this rule is `white.sno` in the february directory, which is in the winter directory.

To specify your own default management class `mgmt_class_name` for files that are not explicitly included, put the following statement at the top of your include list:

```
include /*.* mgmt_class_name
```

When you archive a file using the graphical user interface, you can select a different management class to override the management class assigned to the file.

Related reference:

“Dirmc” on page 356

“Include options” on page 421

Override the management class for archived files

When you archive a file, you can override the assigned management class using the a graphical user interface (GUI), or by using the `archmc` option on the **archive** command.

Overriding the management class using the GUI is equivalent to using the `archmc` option on the **archive** command. To use the GUI, press the **Options** button on the archive tree to override the management class and select a different management class.

On the command line, to associate the file `budget.jan` with the management class `ret2yrs`, enter this command:

```
dsmc archive -archmc=ret2yrs /home/jones/budget.jan
```

Select a management class for directories

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

If there is more than one management class with the longest RETONLY setting, the Tivoli Storage Manager client selects the management class whose name is last in alphabetical order.

If the default management class does not meet your requirements, select a management class with an adequate retention period specified by the `retain only version` parameter. For example, if the management class happens to back up data

directly to tape, but you want your directory backups to go to disk, you must choose a different management class. You should keep directories at least as long as you keep the files associated with those directories.

For backup directories, use the `dirmc` option to specify the management class to which directories are bound.

For archive directories, use the `archmc` option with the **archive** command.

You can use these methods to view the available management classes and their attributes:

- GUI or web client: Select **View Policy Information** from the **Utilities** menu.
- Command-line client: Run `dsmc query mgmtclass -detail`.

Note: During expiration processing on a Tivoli Storage Manager server, if an archived directory is eligible for expiration, the server checks if any existing archived files require the archived directory to remain. If so, the archived directory is not expired and the backup-archive client updates the insert date on the archived directory to ensure that the directory is not expired before the files under it.

Bind management classes to files

Binding associates a file with a management class.

When you back up a file for the first time, Tivoli Storage Manager binds it to either the default management class or the management class specified in your include-exclude list.

If the backup copy group for the management class specifies keeping multiple backup versions of the file, and you request multiple backups, the server always has one active backup version (the current version) and one or more inactive backup versions of the file. All backup versions of a file are bound to the same management class and are managed based on the attributes in the backup copy group.

When you archive a file for the first time, Tivoli Storage Manager binds it to the default management class, to the management class specified in your include-exclude list, or to a management class you specify when modifying your archive options during an archive.

Archived files are never rebound to a different management class. If you change the management class for a file using an `include.archive` statement, the `archmc` option, or through the backup-archive client GUI, any previous copies of the file that you archived remain bound to the management class specified when you archived them.

If a file is deleted on the client system then that inactive objects of the file are not rebound.

For information about how to associate files and directories with management classes, see the Tivoli Storage Manager server documentation.

Rebind backup versions of files

Rebinding associates a file or a logical volume image with a new management class.

Backups of files are bound again to a different management class in the following conditions. In each condition, the files (active and inactive) are not bound again until the next backup.

- You specify a different management class in an Include statement to change the management class for the file. The backups are managed based on the old management class until you run another backup.
- Your administrator deletes the management class from your active policy set. The default management class is used to manage the backup versions when you back up the file again.
- Your administrator assigns your client node to a different policy domain and the active policy set in that domain does not have a management class with the same name. The default management class for the new policy domain is used to manage the backup versions.

For information about how to associate files and directories with management classes, see the Tivoli Storage Manager server documentation.

Retention grace period

Tivoli Storage Manager also provides a *backup retention grace period* and an *archive retention grace period* to help protect your backup and archive data when it is unable to rebind a file to an appropriate management class.

The backup retention grace period is in the following cases:

- You change the management class for a file, but neither the default management class nor the new management class contain a backup copy group.
- The management class to which a file is bound no longer exists, and the default management class does not contain a backup copy group.

The backup retention grace period, defined in your policy domain, starts when you run an incremental backup. The default is 30 days. However, your administrator can lengthen or shorten this period.

When the Tivoli Storage Manager server manages a file using the backup retention grace period, it does not create any new backup versions of the file. All existing backup versions of the file expire 30 days (or the number of days specified in your policy domain) from the day they are marked inactive.

Archive copies are never rebound because each archive operation creates a different archive copy. Archive copies remain bound to the management class name specified when the user archived them. If the management class to which an archive copy is bound no longer exists or no longer contains an archive copy group, the server uses the default management class. If you later change or replace the default management class, the server uses the updated default management class to manage the archive copy. If the default management class does not contain an archive copy group, the server uses the archive retention grace period specified for the policy domain.

Event-based policy retention protection

All management classes with an archive copy group must specify a retention period, for example, the number of days that an archived object is stored on the server before being deleted.

Event-based policy provides the option of beginning the retention period either at the time the object is archived or at a later date when an activation event is sent to the server for that object.

Setting the copy group value `RETINIT=CREATE` starts the data retention period when the file is archived. Using the copy group value `RETINIT=EVENT` starts the data retention period when the server is notified that the event has occurred.

The following example demonstrates this concept:

The user has two files, `create.file` and `event.file`. The user has available two management classes; `CREATE`, with `RETINIT=CREATE`, and `EVENT`, with `RETINIT=EVENT`. Both management classes have a 60-day retention period. The user, on the same day, archives both files:

```
dsmc archive create.file -archmc=CREATE
dsmc archive event.file -archmc=EVENT
```

Ten days later, the user issues the **set event** `-type=hold` command for the `create.file` file, so the file cannot be deleted. On the same day the user issues the **set event** `-type=activate` for the `event.file` file. At this time, `create.file` has 50 days left on its retention period, and `event.file` has 60 days. If no other action is taken, `create.file` remains on the server forever, and `event.file` is expired 70 days after it was created (60 days after its event occurred). However, if 20 days after the initial archive, the user issues **set event** `-type=release` for the `create.file` file. Thirty days of its retention period have passed, so the file is expired in 30 days (the hold does not extend the retention period).

For information about the `RETINIT` copy group value, see the Tivoli Storage Manager server documentation.

Related reference:

“Set Event” on page 731

Archive files on a data retention server

Up to this point, there is no difference between archiving files on a normal server or a data retention server.

The following example demonstrates the differences between the two servers, and what can be done at day 5:

If the files were archived on a non-data retention server, the user can issue the **delete archive** `create.file event.file` command and both files are deleted. If the files were archived on a data retention server, the same command fails both files. The data retention server forces the user to keep archives until the stated retention criteria are met.

Now here is the difference at day 15 (after the hold):

The **delete archive** *create.file event.file* command on the non-data retention server now deletes *event.file*, but returns a *cannot delete* error for *create.file* because it is in hold status. That same command to a data retention server still rejects the deletion of both files.

Chapter 10. Processing options

You can use defaults for processing client options or you can tailor the processing options to meet your specific needs. Read about an overview of processing options and explore the options reference that provides detailed information about each option.

Related concepts:

“Using options with commands” on page 307

Related reference:

“Reading syntax diagrams” on page xiv

Processing options overview

Tivoli Storage Manager uses *processing options* to control communications, backup-archive processing, and other types of processing.

You can specify processing options in the client system-options file (`dsm.sys`), client user-options file (`dsm.opt`), or on the command line.

You can set the following types of options:

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

The backup-archive client also includes a group of client command options that you can enter only on the command line with specific commands. You can override some of the options in your options file by entering them with appropriate backup-archive commands.

Related concepts:

“Entering options with a command” on page 308

Related tasks:

“Creating and modifying the client system-options file” on page 51

Communication options

You use communication options to specify how your client node communicates with a Tivoli Storage Manager server. This topic provides information about the types of communication options you can use.

For UNIX and Linux use one of the following communication protocols:

- TCP/IP
- Shared memory (AIX, HP-UX, Linux, and Solaris)

Use the `commmethod` option to specify the communication protocol.

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

Related reference:

“`Commmethod`” on page 335

TCP/IP options

To use the TCP/IP communication protocol, you must include the `tcpserveraddress` option in your client options file.

The other TCP/IP options have default values that you can modify if you want to change the default value. This topic provides information about the types of communication options you can use.

Table 54. TCP/IP options

Option	Description
<code>httpport</code> “ <code>Httpport</code> ” on page 414	Specifies a TCP/IP port address for the web client.
<code>lanfreetcpport</code> “ <code>Lanfreetcpport</code> ” on page 443	Specifies the TCP/IP port number where the Tivoli Storage Manager storage agent is listening.
<code>lanfreetcpsrveraddress</code> “ <code>Lanfreetcpsrveraddress</code> ” on page 444	Specifies the TCP/IP address for the Tivoli Storage Manager storage agent.
<code>tcpbuffsize</code> “ <code>Tcpbuffsize</code> ” on page 543	Specifies the size, in kilobytes, of the internal TCP/IP communication buffer.
<code>tcpnodelay</code> “ <code>Tcpnodelay</code> ” on page 546	Specifies whether the server or client disables the delay of sending successive small packets on the network. This option is for all UNIX clients.
<code>tcpadminport</code> “ <code>Tcpadminport</code> ” on page 542	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.
<code>tcpcadaddress</code> “ <code>Tcpcadaddress</code> ” on page 544	Specifies a TCP/IP address for <code>dsmcad</code> .
<code>tcpport</code> “ <code>Tcpport</code> ” on page 547	Specifies the TCP/IP port address for a Tivoli Storage Manager server.
<code>tcpserveraddress</code> “ <code>Tcpserveraddress</code> ” on page 548	Specifies the TCP/IP address for a Tivoli Storage Manager server.
<code>tcpwindowsize</code> “ <code>Tcpwindowsize</code> ” on page 548	Specifies the size, in kilobytes, of the TCP/IP sliding window for your client node.

Table 54. TCP/IP options (continued)

Option	Description
webports “Webports” on page 599	Enables the use of the web client outside a firewall by specifying the TCP/IP port number used by the client acceptor daemon and the web client agent service (web client agent service does not apply to Mac OS X) for communications with the web GUI.

Related reference:

“Nfstimeout” on page 461

Shared memory options

This topic provides information on the shared memory options that you can use.

Table 55. Shared memory communication options

Option	Description
lanfreeshmport “Lanfreeshmport” on page 442	Specifies the unique number that is used by the client and the storage agent to identify shared memory area used for communications.
lanfreeshmport “Shmport” on page 516	Specifies the unique number that is used by the client and the server to identify shared memory area used for communications.

Server options

Use the servername option in your dsm.sys file to begin a group of options (stanzas) used to connect to the Tivoli Storage Manager server.

You can set up multiple groups of stanzas in the dsm.sys file to connect to different servers. Each servername stanza must have listed below it all client option stanzas required to establish communication with a server. The stanza list can also contain other options for backup-archive operations.

If your client system-options file contains only one stanza - Your client node contacts the server you specify in that stanza for all services.

If your client system-options file contains more than one stanza - You can specify a default server with the defaultserver option. If you do not specify a default server, Tivoli Storage Manager contacts the server you specify in the first stanza of your dsm.sys file.

Place the defaultserver option at the beginning of your dsm.sys file before any server stanzas. See “Defaultserver” on page 350 for more information.

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

Note: You cannot override the migration server specified in the client system-options file.

Table 56 on page 290 shows a sample dsm.sys file.

Table 56. Sample client system-options file

Sample dsm.sys file	
DEFAULTServer	server2
SErvername	server1
NODename	node1
COMMMethod	TCPip
TCPPort	1500
TCPServeraddress	node.domain.company.com
PASSWORDAccess	generate
GROups	system adsm
USERS	ashton stewart kaitlin
INCLExc1	/adm/adsm/backup1.exc1
SErvername	server2
COMMMethod	SHAREdmem
shmport	1520
PASSWORDAccess	prompt
GROups	system adsm
USERS	danielle derek brant
INCLExc1	/adm/adsm/backup2.exc1

Backup and archive processing options

You can specify client options to control some aspects of backup and archive processing.

Table 57. Backup and archive processing options

Option	Description
afmskipuncachedfiles "Afmskipuncachedfiles" on page 317	Use the afmskipuncachedfiles option to specify whether uncached and dirty files in General Parallel File System (GPFS™) Active File Management file sets are processed for backup, archive, and migration operations.
archmc "Archmc" on page 318	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.
archsymb linkasfile "Archsymb linkasfile" on page 319	Specifies whether you want the client to follow a symbolic link and archive the file or directory to which it points, or archive the symbolic link only.
asnodename "Asnodename" on page 319	Use the asnodename option to allow agent nodes to back up or restore data on behalf of another node (the target node). This option enables concurrent operations from multiple nodes to store data to the same target node and file space in parallel.

Table 57. Backup and archive processing options (continued)

Option	Description
automount "Automount" on page 328	Use this option with the domain option to specify all automounted file systems that the client tries to mount at the following points in time: • When the backup-archive client starts • When the backup is started • When the backup-archive client reaches an automounted file system during backup
autofsrename "Autofsrename" on page 326	Specifies whether to rename an existing file space on a Unicode-enabled server so a Unicode-enabled file space can be created for the current operation.
changingretries "Changingretries" on page 333	Specifies the number of times the client attempts to back up or archive a file that is in use.
compressalways "Compressalways" on page 338	The compressalways option specifies whether to continue compressing an object if it grows during compression. Use this option with the compression option.
compression "Compression" on page 339	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.
createnewbase "Createnewbase" on page 342	The createnewbase option creates a base snapshot and uses it as a source to run a full incremental. Setting this option ensures the backup of any files that might have been skipped during the snapshot difference incremental.
deduplication "Deduplication" on page 349	Specifies whether to eliminate redundant data on the client side when the client transfers data to the Tivoli Storage Manager server during backup or archive processing.
dedupcachepath "Dedupcachepath" on page 348	Specifies the location where the client-side data deduplication cache database is created, if the enablededupcache=yes option is set during backup or archive processing.
dedupcachesize "Dedupcachesize" on page 348	Determines the maximum size of the data deduplication cache file.
enablededupcache "Enablededupcache" on page 378	Specifies whether you want to enable client-side data deduplication cache, so that the backup-archive client gets the changed data from the cache.

Table 57. Backup and archive processing options (continued)

Option	Description
deletefiles "Deletefiles" on page 351	Use the deletefiles option with the archive command to delete files from your workstation after you archive them. You can also use this option with the restore image command and the incremental option to delete files from the restored image if they were deleted after the image was created.
description "Description" on page 351	The description option assigns or specifies a description for files when the client performs archive, delete, retrieve, query archive, or query backupset operations.
detail "Detail" on page 353	Use the detail option to list management class, file space, backup, and archive information, depending on the command with which it is used.
diffsnapshot "Diffsnapshot" on page 354	Use the diffsnapshot option to determine whether the client creates a differential snapshot.
dirmc "Dirmc" on page 356	Specifies the management class to use for directories. If you do not specify this option, the client uses the management class in the active policy set of your policy domain with the longest retention period.
dirsonly "Dirsonly" on page 357	Backs up, restores, archives, retrieves, or queries directories only.
diskcachelocation "Diskcachelocation" on page 359	Specifies the location where the disk cache database is created if the option memoryefficient=diskcachemethod option is set during an incremental backup.
domain "Domain" on page 361	Specifies the file systems to include in your default client domain for an incremental backup.
domain.image "Domain.image" on page 366	Specifies the mounted file systems and raw logical volumes that you want to include in your client domain for an image backup. This option is for AIX, HP-UX, Linux x86_64, Linux on POWER, and Solaris only.
domain.nas "Domain.nas" on page 366	Specifies the volumes to include in your default domain for NAS image backups.
domain.vmfull "Domain.vmfull" on page 369	Specifies the virtual machines to include in full image backups of VMware virtual machines.
efsdecrypt "Efsdecrypt" on page 376	Specifies whether files encrypted by an AIX Encrypted File System (EFS) are read in encrypted or decrypted format.

Table 57. Backup and archive processing options (continued)

Option	Description
enablearchiveretentionprotection "Enablearchiveretentionprotection" on page 377	Allows the client to connect to a data retention server.
enablelanfree "Enablelanfree" on page 382	Specifies whether to enable an available LAN-free path to a storage area network (SAN) attached storage device.
"Exclude options" on page 389 exclude exclude.backup exclude.file exclude.file.backup	Use these options to exclude a file or group of files from backup services and space management services (if the HSM client is installed). The exclude.backup option excludes only files from normal backup, but not from HSM.
encryptiontype "Encryptiontype" on page 383	Select AES-256, AES-128 or DES-56 bit data encryption. AES 256-bit data encryption provides the highest level of data encryption.
encryptkey "Encryptkey" on page 384	Specifies whether to save the encryption key password locally when the client performs a backup-archive operation or whether to prompt for the encryption key password.
exclude.archive "Exclude options" on page 389	Excludes a file or a group of files that match the pattern from archive services only.
exclude.attribute.symlink "Exclude options" on page 389	Excludes a file or a group of files that are symbolic links or aliases (aliases apply to Mac OS X) from backup processing only.
exclude.compression "Exclude options" on page 389	Excludes files from compression processing if you set the compression option to <i>yes</i> . This option applies to backups and archives.
exclude.dir "Exclude options" on page 389	Excludes a directory, its files, and all its subdirectories and their files from backup processing.
exclude.encrypt "Exclude options" on page 389	Excludes specified files from encryption processing.
exclude.fs "Exclude options" on page 389	Excludes file spaces that match a pattern. This option is valid for all UNIX clients.
exclude.fs.nas "Exclude options" on page 389	Excludes file systems on the NAS file server from an image backup when used with the backup nas command. This option is for AIX and Solaris clients only.
exclude.image "Exclude options" on page 389	Excludes mounted file systems and raw logical volumes that match the specified pattern from full image backup operations. This option is valid only for AIX, HP-UX, Solaris, and all Linux clients.
fbbranch "Fbbranch" on page 396	Specifies the branch ID of the remote FastBack server to back up or archive.

Table 57. Backup and archive processing options (continued)

Option	Description
fbclientname "Fbclientname" on page 396	Specifies the name of one or more FastBack clients to back up from the backup proxy.
fbpolicyname "Fbpolicyname" on page 398	Specifies the name of one or more Tivoli Storage Manager FastBack policies that you want to back up from the backup proxy.
fbreposlocation "Fbreposlocation" on page 399	Specifies the location of the Tivoli Storage Manager FastBack repository for the Tivoli Storage Manager client proxy to connect to issue MOUNT DUMP , MOUNT ADD , and MOUNT DEL commands.
fbserver "Fbserver" on page 400	Specifies host name of the FastBack server workstation or the FastBack Disaster Recovery Hub workstation that owns the repository that is specified by the fbreposlocation option.
fbvolumename "Fbvolumename" on page 401	Specifies the name of one or more Tivoli Storage Manager FastBack volumes to back up from the backup proxy.
filelist "Filelist" on page 402	Specifies a list of files to be processed for the command. The client opens the designated file list and processes the files that are listed within according to the command.
filesonly "Filesonly" on page 406	Backs up, restores, retrieves, or queries files only.
groupname "Groupname" on page 412	Use this option with the backup group command to specify the fully qualified name of the group leader for a group.
guitreeviewafterbackup "Guitreeviewafterbackup" on page 413	Specifies whether the client GUI is returned to the Backup, Restore, Archive, or Retrieve window after an operation finishes.
ieobjtype "Ieobjtype" on page 416	Specifies an object type for a client-side data deduplication operation. This option is used with the include.dedup and exclude.dedup options.
imagegapsize "Imagegapsize" on page 417	Specifies the minimum size of empty regions on a volume that you want to skip during image backup. This option is valid for AIX JFS2 clients.
incl excl "Incl excl" on page 420	Specifies the path and file name of an include-exclude options file.
"Include options" on page 421	Use these options to include files or assign management classes for backup processing.
include include.backup include.file	
include.archive "Include options" on page 421	Includes files or assigns management classes for archive processing.

Table 57. Backup and archive processing options (continued)

Option	Description
<code>include.attribute.symlink</code> "Include options" on page 421	Includes a file or a group of files that are symbolic links or aliases (aliases apply to Mac OS X) within broad group of excluded files for backup processing only.
<code>include.compression</code> "Include options" on page 421	Includes files for compression processing if you set the <code>compression</code> option to <i>yes</i> . This option applies to backups and archives.
<code>include.encrypt</code> "Include options" on page 421	Includes the specified files for encryption processing. By default, the client does not perform encryption processing.
<code>include.fs</code> "Include options" on page 421	Use the <code>include.fs</code> option to control how the client processes your file space for incremental backup.
<code>include.fs.nas</code> "Include options" on page 421	Use the <code>include.fs.nas</code> option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether the client saves Table of Contents (TOC) information during a NAS file system image backup by using the <code>toc</code> option with the <code>include.fs.nas</code> option in your <code>dsm.sys</code> file. For more information, see "Toc" on page 552. This option is valid for AIX and Solaris clients only.
<code>include.image</code> "Include options" on page 421	Specifies a file system or logical volume to be included for image backup processing. This option also provides a way to specify an explicit management class assignment for a specified file system or logical volume. The backup image command ignores all other include options. This option is valid for AIX, HP-UX, Solaris, and all Linux clients.
<code>incrbydate</code> "Incrbydate" on page 436	Use with the incremental command to request an incremental backup by date.
<code>incremental</code> "Incremental" on page 438	Use with the restore image command to ensure that any changes that were made to the base image are also applied to the restored image. This option is only valid for AIX, HP-UX, Linux x86_64, Linux on POWER, and Solaris.
<code>memoryefficientbackup</code> "Memoryefficientbackup" on page 452	Specifies a memory-saving backup algorithm for incremental backups when used with the incremental command.

Table 57. Backup and archive processing options (continued)

Option	Description
mode "Mode" on page 454	Use the mode option with these commands, as follows: backup image To specify whether to perform a selective or incremental image backup of client file systems. backup nas To specify whether to perform a full or differential image backup of NAS file systems. backup group To specify whether to perform a full or differential group backup that contains a list of files from one or more file space origins. backup vm To specify whether to perform a selective or incremental backup of VMware systems.
monitor "Monitor" on page 457	Specifies whether you want to monitor an image backup of file systems that belong to a Network Attached Storage (NAS) file server.
noprompt "Noprompt" on page 464	Suppresses the confirmation prompt that is presented by the delete group , delete archive , expire , and set event commands.
noprompt "Noprompt" on page 464	Suppresses the confirmation prompt that is presented by the delete group , delete archive , expire , restore image , and set event commands.
nojournal "Nojournal" on page 463	Use this option with the incremental command to specify that you want to perform the traditional full incremental backup, instead of the default journal-based backup.
optfile "Optfile" on page 467	Specifies the client user-options file that you want to use when you start a backup-archive client session.
postsnapshotcmd "Postsnapshotcmd" on page 476	During a snapshot-based backup, this option allows you to manually open an application after the snapshot is created. This option is valid only for AIX JFS2 or Linux LVM snapshot-based operations.

Table 57. Backup and archive processing options (continued)

Option	Description
preservelastaccessdate "Preservelastaccessdate" on page 479	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 after a backup or archive operation. By default, the client does not reset the last access date of any backed up or archived files to their original value before the backup or archive operation.
presnapshotcmd "Presnapshotcmd" on page 483	During a snapshot-based backup operation, this option allows you to manually quiesce an application before the snapshot is created. This option is valid only for AIX JFS2 or Linux LVM snapshot-based operations.
removeoperandlimit "Removeoperandlimit" on page 488	Specifies that the client removes the 20-operand limit. If you specify the <code>removeoperandlimit</code> option with the incremental , selective , or archive commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits.
skipacl "Skipacl" on page 517	Specifies whether to skip ACL processing completely.
skipaclupdatecheck "Skipaclupdatecheck" on page 518	Specifies whether to perform checksum and size comparisons before and after backup and during incremental processing.
snapdiff "Snapdiff" on page 519	Specifies an incremental backup of the files reported as changed by NetApp, instead of scanning the volume and looking for files that have changed. Use this option with a NAS full volume incremental backup.
snapshotcachesize "Snapshotcachesize" on page 525	Linux and AIX only: Use this option to specify an appropriate snapshot size so that all original data blocks can be stored during file modification and deletion. A snapshot size of 100 percent ensures a valid snapshot. The default value is 100 percent.
snapshotproviderfs "Snapshotproviderfs" on page 526	Use the <code>snapshotproviderfs</code> option to enable snapshot-based file backup and archive operations, and to specify a snapshot provider. You must be a root user to perform a snapshot-based file backup or archive operation. If you are not a root user, the operation fails with an error message.

Table 57. Backup and archive processing options (continued)

Option	Description
snapshotproviderimage "Snapshotproviderimage" on page 527	Use the snapshotproviderimage option to enable snapshot-based image backup, and to specify a snapshot provider. You must be a root user to perform a snapshot-based image backup operation. If you are not a root user, the operation fails with an error message.
snapshotroot "Snapshotroot" on page 528	Use the snapshotroot option with the incremental , selective , or archive commands with an independent software vendor application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. This option is valid for all UNIX and Linux clients.
subdir "Subdir" on page 539	Specifies whether to include subdirectories of a named directory.
tapeprompt "Tapeprompt" on page 541	Specifies whether you want the client to wait for a tape mount if it is required for a backup, archive, restore, or retrieve process, or to be prompted for a choice.
toc "Toc" on page 552	Use the toc option with the backup nas command or the include.fs.nas option to specify whether the client saves Table of Contents (TOC) information for each file system backup. If you save TOC information, you can use the QUERY TOC server command to determine the contents of a file system backup 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.
type "Type" on page 556	Use the type option with the query node command to specify the type of node to query.
v2archive "V2archive" on page 560	Use the v2archive option with the archive command to archive only files to the server. The client does not process directories that exist in the path of the source file specification.
virtualfsname "Virtualfsname" on page 565 (does not apply to Mac OS X)	Use this option with the backup group command to specify the name of the container for the group on which you want to perform the operation.
virtualmountpoint "Virtualmountpoint" on page 566	Defines a virtual mount point for a file system if you want to consider files for backup that begin with a specific directory within that file system.

Table 57. Backup and archive processing options (continued)

Option	Description
vmchost "Vmchost" on page 570	Used with the backup VM , restore VM , or query VM commands to specify the host name of the VMware VirtualCenter or ESX server where the commands are directed.
vmcpw "Vmcpw" on page 571	Used with the backup VM , restore VM , or query VM commands to specify the password of the VirtualCenter or ESX user that is specified with the vmcuser option.
vmcuser "Vmcuser" on page 573	Used with the backup VM , restore VM , or query VM commands to specify the user name for the VMware VirtualCenter or ESX server where the commands are directed.
vmmxvirtualdisks "Vmmxvirtualdisks" on page 582	Used with the backup VM command to specify the maximum size of the VMware virtual machine disks (VMDKs) to include in a backup operation.
vmskipmaxvirtualdisks "Vmskipmaxvirtualdisks" on page 590	Used with the backup VM command to specify how the backup operation processes VMware virtual machine disks (VMDKs) that exceed the maximum disk size. In V7.1.3 and earlier, the vmskipmaxvirtualdisks option was named vmskipmaxvmdks.

Restore and retrieve processing options

You can use client options to control some aspects of restore and retrieve processing.

Table 58 lists the restore and retrieve processing options that are available.

Table 58. Restore and retrieve processing options

Option	Description
dironly "Dironly" on page 357	Qualifies the operation (backup, archive, restore, retrieve) to process directories alone.
disablenqr "Disablenqr" on page 358	Specifies whether the backup-archive client can use the no-query restore method for restoring files and directories from the server.
filelist "Filelist" on page 402	Specifies a file that contains a list of files to be processed by the specified command.
filesonly "Filesonly" on page 406	Qualifies the operation (backup, archive, restore, retrieve) to process files alone.
followsymbolic "Followsymbolic" on page 407	Specifies whether you want to restore files to symbolic links or use a symbolic link as a virtual mount point.

Table 58. Restore and retrieve processing options (continued)

Option	Description
fromdate "Fromdate" on page 409	Use the fromdate option with the fromtime option to specify a date and time from which you want to search for backups or archives during a restore, retrieve, or query operation.
fromnode "Fromnode" on page 410	Permits one node to perform commands for another node. A user on another node must use the set access command to give you permission to query, restore, or retrieve files or images for the other node.
fromowner "Fromowner" on page 411	Displays file spaces for an alternative owner. Also specifies an alternative owner from which to restore or retrieve files.
fromtime "Fromtime" on page 412	Use the fromtime option with the fromdate option to specify a beginning time from which you want to search for backups or archives during a restore, retrieve, or query operation.
guitreeviewafterbackup "Guitreeviewafterbackup" on page 413	Specifies whether the client GUI is returned to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.
ifnewer "Ifnewer" on page 417	Replaces an existing file with the latest backup version only if the backup version is newer than the existing file.
imagetofile "Imagetofile" on page 418	Use the imagetofile option with the restore image command to specify that you want to restore the source image to a file. You might need to restore the image to a file in the event of bad sectors present on the target volume, or if you want to do some manipulations with the image data. This option is valid for AIX, HP-UX, Linux, and Solaris clients.
inactive "Inactive" on page 419	Displays a list of active and inactive files when used with the pick option.
latest "Latest" on page 445	Restores the most recent backup version of a file whether it is active or inactive.
localbackupset "Localbackupset" on page 446	Specifies whether the backup-archive client GUI bypasses initial logon with the server to restore a local backup set on a stand-alone workstation.
makesparsefile "Makesparsefile" on page 447 (does not apply to Mac OS X)	Use the makesparsefile option with the restore or retrieve commands to specify how sparse files are re-created.
monitor "Monitor" on page 457	Specifies whether you want to monitor an image restore of one or more file systems that belong to a network-attached storage (NAS) file server.
noprompt "Noprompt" on page 464	suppresses the confirmation prompt that is presented by the delete group , delete archive , expire , and set event commands.
noprompt "Noprompt" on page 464	suppresses the confirmation prompt that is presented by the delete group , delete archive , expire , restore image , and set event commands.

Table 58. Restore and retrieve processing options (continued)

Option	Description
optfile "Optfile" on page 467	Specifies the client user-options file that you want to use when you start a backup-archive client session.
pick "Pick" on page 472	Creates a list of backup versions, images, or archive copies that match the file specification you enter. From the list, you can select the versions to process. Include the inactive option to view both active and inactive objects.
pitdate "Pitdate" on page 473	Use the pitdate option with the pittime option to establish a point in time for which you want to display or restore the latest version of your backups.
pittime "Pittime" on page 474	Use the pittime option with the pitdate option to establish a point in time for which you want to display or restore the latest version of your backups.
preservepath "Preservepath" on page 480	Specifies how much of the source path to reproduce as part of the target directory path when you restore or retrieve files to a new location.
replace "Replace" on page 489	Specifies whether to overwrite an existing file, or to prompt you for your selection when you restore or retrieve files.
showmembers "Showmembers" on page 517 (does not apply to Mac OS X)	Displays all members of a group.
subdir "Subdir" on page 539	Specifies whether you want to include subdirectories of a named directory.
tapeprompt "Tapeprompt" on page 541	Specifies whether you want the backup-archive client to wait for a tape that is required for a restore or retrieve to be mounted, or to prompt you for your choice.
todate "Todate" on page 554	Use the todate option with the totime option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation.
totime "Totime" on page 554	Use the totime option with the todate option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation.
type "Type" on page 556	Use the type option with the query node command to specify the type of node to query.
verifyimage "Verifyimage" on page 565	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, the client issues a warning message on the console and in the error log.

The following options are backup-archive client options that apply to Tivoli Storage Manager HSM for Windows migrated files. For more information about these options, see the IBM Knowledge Center topics at <https://www.ibm.com/support/knowledgecenter/SSERFH>.

- Checkreparsecontent
- Restorecheckstubaccess
- Restoremigstate
- Skipmigrated

The following options are backup-archive client options that apply to Tivoli Storage Manager for Space Management migrated files. For more information about these options, see the IBM Knowledge Center topics at <http://www.ibm.com/support/knowledgecenter/SSSR2R>.

- Restoremigstate
- Skipmigrated

Scheduling options

This topic discusses the options that you can use to regulate central scheduling. The backup-archive client uses scheduling options only when the Scheduler is running.

Table 59 lists the scheduling options that are available.

Table 59. Scheduling options

Option	Description
cadlistenonport “Cadlistenonport” on page 332	Specifies whether to open listening ports for the client acceptor when the client acceptor is used to manage schedules in polling mode.
managedservices “Managedservices” on page 448	Specifies whether the client acceptor manages the web client, the scheduler, or both.
maxcmdretries “Maxcmdretries” on page 450	Specifies the maximum number of times the client scheduler attempts to process a scheduled command that fails.
postschedulecmd/postnschedulecmd “Postschedulecmd/Postnschedulecmd” on page 475	Specifies a command to process after running a schedule.
preschedulecmd/prenschedulecmd “Preschedulecmd/Prenschedulecmd” on page 478	Specifies a command to process before running a schedule.
querschedperiod “Querschedperiod” on page 484	Specifies the number of hours the client scheduler waits between attempts to contact the server for scheduled work.
retryperiod “Retryperiod” on page 501	Specifies the number of minutes the client scheduler waits between attempts to process a scheduled command that fails or between unsuccessful attempts to report results to the server.
schedcmddisabled “Schedcmddisabled” on page 503	Specifies whether to disable the scheduling of generic commands specified by your Tivoli Storage Manager administrator.

Table 59. Scheduling options (continued)

Option	Description
schedcmduser (server defined only) “Scheduler options for commands” on page 269	The scheduler executes commands under a uid of 0, however, there might be some users who have a different user ID. In this case, your Tivoli Storage Manager administrator can define schedules and allow these schedules to be executed under a uid other than 0, using this option. The Tivoli Storage Manager Client API does not support this option.
schedlogmax “Schedlogmax” on page 504	Specifies the maximum size of the scheduler log and web client log, in megabytes.
schedlogname “Schedlogname” on page 506	Specifies the path and file name where you want to store schedule log information.
schedlogretention “Schedlogretention” on page 507	Specifies the number of days to keep log file entries in the schedule log and the web client log, and whether to save pruned entries.
schedmode “Schedmode” on page 509	Specifies which schedule mode to use, <i>polling</i> or <i>prompted</i> .
schedrestretrdisabled “Schedrestretrdisabled” on page 510	Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing restore or retrieve schedule operations.
sessioninitiation “Sessioninitiation” on page 514	Use the sessioninitiation option to control whether the server or client initiates sessions through a firewall. The default is that the client can initiate sessions.
srvprepostscheddisabled “Srvprepostscheddisabled” on page 531	Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing pre-schedule and post-schedule commands when performing scheduled operations.
srvprepostsnapdisabled “Srvprepostsnapdisabled” on page 532	Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing pre-snapshot and post-snapshot commands when performing scheduled image snapshot backup operations.
tcpclientaddress “Tcpclientaddress” on page 545	Specifies a TCP/IP address if your client node has more than one address, and you want the server to contact an address other than the one that was used to make the first server contact. The server uses this address when it begins the server prompted scheduled operation. See schedmode <i>prompted</i> (“Schedmode” on page 509) for details.
tcpclientport “Tcpclientport” on page 545	Specifies a TCP/IP port number for the server to contact the client when the server begins the server prompted scheduled operation. See schedmode <i>prompted</i> (“Schedmode” on page 509) for details.

Format and language options

Format and language options allow you to select different formats for date, time and numbers for different languages.

Format options allow you to select different formats for date, time, and numbers.

Table 60. Format and language options

Option	Description
dateformat "Dateformat" on page 344	Specifies the format for displaying dates.
numberformat "Numberformat" on page 465	Specifies the format for displaying numbers.
timeformat "Timeformat" on page 549	Specifies the format for displaying time.

Command processing options

This topic explains the options that you can use with the backup-archive client commands.

Command processing options allow you to control some of the formatting of data on your terminal screen.

Table 61. Command processing options

Option	Description
quiet "Quiet" on page 486	Limits the number of messages that are displayed on your screen during processing. This option can be overridden by the server.
scrolllines "Scrolllines" on page 511	Specifies the number of lines of information that are displayed on your screen at one time. Use this option only when scrollprompt is set to <i>yes</i> .
scrollprompt "Scrollprompt" on page 512	Specifies whether you want the backup-archive client to stop and wait after displaying the number of lines of information you specified with the scrolllines option, or scroll through and stop at the end of the information list.
verbose "Verbose" on page 564	Specifies that processing information should be displayed on your screen. The alternative is quiet. This option can be overridden by the server.

Authorization options

Authorization options control access to the Tivoli Storage Manager server.

Table 62 lists the authorization options that are available.

Table 62. Authorization options

Option	Description
autodeploy "Autodeploy" on page 325	Specifies whether you want to enable or disable an automatic deployment of the client if a restart is required.

Table 62. Authorization options (continued)

Option	Description
groups "Groups (deprecated)" on page 413	Specifies the groups on your workstation that you want to authorize to request Tivoli Storage Manager services from the server.
password "Password" on page 468	Specifies a Tivoli Storage Manager password.
passwordaccess "Passwordaccess" on page 470	Specifies whether you want to use a generated password or be prompted for a password each time you start the client.
passworddir "Passworddir" on page 471	Specifies the directory in which you want to store the automatically generated password for your client node. The encryption key and password are encrypted and stored in the TSM.sth file.
revokeremoteaccess "Revokeremoteaccess" on page 502	Restricts an administrator with client access privileges from accessing your workstation through the web client.
users "Users (deprecated)" on page 560	Authorizes specific users on your workstation to request services from a server.

Error processing options

Error processing options specify the name of the error log file and how the backup-archive client treats the entries in the log file.

Table 63 lists the error processing options that are available.

Table 63. Error processing options

Option	Description
errorlogmax "Errorlogmax" on page 385	Specifies the maximum size of the error log, in megabytes.
errorlogname "Errorlogname" on page 387	Specifies the fully qualified path and file name of the file where you want to store information about errors that occur during processing.
errorlogretention "Errorlogretention" on page 388	Specifies how many days to maintain error log entries before pruning, and whether to save the pruned entries.

Transaction processing options

Transaction processing options control how transactions are processed between the Tivoli Storage Manager client and server.

Table 64 lists the transaction processing options that are available.

Table 64. Transaction processing options

Option	Description
collocatebyfilespec "Collocatebyfilespec" on page 334	Specifies that you want the backup-archive client to use only one server session to send objects generated from one file specification. Setting the collocatebyfilespec 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).
commrestartduration "Commrestartduration" on page 337	Specifies the maximum number of minutes you want the client to try to reconnect to a Tivoli Storage Manager server after a communication error occurs.
commrestartinterval "Commrestartinterval" on page 337	Specifies the number of seconds you want the client to wait between attempts to reconnect to a Tivoli Storage Manager server after a communication error occurs.
diskbuffsize "Diskbuffsize" on page 359	Specifies the maximum disk I/O buffer size (in kilobytes) that the client can use when reading files.
largecommbuffers "Diskbuffsize" on page 359	This option has been replaced by the diskbuffsize option. At this time, largecommbuffers is still accepted by the Tivoli Storage Manager client in order to ease the transition to the new option. However, the value specified by largecommbuffers is ignored in favor of the diskbuffsize setting. Important: Discontinue the use of largecommbuffers because future releases of Tivoli Storage Manager might not accept this option.
nfstimeout "Nfstimeout" on page 461	Specifies the number of seconds the server waits for a status system call on an NFS file system before it times out.
resourceutilization "Resourceutilization" on page 499	Use the resourceutilization option in your dsm.sys file to regulate the level of resources the Tivoli Storage Manager server and client can use during processing.
txnbytelimit "Txnbytelimit" on page 555	Specifies the number of kilobytes the client program buffers before it sends a transaction to the server.

Web client options

Several backup-archive client options are used to configure the Tivoli Storage Manager web client.

Table 65 lists the web client options that are available.

Table 65. Web client options

Option	Description
httpport "Httpport" on page 414	Specifies a TCP/IP port address for the web client.
managedservices "Managedservices" on page 448	Specifies whether the client acceptor daemon manages the web client, the scheduler, or both.

Table 65. Web client options (continued)

Option	Description
revokeremoteaccess “Revokeremoteaccess” on page 502	Restricts administrator access on a client workstation through the web client.
webports “Webports” on page 599	Enables the use of the web client outside a firewall by specifying the TCP/IP port number used by the client acceptor daemon and the web Client Agent service for communications with the web client.

Diagnostics options

Use the **query systeminfo** command to gather Tivoli Storage Manager system information and output this information to a file or the console.

The **query systeminfo** command is intended primarily as a diagnostic aid. You can submit the resulting information to technical support personnel for problem diagnosis.

Table 66 lists the diagnostics options that are available.

Table 66. Diagnostics options

Option	Description
console “Console” on page 341	Use the console option with the query systeminfo command to output system information to the console.
filename “Filename” on page 405	Use the filename option with the query systeminfo command to specify a file name in which to store the system information.

Related reference:

“Query Systeminfo” on page 686

Using options with commands

You can override some of the options in your client options file (dsm.opt) file by entering them with appropriate backup-archive client commands.

You can override some of the options in your dsm.sys file or client user-options file (dsm.opt) by entering them with appropriate backup-archive client commands.

The client 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.

The client also includes a group of client command options that you can enter *only* on the command line with specific commands. For a complete list of command-line options, a description, and where to go for more information, see Table 67 on page 309.

Entering options with a command

You must follow the general rules for entering options with a command.

- Enter a command, a dash (-), the option name, an equal sign (=), and the option value or parameter. Do not include spaces on either side of the = sign.

Here are examples of this syntax on different clients:

```
dsmc archive -description="year end 1999" /home/
```

- For options that do not include parameters, enter a command, a dash (-), and the option name. For example,

```
dsmc incremental -quiet
```

Note: Use a leading dash (-) to indicate that the following text is the name of an option. If an object name begins with a dash, you must surround it in either single quotation marks (') or quotation marks ("). Most operating system command line processors strip the quotation marks before the command-line arguments are submitted to the Tivoli Storage Manager client application. In such cases, by using escape characters or doubling the quotation marks allows the client to receive the quoted object name. In loop mode, surround such objects in either single quotation marks (') or quotation marks (").

- Enter either the option name, or an abbreviation for the option name. For example, to enter the latest option, enter either -lat or -latest. The capital letters in the syntax of each option indicate the minimum abbreviation for that option name.
- Enter options before or after command parameters. For example, you can enter the option before or after a file specification:

```
dsmc selective -subdir=yes "/home/devel/proj1/*"  
dsmc selective "/home/devel/proj1/*" -subdir=yes
```

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

```
dsmc archive -description="Project A" "/home/devel/proj1/*"
```

- Most options that you enter on the command line override the value that is 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.
- On AIX, HP-UX, Solaris, Linux on z, and Mac: The maximum number of characters for a file name is 255. The maximum combined length of the file name and path name is 1024 characters. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.
- On Linux: The maximum length for a file name is 255 bytes. The maximum combined length of both the file name and path name is 4096 bytes. This length matches the PATH_MAX that is supported by the operating system. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that comprises a path and file name can vary. The actual limitation is the number of bytes in the path and file components, which might or might not correspond to an equal number of characters.

On Linux: For archive or retrieve operations, the maximum length that you can specify for a path and file name (combined) remains at 1024 bytes.

- For Mac OS X, the maximum length of a file name is limited to 504 bytes (not characters). The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name contains can vary.

Table 67 lists client command options that you can enter only on the command line with specific commands.

Table 67. Client command options

Command option	Description	Commands
archmc "Archmc" on page 318	Use the archmc option with the archive command to specify the available management class for your policy domain to which you want to bind your archived files.	archive
class "Class" on page 333	Specifies whether to display a list of NAS objects or client objects when you use the following commands.	query backup delete filesystem query filesystem
console "Console" on page 341	Use the console option with the query systeminfo command to output system information to the console.	query systeminfo
deletefiles "Deletefiles" on page 351	Deletes the local copy of files from your workstation after they are archived on the server.	archive
deletefiles "Deletefiles" on page 351	Deletes the local copy of files from your workstation after they are archived on the server. Can also be used with the restore image command and the incremental option to delete files from the restored image that are deleted from the file space after the image is created.	archive restore image
description "Description" on page 351	Assigns or specifies a description for files when archive, delete, retrieve, or query archive operations are performed.	archive delete archive query archive query backupset retrieve
detail "Detail" on page 353	Displays management class, file space, backup, and archive information, depending on the command with which it is used.	delete filesystem query archive query backup query filesystem query mgmtclass
dirsonly "Dirsonly" on page 357	Backs up, restores, archives, retrieves, or queries directories only.	archive incremental query archive query backup restore restore backupset retrieve selective
dynamicimage "Dynamicimage" on page 375	Performs a dynamic image backup.	backup image

Table 67. Client command options (continued)

Command option	Description	Commands
filelist "Filelist" on page 402	Specifies a list of files to be processed for the command. The backup-archive client opens the designated file list and processes the files that are listed within according to the command.	archive backup group delete archive delete backup expire incremental query archive query backup restore retrieve selective
filename "Filename" on page 405	Use the filename option with the query systeminfo command to specify a file name in which to store the system information.	query systeminfo
filesonly "Filesonly" on page 406	Backs up, restores, retrieves, or queries files only.	archive incremental query archive query backup restore restore backupset retrieve selective
fromdate "Fromdate" on page 409	Use the fromdate option with the fromtime option to specify a date and time from which you want to search for backups or archives during a restore, retrieve, or query operation.	delete backup query archive query backup restore restore group retrieve
fromnode "Fromnode" on page 410	Permits one node to perform commands for another node. A user on another node must use the set access command to permit you to query, restore, or retrieve files or images for the other node.	query archive query backup query filespace query group query image query mgmtclass restore restore group restore image retrieve
fromowner "Fromowner" on page 411	Displays file spaces for another owner. Also specifies another owner from which to restore or retrieve files.	query archive query backup query group query image restore restore group restore image retrieve

Table 67. Client command options (continued)

Command option	Description	Commands
fromtime "Fromtime" on page 412	Specifies a beginning time on the specified date. Use with the fromdate option. This option is ignored if the fromdate option is absent.	query archive query backup restore restore group retrieve
groupname "Groupname" on page 412	Specifies the fully qualified name for a group.	backup group
ifnewer "Ifnewer" on page 417	Replaces existing files with the latest backup version only if the backup version is newer than the existing version.	restore restore backupset restore group retrieve
imagetofile "Imagetofile" on page 418	Use the imagetofile option with the restore image command to specify that you want to restore the source image to a file. You might need to restore the image to a file in the event of bad sectors present on the target volume, or if you want to do some manipulations with the image data. This option is valid with AIX, HP-UX, Linux, and Solaris clients.	restore image
inactive "Inactive" on page 419	Displays a list of active and inactive files when used with the pick option.	delete group query backup query group query image query nas restore restore group restore image restore nas
incrbydate "Incrbydate" on page 436	Requests an incremental backup by date.	incremental
incremental "Incremental" on page 438	Applies changes to the base image by using information from incremental backups that are made after the original image backup. This option is valid for AIX, HP-UX, Linux x86_64, Linux on POWER, and Solaris only.	restore image
latest "Latest" on page 445	Restores the most recent backup version of a file whether it is active or inactive.	restore restore group
mode "Mode" on page 454	Use the mode option with these commands, as follows: backup image To specify whether to perform a selective or incremental image backup of client file systems. backup nas To specify whether to perform a full or differential image backup of NAS file systems. backup group To specify whether to perform a full or differential group backup that contains a list of files from one or more file space origins.	backup group backup nas backup image restore nas

Table 67. Client command options (continued)

Command option	Description	Commands
monitor "Monitor" on page 457	Specifies whether you want to monitor an image backup or restore of one or more file systems that belong to a network-attached storage (NAS) file server. Specifies whether you want to monitor a restore of one or more file systems that belong to a network-attached storage (NAS) file server.	backup nas restore nas
nojurnal "Nojournal" on page 463	Use this option with the incremental command to specify that you want to perform the traditional full incremental backup, instead of the default journal-based backup.	incremental
noprompt "Noprompt" on page 464	Suppresses the confirmation prompt that is presented by the delete group , delete archive , expire , and set event commands.	delete archive delete backup delete group expire
noprompt "Noprompt" on page 464	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
optfile "Optfile" on page 467	Specifies the client user-options file that you want to use when you start a backup-archive client session.	dsmc
pick "Pick" on page 472	Creates a list of backup versions, images, or archive copies that match the file specification you enter. From the list, you can select the versions to process. Include the inactive option to view both active and inactive objects.	delete archive delete group expire query nas restore restore group restore image restore nas retrieve
pitdate "Pitdate" on page 473	Use the pitdate option with the pittime option to establish a point in time for which you want to display or restore the latest version of your backups.	query backup query group query image query nas restore restore group restore image restore nas
pittime "Pittime" on page 474	Use the pittime option with the pitdate option to establish a point in time for which you want to display or restore the latest version of your backups.	query backup query image query nas restore restore image restore nas
preservepath "Preservepath" on page 480	Specifies how much of the source path to reproduce as part of the target directory path when you restore or retrieve files to a new location.	restore restore backupset restore group retrieve

Table 67. Client command options (continued)

Command option	Description	Commands
removeoperandlimit “Removeoperandlimit” on page 488	Specifies that Tivoli Storage Manager removes the 20-operand limit. If you specify the removeoperandlimit option with the incremental , selective , or archive commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits.	incremental selective archive
showmembers “Showmembers” on page 517	Displays all members of a group.	query group restore group
todate “Todate” on page 554	Use the todate option with the totime option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation.	query archive query backup restore restore group retrieve
totime “Totime” on page 554	Use the totime option with the todate option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation.	query archive query backup restore restore group retrieve
type “Type” on page 556	Use the type option with the query node command to specify the type of node to query.	query node
v2archive “V2archive” on page 560	Use the v2archive option with the archive command to archive only files to the server. The client will not process directories that exist in the path of the source file specification.	archive
verifyimage “Verifyimage” on page 565	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, the client issues a warning message on the console and in the error log. This option is valid with AIX, HP-UX, Linux, and Solaris clients.	restore image
virtualfsname “Virtualfsname” on page 565	Specifies the name of the virtual file space for the group on which you want to run the operation.	backup group

Initial command-line-only options

A subset of client options is valid on the initial command line only. Many of these options establish the runtime environment, such as the **commethod** 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 68 on page 314 lists the options that are valid only on the initial command line.

Table 68. Options that are valid on the initial command line only

Options valid on the initial command line

commethod	preschedulecmd/prenschedulecmd (can be included in the schedule definition)
deduplication	
diskbuffsize	querschedperiod
editor	resourceutilization
enablededupcache	retryperiod
enablelanfree	schedlogmax
errorlogmax	schedlogname
errorlogname	schedlogretention
errorlogretention	schedmode
lanfreecommethod	servername
lanfreeshmport	sessioninitiation
lanfreetcppport	tcpbuffsize
maxcmdretries	tcpadaddress
nfstimeout	tcpclientaddress
nodename	tcpclientport
optfile	tcpwindowsize
password	txnbytelimit
postschedulecmd/postnschedulecmd (can be included in the schedule definition)	virtualnodename

Client options that can be set by the Tivoli Storage Manager server

Some client options can be set by the Tivoli Storage Manager server.

Table 69 on page 315 lists the options that can be set by the server.

Table 69. Options that can be set by the Tivoli Storage Manager server

Options that can be set by the Tivoli Storage Manager server

- “AfmSKIPuncachedfiles” on page 317
- “Archsymlinkasfile” on page 319
- “Changingretries” on page 333
- “Collocatebyfilespec” on page 334
- “Compressalways” on page 338
- “Compression” on page 339
- “Deduplication” on page 349
- “Dirmc” on page 356
- “Disablenqr” on page 358
- “Diskcachelocation” on page 359
- “Domain” on page 361
- “Domain.image” on page 366
- “Domain.nas” on page 366
- “Encryptiontype” on page 383
- “Encryptkey” on page 384
- “Exclude options” on page 389
- “Incl excl” on page 420
- “Include options” on page 421
- maxcandprocsmaxcandprocs
- maxmigratorsmaxmigrators
- “Memoryefficientbackup” on page 452
- “Nfstimeout” on page 461
- “Postschedulecmd/Postnschedulecmd” on page 475
- “Postsnapshotcmd” on page 476
- “Preschedulecmd/Prenschedulecmd” on page 478
- “PreserveLastAccessDate” on page 479
- “Presnapshotcmd” on page 483
- “Queryschedperiod” on page 484
- “Quiet” on page 486
- “Resourceutilization” on page 499
- “Retryperiod” on page 501
- “Schedmode” on page 509
- “Scrolllines” on page 511
- “Scrollprompt” on page 512
- “Snapshotcachesize” on page 525
- “Snapshotproviderfs” on page 526
- “Snapshotproviderimage” on page 527
- “Stagingdirectory” on page 538
- “Subdir” on page 539
- “Tapeprompt” on page 541
- “Txnbytelimit” on page 555
- “Verbose” on page 564
- “Vmchost” on page 570
- “Vmcuser” on page 573
- “Vmprocessvmwithindependent” on page 587
- “Vmprocessvmwithprdm” on page 588

Note:

1. See IBM Tivoli Storage Manager for Space Management product documentation on IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SSSR2R/welcome>.
2. See Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server product documentation on IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SSTG2D/welcome>.

Related tasks:

-  Controlling client operations through client option sets

Client options reference

The following sections contain detailed information about each of the Tivoli Storage Manager processing options.

Information for each option includes the following information:

- A description
- A syntax diagram
- Detailed descriptions of the parameters
- Examples of using the option in the client options file (if applicable)
- Examples of using the option on the command line (if applicable)

Options with a command-line example of **Does not apply** cannot be used with command line or scheduled commands.

Note:

1. Do not enclose an option value with single or quotation marks, unless the value is a file specification that contains spaces or wildcard characters. For example, the following option is not valid:

```
passwordaccess "generate"
```
2. All options in the dsm.sys file, except for the defaultserver option, must be placed within a server stanza. A server stanza is a collection of options statements in dsm.sys that begins with a SERVERName option and ends either at the next SERVERName option or the end of the file.

Absolute

Use the **absolute** option with the **incremental** command to force a backup of all files and directories that match the file specification or **domain**, even if the objects were not changed since the last incremental backup.

This option overrides the management class copy group mode parameter for backup copy groups; it does not affect the frequency parameter or any other backup copy group parameters. This option does not override **exclude** statements, so objects that are excluded from backup are not eligible for backup even when the **absolute** option is specified.

Important: Before you use the absolute option, consider the following effects that this option can have on backup and Tivoli Storage Manager server operations:

- Backups consume more server storage and database resources.
- Backups consume more network bandwidth.
- Server operations, such as inventory expiration, storage pool backup, storage pool migration, reclamation, and node replication, require more time to complete. Data deduplication might help mitigate some of these effects, but it does not avoid the processing that is required to reconstitute the deduplicated data back to its original form when the storage pool is migrated or backed up to non-deduplicated storage.

This option is valid only as a command-line parameter for the **incremental** command when you are performing the following operations:

- Full or partial progressive incremental backups of file systems or disk drives.
- Snapshot differential backups when createnewbase=yes is also specified.

To force a full backup of a file system that uses journal-based backup, specify both the `nojournal` and `absolute` options on the **incremental** command.

To use the `absolute` option on scheduled incremental backups, the Tivoli Storage Manager server administrator must create a separate backup schedule that includes the `absolute` option on the schedule's options parameter.

Supported Clients

This option is valid for all clients as a command-line parameter for the **incremental** command. This option cannot be added to a client option set on the Tivoli Storage Manager server.

Syntax

►—ABSolute—◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc incr -absolute "/Users/sparky/source/*.c"
```

Afmskipuncachedfiles

The `afmskipuncachedfiles` option specifies whether uncached and dirty files in General Parallel File System (GPFS) Active File Management file sets are processed for backup, archive, and migration operations.

GPFS Active File Management and *uncached* and *dirty* file states are explained in GPFS product information.

Running HSM on GPFS file systems that use Active File Management file sets is explained in Guidance for integrating IBM Spectrum Scale AFM with IBM Spectrum Protect .

If you back up, archive, or migrate files from a file system that contains Active File Management file sets, set `afmskipuncachedfiles=yes`.

Restriction: If Active File Management is running in Local Update (LU) mode, the **afmskipuncachedfiles** option in the cache file set must be set to **No**.

Supported Clients

This option is valid for backup-archive clients that run on AIX and Linux systems.

Options File

Place this option in the `dsm.sys` file before any server stanzas.

Syntax



Parameters

NO The Active File Management file state is ignored during backup, archive, and migration operations. Migration operations on uncached or dirty files fail and yield error message ANS9525E. Backup and archive operations on uncached files require Active File Management fetch operations. The fetch operations can cause significant network traffic between the Active File Management home and cache. This parameter is the default.

YES

Uncached or dirty files in Active File Management file sets are skipped during backup, archive, and migration processing.

Archmc

Use the **archmc** option with the **archive** command to specify the available management class for your policy domain to which you want to bind your archived files and directories.


When you archive a file, you can override the assigned management class using the **archmc** option on the **archive** command or by using the web client. Overriding the management class using the web client is equivalent to using the **archmc** option on the **archive** command.

If you do not use the **archmc** option, the server binds archived directories to the default management class. If the default management class has no archive copy group, the server binds archived directories to the management class with the shortest retention period.

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager API does not support this option.

Syntax



Parameters

managementclass

Specifies an available management class in the active policy set of your policy domain. This management class overrides the default management class and any include statements for the files and directories you are archiving.

Examples

Command line:

```

dsmc archive -archmc=ret2yrs /Users/van/Documents/budget.jan
dsmc archive -archmc=ret2yrs /home/plan/proj1/budget.jan

```

Archsymlinkasfile

The `archsymlinkasfile` option specifies whether the backup-archive client follows a symbolic link and archives the file or directory to which it points, or archives the symbolic link only. Use this option with the **archive** command.

Supported Clients

This option is valid for all UNIX clients except Mac OS X. The server can also define this option.

Options File

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

Syntax



Parameters

Yes

Specifies that the client follows a symbolic link and archives the associated file or directory. This is the default.

No Specifies that the client archives the symbolic link and not the associated file or directory.

Examples

Options file:

```
archsymlinkasfile no
```

Command line:

```
-archsyml=no
```

Asnodename

Use the `asnodename` option to allow agent nodes to back up or restore data on behalf of another node (the target node). This enables concurrent operations from multiple nodes to store data to the same target node and file space in parallel.

Your client node must be granted access to the target node by the Tivoli Storage Manager server administrative client **grant proxynode** command, and you must be a root user to use the `asnodename` option.

When the Tivoli Storage Manager administrator grants a node proxy authority, and you use the `asnodename` option to become that node, you can query and restore all files as if you had root authority.

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

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

A proxy operation uses the settings for the target node (such as **maxnummp** and **deduplication**) and schedules that are defined on the server. The Tivoli Storage Manager server node settings and schedules for the agent node are ignored.

For example, you can use the following command to back up shared data for file space stored under the node name MyCluster:

```
/cluster1/mydata  
dsmc incremental /Users -asnodename=MyCluster
```

You can also use the `asnodename` option to restore data under another node name on the server. You can only restore the data that you own.

The `asnodename` option differs from the `nodename` option as follows:

- When using the `nodename` option, you must enter the password for the node name you specify.
- When using the `asnodename` option, you must enter the password for your client agent node to access the data stored for the client target node.

Restrictions: You cannot use the `asnodename` option with `-fromnode` and you cannot perform NAS backup using `asnodename`. Also, `asnodename` can be used for clustered systems, although no specific cluster software is supported.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

Place this option in the `dsm.sys` file *within* a server stanza. You can set this option on the **General** tab of the Preferences editor.

Syntax

►► `ASNODENAME=` *—targetnode—* ◀◀

Parameters

targetnode

Specifies the node name on the Tivoli Storage Manager server under which you want to back up or restore data.

Examples

Options file:

```
asnodename mycluster
```

Command line:

```
-asnodename=mycluster
```

This option is not valid in interactive mode, but it can be defined in the options portion of a schedule definition.

Session settings and schedules for a proxy operation

A proxy operation occurs when an agent node uses the `asnodename` *target_node_name* option to complete operations on behalf of the specified target node.

A proxy operation uses the settings for the target node (such as **maxnummp**, **cloptset**, and **deduplication**) and schedules that are defined on the Tivoli Storage Manager server. The server node settings and schedules for the agent node are ignored.

The following considerations apply to proxy operations.

- All operations use the policy domain settings and constructs of the target node, even if the agent node belongs to a different domain. The policy domain settings and constructs of the agent node are ignored.
- The agent node authenticates to the Tivoli Storage Manager server by using the agent node's password.
- In order to run proxy operations, the agent node and target node must not be locked on the server.
- Proxy node relationships are not transitive. If a target node is itself defined as a proxy node for some other node, the agent node cannot be used to run operations on that other node unless the agent is also defined as a proxy node for that other node.

For example, assume the following proxy definitions among nodes TAURUS, SCORPIO, and GEMINI:

- TAURUS is a proxy node for SCORPIO.
- TAURUS is not a proxy node for GEMINI.
- SCORPIO is a proxy node for GEMINI.

The proxy definitions yield the following results:

- TAURUS can run operations on behalf of SCORPIO.
- SCORPIO can run operations on behalf of GEMINI.
- TAURUS cannot run operations on behalf of GEMINI.

Audit logging

Use the `audit logging` option to generate an audit log that contains an entry for each file that is processed during an incremental, selective, archive, restore, or retrieve operation.

The audit log can be configured to capture either a basic level of information or a more inclusive (full) level of information.

The basic level of the audit logging feature captures the information that is in the schedule log and it records information that a file has been backed up, archived, updated, restored, retrieved, expired, deleted, skipped or failed during an incremental backup, selective backup, archive, restore or retrieve operation. In addition, the basic level of audit logging captures the input command for commands run through the backup-archive command line or scheduler clients.

The full level of audit logging records an action for each file that is processed by the backup-archive client. In addition to all of the events recorded by the basic level of audit logging, the full level of audit logging records information for a file that has been excluded or not sent during a progressive incremental backup operation because the file had not changed.

The following is an example of the messages that are issued when the audit log is configured to capture the basic level of information:

```

04/21/07 15:25:05 ANS1650I Command:
    sel /home/spike/test/*
04/21/07 15:25:05 ANS1651I Backed Up:
    /home/spike/test/file.txt
04/21/07 15:25:05 ANS1652I Archived:
    /home/spike/test/file.txt
04/21/07 15:25:05 ANS1653I Updated:
    /home/spike/test/file.txt
04/21/07 15:25:05 ANS1654E Failed:
    /home/spike/test/file.txt
04/21/07 15:25:05 ANS1655I Restored:
    /home/spike/test/file.txt
04/21/07 15:25:05 ANS1656I Retrieved:
    /home/spike/test/file.txt
04/21/07 15:25:05 ANS1657I Expired:
    /home/spike/test/file.txt
04/21/07 15:25:05 ANS1658I Deleted:
    /home/spike/test/file.txt
04/21/07 15:25:05 ANS1659I Skipped:
    /home/spike/test/file.txt

```

The following messages can be issued when the audit log is configured to capture the full level of information (in addition to all messages issued for the basic level of audit logging):

```

04/21/07 15:25:05 ANS1660I Excluded:
    /home/spike/test/file.txt
04/21/07 15:25:05 ANS1661I Unchanged:
    /home/spike/test/file.txt

```

The audit log is not a substitute or a replacement for the standard error log (`dsmerror.log`) or for the schedule log (`dsmsched.log`). If an error occurs that prevents a file from being processed, a message indicating that an error has occurred is written to the audit log, but the message will not indicate the nature of the error. For problem diagnostics the standard error log must still be used.

The audit log entries only contain a time stamp and object name. There is no information to distinguish between files and directories or any information about the size of an object.

The Mac OS X backup-archive client creates the audit log as a Unicode (UTF-16) file.

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

The **auditlogging** command is supported with backup commands that interact with file-level objects such as **backup groups**.

The **auditlogging** command is not supported with backup commands which interact with image-level objects such as **backup image** or **restore image**. The **auditlogging** command is supported with backup commands that interact with file-level objects such as **backup groups**.

If you have enabled audit logging for an operation and there is a failure trying to write to the audit log (for example, the disk on which the audit log resides is out of space), the audit logging is disabled for the rest of the operation and the return code for the operation is set to 12, regardless of the outcome of the operation.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza.

Syntax



Parameters

off

Specifies that the audit logging facility is not engaged. This is the default.

basic

Specifies that the audit log captures a basic level of information.

full

Specifies that the audit log captures a more extensive level of information.

Examples

Run an incremental backup with audit logging enabled.

Command line:

```
dsmc i -auditlogging=basic
```

Back up a list of files using the maximum level of auditing, which enables a separate application, such as a Perl script, to verify the results.

Auditlogname

The `auditlogname` option specifies the path and file name where you want to store audit log information. This option applies when audit logging is enabled.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza.

Syntax



Parameters

filespec

Specifies the path and file name where you want the backup-archive client to store audit log information.

If you specify a file name only, the file is stored in your current directory. The default is the installation directory with a file name of dsmaudit.log. The dsmaudit.log file cannot be a symbolic link.

Examples

Run an incremental backup with audit logging enabled.

Sample output

The following is a sample execution and output file:

```
> dsmc inc /SMSVT/mfs1 -auditlogging=full
-auditlogname=/home/cliv3/audit.log
IBM Tivoli Storage Manager
Command Line Backup/Archive Client Interface
  Client Version 5, Release 5, Level 0.0
  Client date/time: 07/03/07 12:05:13
(c) Copyright by IBM Corporation and other(s)
    1990, 2007. All Rights Reserved.
```

```
Node Name: NAXOS_CLUSTER
Session established with server
  ODINHSMSERV: AIX-RS/6000
  Server Version 5, Release 4, Level 0.0
  Server date/time: 07/03/07 12:05:18
  Last access: 07/03/07 12:01:57
```

```
Incremental backup of volume '/SMSVT/mfs1'
Directory-->          4,096 /SMSVT
  /mfs1/ [Sent]
Normal File-->      32,768 /SMSVT
  /mfs1/test0 [Sent]
Normal File-->      32,768 /SMSVT
  /mfs1/test1 [Sent]
Normal File-->      32,768 /SMSVT
  /mfs1/test2 [Sent]
Normal File-->      32,768 /SMSVT
  /mfs1/test3 [Sent]
Normal File-->      32,768 /SMSVT
  /mfs1/test4 [Sent]
Normal File-->      32,768 /SMSVT
  /mfs1/test5 [Sent]
Normal File-->      32,768 /SMSVT
  /mfs1/test6 [Sent]
Normal File-->      32,768 /SMSVT
  /mfs1/test7 [Sent]
Normal File-->      32,768 /SMSVT
  /mfs1/test8 [Sent]
Normal File-->      32,768 /SMSVT
  /mfs1/test9 [Sent]
Successful incremental backup of '/SMSVT/mfs1'
```

```
Total number of objects inspected:      11
Total number of objects backed up:      11
Total number of objects updated:         0
Total number of objects rebound:        0
Total number of objects deleted:         0
Total number of objects expired:         0
Total number of objects failed:          0
Total number of bytes transferred:    320.31 KB
```

```

Data transfer time:                0.01 sec
Network data transfer rate:        17,141.84 KB/sec
Aggregate data transfer rate:      297.43 KB/sec
Objects compressed by:             0%
Elapsed processing time:           00:00:01

```

The following are the audit log contents:

```

07/03/07 12:05:14 ANS1650I Command:
      inc /SMSVT/mfs1
07/03/07 12:05:15 ANS1651I Backed Up:
      /SMSVT/mfs1/
07/03/07 12:05:15 ANS1651I Backed Up:
      /SMSVT/mfs1/test0
07/03/07 12:05:15 ANS1651I Backed Up:
      /SMSVT/mfs1/test1
07/03/07 12:05:15 ANS1651I Backed Up:
      /SMSVT/mfs1/test2
07/03/07 12:05:15 ANS1651I Backed Up:
      /SMSVT/mfs1/test3
07/03/07 12:05:15 ANS1651I Backed Up:
      /SMSVT/mfs1/test4
07/03/07 12:05:15 ANS1651I Backed Up:
      /SMSVT/mfs1/test5
07/03/07 12:05:15 ANS1651I Backed Up:
      /SMSVT/mfs1/test6
07/03/07 12:05:15 ANS1651I Backed Up:
      /SMSVT/mfs1/test7
07/03/07 12:05:15 ANS1651I Backed Up:
      /SMSVT/mfs1/test8
07/03/07 12:05:15 ANS1651I Backed Up:
      /SMSVT/mfs1/test9

```

Related information

For more information about the audit logging facility refer to “Audit logging” on page 321.

Autodeploy

Use the autodeploy option to enable or disable an automatic deployment of the client if a restart is required.

Supported Clients

This option is valid for AIX, HP-UX, Linux, Mac, and Solaris clients.

Options File

You can set this option by including it in your client options file. You can also set in using the Java GUI by clicking **Edit > Client Preferences** and selecting the appropriate option on the **General** tab.

Syntax



Parameters

Yes

Specifies that the client is automatically deployed from the server. Yes is the default.

No Specifies that the client is not automatically deployed from the server.

Examples

Options file:

autodeploy no

Command line:

Does not apply.

Important: Use schedmode prompted with the autodeploy option, to enable the scheduler to process the client deployment schedule immediately.

Related concepts:

“Automatic backup-archive client deployment” on page 2

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, the Tivoli Storage Manager server generates a unique name by appending _OLD to the file space name you specify in the current operation. For example, the server renames the file space Jaguar to Jaguar_OLD. If the new file space name is too long, the suffix replaces the last characters of the file space name. For example, the mylongfilesystemname file space name is renamed to:

mylongfilesystem_OLD

If the new file space name already exists on the server, the server renames the new file space Jaguar_OLDx, where x is a unique number.

The server creates new Unicode-enabled file spaces that contain only the data specified in the current operation. For example, assume that Jaguar is the name of your startup disk and you archive all of the .log files in the /Users/user5/Documents directory. Before the archive takes place, the server renames the file space to Jaguar_OLD. The archive places the data specified in the current operation into the Unicode-enabled file space named Jaguar. The new Unicode-enabled file space now contains only the /Users/user5/logs directory and the *.log files specified in the operation. The server stores all subsequent full and partial incremental, selective backup, and archive data in the new Unicode-enabled file spaces.

For example, assume that Jaguar is the name of your startup disk and you archive all of the .log files in the /Users/user5/Documents directory. Before the archive takes place, the server renames the file space to Jaguar_OLD. The archive places the data specified in the current operation into the Unicode-enabled file space named Jaguar. The new Unicode-enabled file space now contains only the /Users/user5/logs directory and the *.log files specified in the operation. All

subsequent full and partial incremental, selective backup, and archive data are stored 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 the files were backed up. A workstation running in a different locale might be unable to view or restore from these file spaces. Unicode-enabled file spaces that are backed up in one locale are visible in all other locales, provided that the workstation has the proper fonts installed.

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


Supported Clients

This option is valid for Mac OS X only. The server can define the `autofsrename` option and override the `autofsrename` setting on the client. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **General** tab, **Rename non-Unicode files** during backup/archive drop-down list box of the Preferences editor.

Syntax



Parameters

Yes

Specifies that the Tivoli Storage Managerserver automatically renames all file spaces that are not Unicode-enabled in the current backup or archive operation.

No Specifies that the server does not rename file spaces that are not Unicode-enabled in the current backup or archive operation.

Prompt

Specifies that you are prompted 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
```

Automount

The `automount` option adds an automounted file system into the domain by mounting it. Use this option with the `domain` option.

Use this option to specify all automounted file systems that the backup-archive client tries to mount at the following points in time:

- When the client starts
- When the backup is started
- When the the client has reached an automounted file system during backup

Mount the file system before the client does a backup of that file system. If the file system is always mounted before the backup is done, it is unnecessary to explicitly specify an automounted file system in the `automount` option. However, add this file system in the `automount` option to ensure that the file system has been mounted at all the points in time mentioned previously. The automounted file systems are remounted if they have gone offline in the meantime during a backup.

Supported Clients

This option is valid for all UNIX platforms except Mac OS X. The Tivoli Storage Manager API does not support this option.

Options File

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

Syntax



Parameters

*filespace**name*

Specifies one or more fully qualified automounted file systems that are mounted and added into the domain.

Examples

Options file:

automount /home/Fred /home/Sam

Command line:

Does not apply.

Related information

See “Domain” on page 361 for more information about working with automounted file systems and the domain option.

Backmc

The backmc option specifies the management class to apply to the **backup fastback** command for retention purposes.

Use the backmc option with the **backup fastback** command.

If you back up an object more than once and specify a different management class for each backup, all backup versions of the object are rebound to the last management class specified.

Supported Clients

This option is valid for Linux x86_64 clients.

Options File

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

Syntax

►►—BACKMc=—*management_class_name*————►►

Parameters

management_class_name

Specifies the management class name.

Examples

Command line:

dsmc backup fastback -fbpolicyname=policy1 -fbserver=server1
-backmc=ret2yrs

Backupsetname

The backupsetname option specifies the name of a backup set from the Tivoli Storage Manager server.

You can use backupsetname option with the following commands:

- **query backup**
- **query filesystem**
- **query image**
- **restore image**

Note: The following commands take backupsetname as a positional parameter. The backupsetname positional parameter behaves differently from the backupsetname option. See the command explanations for a discussion of how the backupsetname positional parameter affects each of these commands:

query backupset
restore
restore backupset

Supported Clients

This option is valid for all UNIX and Linux clients. The Tivoli Storage Manager API does not support this option.

Options File

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

Syntax

►►—BACKUPSETName—*backupsetname*—————►►

Parameters

backupsetname

Specifies the name of a backup set from the Tivoli Storage Manager server. You cannot use wildcards.

Examples

Command line:

```
dsmc query backup /Volumes/bkSets/file.1
-backupsetname=YEAR_END_ACCOUNTING.12345678
dsmc query backup /usr/projects -subdir=yes
-backupsetname=YEAR_END_ACCOUNTING.12345678
dsmc restore image /home/proj
-backupsetname=ACCOUNTING_2007.12345678
dsmc query image -backupsetname=WEEKLY_BSET.21435678
```

Related information

“Restore data from a backup set” on page 224

Basesnapshotname

The basesnapshotname option specifies the snapshot to use as the base snapshot, when you perform a snapshot differential (snapdiff) backup of a NetApp filer volume. If you specify this option, you must also use the snapdiff option or an error occurs. If basesnapshotname is not specified, the useexistingbase option selects the most recent snapshot on the filer volume as the base snapshot.

If the specified snapshot cannot be found, an error is reported and the backup operation fails.

Supported Clients

This option can be used with supported x86_64 Linux and AIX clients.

Options File

This option can be specified in the client options file or on the command line.

Syntax

►►—BASESNAPSHOTName— —*snapshot_name*—►►

Parameters

snapshot_name

Specifies the name of an existing snapshot to use as the base snapshot. The name specified can be a snapshot name, such as `vol1_snap`, or it can be the name of a scheduled NetApp backup that has a name like `nightly.x`, where `x` is the sequence number (where `nightly.0` is the oldest snapshot).

You can also use a pattern with wildcard characters to select a snapshot. The wildcard characters can be either of the following:

* An asterisk (*) matches any character.

? A question mark (?) matches a single character.

The wildcards are useful if your snapshots follow a pattern, such as including the date or data and time as part of the snapshot name. For example, a snapshot created on November 12 2012 at 11:10:00 AM could be saved as `UserDataVol_121103111000_snapshot`. The most recent snapshot that matches the pattern is selected as the existing base. For example, if there are two saved snapshots (`UserDataVol_121103111000_snapshot` and `UserDataVol_121103231000_snapshot`), the `UserDataVol_121103231000_snapshot` is selected because it is 12 hours newer than the other snapshot.

```
-basesnapshotname="UserDataVol_*_snapshot"
```

Question marks work well for scheduled backups that follow a consistent name pattern. This syntax selects the latest “nightly” backup as the snapshot to use as the existing base.

```
-basenameshotname="nightly.?"
```

Examples

Options file:

```
basesnapshotname nightly.?
basesnapshotname volum_base_snap
```

Command line:

```
dsmc incr \\DRFiler\UserDataVol_Mirror_Share -snapdiff
-useexistingbase -basesnapshotname="nightly.?"
```

Related information

Useexistingbase

Cadlistenonport

The cadlistenonport option specifies whether to open a listening port for the client acceptor.

When a listening port is open, it can accept any inbound connections. However, the port is not used when the client acceptor manages only the scheduler and the scheduler runs in polling mode. You can use this option to prevent the acceptor from opening the unused port.

The default setting for this option is yes. Use cadlistenonport no only when managedservices schedule and schedmode polling are used.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza.

Syntax



Parameters

Yes

Specifies that the client acceptor opens a listening port. This parameter is the default.

No

Specifies that the client acceptor does not open a listening port. Use this setting when you use the client acceptor only to manage the scheduler in polling mode.

This setting effectively disables other client features that depend on the client acceptor, such as web client backup and restore operations, IBM Tivoli Storage Manager for Virtual Environments: Data Protection for VMware vSphere GUI operations, and IBM Tivoli Storage FlashCopy Manager backup and restore operations.

Example

Options file:

```
cadlistenonport no
```

Command line:

Does not apply.

Related reference:

"Managedservices" on page 448

"Schedmode" on page 509

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 copy serialization, an attribute in a management class copy group, is shared static or shared dynamic.

With shared static serialization, if a file is open during an operation, the operation repeats the number of times that you specify. If the file is open during each attempt, the operation does not complete.

With shared dynamic serialization, if a file is open during an operation, the operation repeats the number of times that you specify. The backup or archive occurs during the last attempt whether the file is open or not.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **Backup** tab, **Number of retries if file is in use** field of the Preferences editor.

Syntax

►►—CHAngingretries— *numberretries* —◄◄

Parameters

numberretries

Specifies the number of times a backup or archive operation is attempted if the file is in use. The range of values is zero through 4; the default is 4.

Examples

Options file:

```
changingretries 3
```

Command line:

```
-cha=3
```

Class

The `class` option specifies whether to display a list of NAS or client objects when using the **delete filesystem**, **query backup**, and **query filesystem** commands.

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 only for AIX, Linux, and Oracle Solaris clients. The Tivoli Storage Manager API does not support this option.

Options File

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

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

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

Command line:

```
q backup -nasnodename=nodename -class=nas
```

Collocatebyfilespec

Use the `collocatebyfilespec` option to specify whether the backup-archive client uses only one server session to send objects generated from one file specification.

Setting the `collocatebyfilespec` option to `yes` attempts to eliminate interspersing of files from different file specifications, by limiting the client to one server session per file specification. Therefore, if you store the data to tape, files for each file specification are stored together on one tape (unless another tape is required for more capacity).

Considerations:

- Use the `collocatebyfilespec` option only if the storage pool is going directly to tape. If you use this option going to a disk storage pool, you could affect some load balancing, and therefore, performance.

Supported Clients

This option is valid for all UNIX and Linux clients. The server can also define this option.

Options File

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

Syntax



Parameters

Yes

Specifies that you want the 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 client can (depending on the execution dynamics and on the setting of the resourceutilization option of 3 or higher) use more than one server session to send the files from one file specification. This is the default.

Backup performance might increase as a result. If the files are backed up to tape, files are stored on multiple tapes. Generally, the files specified in the file specification are still contiguous.

Examples

Options file:

```
collocatebyfilespec yes
```

Command line:

```
-collocatebyfilespec=yes
```

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

Commmethod

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


Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza. You can set this option on the **Communication** tab of the Preferences editor.

Syntax



Parameters

TCPip

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

V6Tcpip

Indicates that either TCP/IP V4 or V6 should be used, depending on the system configuration and the results of a domain name service lookup. A valid DNS environment must be available.

SHAREdmem

Use the shared memory communication method when the client and server are running on the same system. This provides better performance than the TCP/IP protocol. This communication method is valid only for AIX, HP-UX, Linux, and Solaris clients.

When specifying this communication method on AIX, the client can be logged in as root or non-root, as long as the server is running as root. If the server is not running as root, the user ID running the client must match the user ID running the server.

Important: When using `commethod sharedmem` on Linux, you might receive error message: ANR8294W Shared Memory Session unable to initialize on the server or storage agent console. By default, Linux is not set up with sufficient system resources to create the message queues. You must increase the kernel parameter, `MSGMNI`, to 128 (the default is 16). You can modify this parameter by performing the following command:

```
echo 128 > /proc/sys/kernel/msgmni
```

To enable this parameter to remain persistent after rebooting the system, you can instead add the following line to the file `/etc/sysctl.conf`, then reboot the system:

```
kernel.msgmni=128
```

To view the current ipc settings, run this command:

```
ipcs -l
```

Now look at the max queues system wide value. The default is 16.

Examples

Options file:

```
commethod tcpip
```

Use only TCP/IP V4

```
commethod V6Tcpip
```

Use both TCP/IP V4 and V6, depending on how the system is configured, and the results of a domain name service lookup.

Note: The `dsmc schedule` command cannot be used when both `SCHEDMODE` prompt and `commethod V6Tcpip` are specified.

Command line:

```
-commm=tcpip
```

```
-commm=V6Tcpip
```

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

Commrestartduration

The `commrestartduration` option specifies the maximum number of minutes you want the client to try to reconnect to a Tivoli Storage Manager server after a communication error occurs.

Note: A scheduled event continues if the client reconnects with the server before the `commrestartduration` value elapses, even if the startup window of the event has elapsed.

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

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **Communication** tab, **Common Options** section of the Preferences editor.

Syntax

►►—`COMMRESTARTDuration— minutes`—►►

Parameters

minutes

The maximum number of minutes you want the client to attempt to reconnect with a server after a communication failure occurs. The range of values is zero through 9999; the default is 60.

Examples

Options file:

```
commrestartduration 90
```

Command line:

Does not apply.

Commrestartinterval

The `commrestartinterval` option specifies the number of seconds you want the client to wait between attempts to reconnect to a Tivoli Storage Manager server after a communication error occurs.

Note: Use this option only when `commrestartduration` is a value greater than zero.

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

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza. You can set this option on the **Communication** tab, **Common Options** section of the Preferences editor.

Syntax

►►—COMMRESTARTInterval— *seconds* —————►►

Parameters

seconds

The number of seconds you want the client to wait between attempts to reconnect with a server after a communication failure occurs. The range of values is zero through 65535; the default is 15.

Examples

Options file:

commrestartinterval 30

Command line:

Does not apply.

Compressalways

The compressalways option specifies whether to continue compressing an object if it grows during compression.

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

The compressalways option is ignored when client-side deduplication is enabled.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

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

Syntax

►►—COMPRESSAlways—

Yes
No

 —————►►

Parameters

Yes

File compression continues even if the file grows as a result of compression. This is the default.

No Backup-archive client objects are resent uncompressed if they grow during compression. API behavior depends on the application. Application backups might fail.

Examples

Options file:

```
compressalways yes
```

Command line:

```
-compressa=no
```

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

Compression

The compression option compresses files before you send them to the server.

Compressing your files reduces data storage for backup versions and archive copies of your files. It can, however, affect Tivoli Storage Manager throughput. A fast processor on a slow network connection benefits from compression, but a slow processor on a fast network connection does not.

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

The **backup image** command uses the compression option value specified in the dsm.sys file. This option is valid on the initial command line and in interactive mode. The server can also define this option which overrides the client value.

The backup-archive client backs up a sparse file as a regular file if client compression is off. Set compression yes to enable file compression when backing up sparse files to minimize network transaction time and maximize server storage space.

If you set compressalways yes, compression continues even if the file size increases. To stop compression if the file size grows, and resend the file uncompressed, set compressalways no.

If you set compression yes, you can control compression processing in the following ways:

- Use the exclude.compression option in your client system-options file (dsm.sys) to exclude specific files or groups of files from compression processing.
- Use the include.compression option in your client system-options file (dsm.sys) to include files within a broad group of excluded files for compression processing.

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

The type of compression that the client uses is determined by the combination of compression and client-side data deduplication that is used during backup or archive processing. The following types of compression are used:

LZ4 A faster and more efficient compression method that the client uses when client-deduplicated data is sent to an LZ4-compatible container storage

pool on the Tivoli Storage Manager server. The server must be at version 7.1.5 or later, and must use container storage pools. Client-side LZ4 compression is used only when client-side data deduplication is enabled.

- LZW** A traditional type of compression that the client uses in any of the following situations:
- Client-deduplicated data is sent to traditional (non-container) storage pools on the server.
 - The client data does not undergo client-side data deduplication. (Does not apply to Data Protection for VMware and Data Protection for Microsoft Hyper-V, in which only client-deduplicated data can be compressed.)
 - The client data undergoes only traditional server-side data deduplication. (Does not apply to Data Protection for VMware and Data Protection for Microsoft Hyper-V, in which only client-deduplicated data can be compressed.)
- None** The object is not compressed by the client. The object is not compressed because the compression option is set to *no*, or the option is not specified during backup or archive processing. Although the object is not compressed by the client, it might be compressed by the server.

You do not need to set the compression type. It is determined by the backup-archive client at the time of backup or archive processing.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

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

Syntax



Parameters

No Files are not compressed before they are sent to the server. This is the default.

Yes

Files are compressed before they are sent to the server.

Examples

Options file:

`compression yes`

Command line:

`-compressi=no`

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

Related reference:

“Deduplication” on page 349

“Exclude options” on page 389

“Include options” on page 421

Console

Use the `console` option with the **query systeminfo** command to output information to the console.

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

Note: The **query systeminfo** command is intended primarily as an aid for IBM support to assist in diagnosing problems, although users who are familiar with the concepts addressed by this information might also find it useful. If you use the `console` option, no special formatting of the output is performed to accommodate screen height or width. Therefore, the console output might be difficult to read due to length and line-wrapping. In this case, use the `filename` option with the **query systeminfo** command to allow the output to be written to a file that can subsequently be submitted to IBM support.

Supported Clients

This option is valid for all clients.

Syntax

►►—CONSOLE—►►

Parameters

There are no parameters for this option.

Examples

Command line:

```
query systeminfo dsmoptfile errorlog -console
```

Related information

“Filename” on page 405

Createnewbase

The createnewbase option creates a base snapshot and uses it as a source to run a full incremental backup.

Some files might not be backed up when the snapshot difference incremental backup command is run. If the files are skipped, you can run a snapshot difference incremental backup with the createnewbase option to back up these files. See “Snapdiff” on page 519 for a list of reasons why a file might not be backed up when the snapshot difference command is run.

One reason that a file can be skipped during backup processing is because the file name is not supported by NetApp Data ONTAP. NetApp Data ONTAP Versions 8.0 and versions lower than 7.3.3 only support file names that are within the 7 bit ASCII character set. NetApp Data ONTAP Version 7.3.3 and versions greater than 8.0.0 support Unicode file names. If you upgraded NetApp Data ONTAP from a version that does not support Unicode file names to a version that does support Unicode files names, run a full incremental backup with the createnewbase=migrate option.

Supported Clients

This option is valid for the following clients:

- AIX 64-bit clients
- Linux x86_64 clients

Enter the createnewbase option on the command line. Specify this option with the snapdiff option.

Syntax



Parameters

No Specifies that a snapshot difference incremental is run. If the backup-archive client detects that the NetApp Data ONTAP file server has been migrated from a version that does not support Unicode file names to a file server that does, a warning message is recorded to the error log and the Tivoli Storage Manager server activity log. The warning message indicates that you must run a full incremental backup and logs a return code of 8 even if the operation completed successfully.

This parameter is the default value.

Yes

Specifies that a full incremental is run by creating a new base snapshot and is using it to run a scan-based incremental backup. Use this option to back up any file changes that might not have been detected by the snapshot difference API.

If the operation finished successfully, the command ends with a return code of 0.

Do not set `createnewbase=yes` for any schedule that runs a daily snapshot difference backup. Instead, create a separate, monthly schedule that has the `createnewbase=yes` option.

IGNore

Specifies that a snapshot difference incremental backup is run when the backup-archive client detects that the NetApp Data ONTAP file server was upgraded to support Unicode file names.

The ignore option is different from the no parameter because the ignore option suppresses the warning message. Instead, an informational message is recorded in the error log and the Tivoli Storage Manager activity log that informs you to run a full incremental backup.

If the command finishes successfully, it returns a code of 0.

Use the ignore option if you have upgraded the NetApp Data ONTAP file server to support Unicode but you have not yet run a full incremental backup. This option is used only when the backup-archive client has detected that the file server was migrated and a full incremental has not yet been run. The option is ignored for all other times.

MIGRate

Specifies that if the NetApp Data ONTAP file server was upgraded to a version that supports Unicode file names, a base snapshot is taken and a scan-based incremental backup is run. The migrate option is different from the yes option because the migrate option creates a base snapshot only when the client detects that the NetApp Data ONTAP file server version was updated. The yes option creates a base snapshot every time the command is run.

After the incremental backup finishes, no additional migration-related messages are recorded to the error log or the Tivoli Storage Manager server activity log. When the operation finishes, the command ends with a return code of 0.

Use the migrate option if you have upgraded the NetApp Data ONTAP file server to support Unicode but you have not yet run a full incremental backup. The migrate option is ignored if the NetApp Data ONTAP file server has not been upgraded.

Examples

Command line:

```
dsmc incremental -snapdiff -createnewbase=yes /net/home1
```

Related tasks:

“Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 106

Related reference:

“Snapdiff” on page 519

Datacenter

Specifies the target location of the data center that will contain the restored machine data.

Use this option on vm restore commands.

If folders are used within virtual center to organize datacenters, then the folder name needs to be included in the datacenter specification, separated by a slash.

If you are restoring through a ESX server rather than a virtual center, the `-datacenter=ha-datacenter` option should be used.

The default target location is the datacenter which the virtual machine was stored at the time of backup.

Examples

Restore a virtual machine to USEast datacenter which is organized under a folder named Production in the virtual center.

```
dsmc restore vm my_vm -datacenter=Production/USEast
```

Restore a virtual machine backup taken from a virtual center, but using a ESX server at the time of restore.

```
restore vm my_vm -datacenter=ha-datacenter
```

Restore the virtual machine into the USWest datacenter.

```
restore vm my_vm -datacenter=USWest
```

Datastore

Specifies the datastore target to be used during VMware restore operation.

Example

Restore the virtual machine to a datacenter named ds8k_prod1

```
restore vm my_vm -datacenter=ds8k_prod1
```

Dateformat

The `dateformat` option specifies the format you want to use to display or enter dates.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time you start the client. Consult the documentation on your local system for details about setting up your locale definition.

Note:

1. The `dateformat` option does not affect the web client. The web client uses the date format for the locale that the browser is running in. If the browser is not running in a locale that is supported, 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, the client 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, the client removes all entries in the error log with a different date when pruning the log. When changing the date format, copy the schedule log and error log if you want to preserve log entries that contain a different date format.

You can use the `dateformat` option with the following commands.

- **delete archive**

- delete backup
- expire
- query archive
- query backup
- query filespace
- query image
- restore
- restore image
- restore nas
- retrieve
- set event

When you include the `dateformat` option with a command, it must precede the `fromdate`, `pitdate`, and `todate` options.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client user-options file (`dsm.opt`). You can set this option on the **Regional Settings** tab, **Date Format** drop-down list of the Preferences editor.

Syntax

►►—`DATEformat`— *—format_number—*—————►►

Parameters

format_number

Displays the date using one of the following formats. Select the number that corresponds to the date format you want to use:

- 0** Use the locale-specified date format (does not apply to Mac OS X).

For AIX, HP-UX, and Solaris: This is the default if the locale-specified date format consists of digits and separator characters.

- 1** MM/DD/YYYY

For AIX, HP-UX, and Solaris: This is the default if the locale-specified date format consists of anything but digits and separator characters.

This is the default for the following available translations:

- US English
- Chinese (Traditional)
- Korean

- 2** DD-MM-YYYY

This is the default for the following available translations:

- Brazilian Portuguese
- Italian

- 3** YYYY-MM-DD

This is the default for the following available translations:

- Japanese
- Chinese (Simplified)
- Polish

- 4** DD.MM.YYYY

This is the default for the following available translations:

- German
- French
- Spanish
- Czech
- Russian

5 YYYY.MM.DD

This is the default for the following available translations:

- Hungarian

6 YYYY/MM/DD

7 DD/MM/YYYY

Examples

Options file:

```
dateformat 3
```

Command line:

```
-date=3
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the `dsm.opt` file unless overridden by the initial command line or by an option forced by the server.

Additional considerations for specifying time and date formats

The date or time format you specify with this option must be used when using options that take date and time as input. Examples are: `totime`, `fromtime`, `today`, `fromdate`, and `pittime`.

For example, if you specify the `timeformat` option as `TIMEFORMAT 4`, the value that you provide on the `fromtime` or `totime` option must be specified as a time such as `12:24:00pm`. Specifying `13:24:00` would not be valid because `TIMEFORMAT 4` requires an hour integer that is 12 or less. If you want to specify up to 24 hour values on an option, and if you want to use commas as separators, you must specify `TIMEFORMAT 2`.

Configuring date and time formats in the system locale configuration file

You can specify date and time formats by configuring them in your system's locale file. If you specify time and date formats in the locale file, they must be defined using a subset of number-producing format specifiers that are supported by the C language `strftime()` function. You can use the following specifiers to set date and time formats in configuration settings for your locale.

Date specifiers

- `%Y` - the year, in four digits. For example, 2011.
- `%y` - the year, last two digits only. For example, 11 not 2011.
- `%m` - the month, as a decimal number (1-12).
- `%d` - the day of the month (1-31).

In the date specifiers, you can specify only one year specifier. Do not specify both `%Y` and `%y`. The `E` modifier (a capital `E`) can precede the date

specifiers to produce the locale's alternative form for the year, month, or day. If no alternative form exists, the E modifier is ignored. Separate the specifiers with a single 7-bit ASCII character. Commonly used separators include colons (:), commas (,), periods (.), hyphens (-), or forward slash (/) characters. Do not use multibyte characters as separators.

Time specifiers

- %H - the hour, in 24-hour form (00-23).
- %I - the hour, in 12-hour form (00-12).
- %M - minutes after the hour (00-59).
- %S - seconds after the minute (00-59)
- %p - adds the AM (before noon) or PM (after noon) indicator.

In the time specifiers, you can specify only one hour specifier. Do not specify both %I and %H.

The O modifier (a capital O) can precede the time specifiers to produce the locale's alternative form for the hour, minutes, or seconds. The O modifier cannot precede the %p specifier. Separate the specifiers with a single 7-bit ASCII character. Commonly used separators include colons (:), commas (,), or periods (.). Do not use multibyte characters as separators. Do not specify a separator between the %p specifier and the separator that precedes or follows it.

Time format examples, configured in the locale settings

To set a particular time format, edit the configuration file for your locale and modify the t_fmt line to support your needs. Whatever time format you select applies both to output and to input. After the locale configuration file has been edited, the **localedef** command must be run to create the final locale file.

Table 70. Sample time format settings in the locale configuration (t_fmt line)

Example	Result
"%H:%M:%S"	Displays time in the form <i>hh:mm:ss</i> with <i>hh</i> ranging from 0 through 23.
"%H,%M,%S"	Displays time in the form <i>hh,mm,ss</i> with <i>hh</i> ranging from 0 through 23.
"%I,%M,13p"	Displays time in the form <i>hh,mm,ssA/P</i> with <i>hh</i> ranging from 1 through 12 and <i>A/P</i> is the local abbreviation for ante-meridian (AM in English) or post-meridian (PM in English).

Date format examples, configured in the locale settings

To set a particular date format, edit the configuration file and modify the d_fmt line as needed to support your needs. Whatever date format you select applies both to output and to input.

Table 71. Sample date format settings in the locale configuration (d_fmt line)

Example	Result
"%m/%d/%y"	Displays the date in the form <i>MM/DD/YY</i> .
"%d.%m.%Y"	Displays the date in the form <i>DD.MM.YYYY</i> .

Dedupcachepath

Use the dedupcachepath option to specify the location where the client-side data deduplication cache database is created.

This option is ignored if the enablededupcache=no option is set during backup or archive processing.

Supported Clients

This option is valid for all clients. This option is also valid for the Tivoli Storage Manager API.

Options File

Place this option in the system-options file (dsm.sys). You can set this option on the **Deduplication Cache Location** field of the Preferences editor. The option can be set in the client option set on the Tivoli Storage Manager server.

Syntax

►►—DEDUPCACHEPath—*path*—————►►

Parameters

path

Specifies the location where the client-side data deduplication cache database is created if the enablededupcache option is set to yes. The default location is to create the data deduplication cache file in the backup-archive client or API installation directory.

Examples

Options file:

dedupcachepath /volumes/temp

Command line:

Does not apply.

Related reference:

“Enablededupcache” on page 378

Dedupcachesize

Use the dedupcachesize option to determine the maximum size of the data deduplication cache file. When the cache file reaches its maximum size, the contents of the cache are deleted and new entries are added.

Supported Clients

This option is valid for all clients. This option is also valid for the Tivoli Storage Manager API.

Options File

Place this option in the system-options file (dsm.sys). You can set this option on the **Deduplication > Deduplication Cache > Maximum Size** field of the

Preferences editor. The option can be set in the client option set on the Tivoli Storage Manager server.

Syntax

►►—DEDUPCACHESize—*dedupcachesize*—◀◀

Parameters

dedupcachesize

Specifies the maximum size, in megabytes, of the data deduplication cache file. The range of values is 1 - 2048; the default is 256.

Examples

Options file:

dedupcachesize 1024

Command line:

Does not apply.

Related reference:

"Deduplication"

Deduplication

Use the deduplication option to specify whether to enable redundant client-side data elimination when data is transferred to the Tivoli Storage Manager server during backup and archive processing.

Data deduplication is disabled if the enablelanfree or subfilebackup options are set. Backup-archive client encrypted files are excluded from client-side data deduplication. Files from encrypted file systems are also excluded.

To support client-side data deduplication, the following criteria must be met:

- Client-side data deduplication for the node is enabled on the server.
- The storage pool destination for the data must be a storage pool that is enabled for data deduplication. The storage pool must have a device type of "file".
- A file can be excluded from client-side data deduplication processing (by default all files are included).
- The server can limit the maximum transaction size for data deduplication by setting the CLIENTDEDUPTXNLIMIT option on the server. For more information about the option, refer to the Tivoli Storage Manager server documentation.
- The file size must be larger than 2 KB.

Supported Clients

This option is valid for all clients; it can also be used by the Tivoli Storage Manager API.

Options File

Place this option in the system-options file (dsm.sys) within a server stanza. You can set this option by selecting the **Deduplication > Enable Deduplication** check box of the Preferences editor. The option can be set in the client option set on the Tivoli Storage Manager server.

Syntax



Parameters

No Specifies that you do not want to enable client-side data deduplication for backup and archive processing. No is the default.

Yes
Specifies that you want to enable client-side data deduplication for backup and archive processing.

Examples

Options file:

```
deduplication yes
```

Command line:

```
-deduplication=yes
```

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

Related reference:

“Include options” on page 421

“Exclude options” on page 389

Defaultserver

Use the `defaultserver` option to specify the name of the Tivoli Storage Manager server to contact for backup-archive services if more than one server is defined in the `dsm.sys` file.

By default, the backup-archive contacts the server defined by the first stanza in the `dsm.sys` file. This option is only used if the `servername` option is not specified in the client user-options file (`dsm.opt`).

If you have the HSM client installed on your workstation, and you do not specify a migration server with the `migrateserver` option, use this option to specify the server to which you want to migrate files. For more information, see the Tivoli Storage Manager for Space Management product documentation on IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SSSR2R/welcome>.

Supported Clients

This option is valid for all UNIX clients.

Options File

Place this option *at the beginning* of the `dsm.sys` file *before* any server stanzas.

Syntax

►►—DEFAULTServer— —servername—►►

Parameters

servername

Specifies the name of the default server to which you back up or archive files. The server to which files are migrated from your local file systems can also be specified with this option.

Examples

Options file:

defaults server_a

Command line:

Does not apply.

Deletefiles

Use the deletefiles option with the **archive** command to delete files from your workstation after you archive them.

You can also use this option with the **restore image** command and the incremental option to delete files from the restored image if they were deleted after the image was created. Deletion of files is performed correctly if the backup copy group of the Tivoli Storage Manager server has enough versions for existing and deleted files.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—DELetefiles—►►

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc archive "/Users/dgordon/Documents/*.c" -deletefiles
dsmc archive "/home/foo/*.c" -deletefiles
dsmc restore image /local/data -incremental -deletefiles
```

Description

The description option assigns or specifies a description for files when performing archive, delete archive, retrieve, query archive, or query backupset.

For example, if you want to archive a file named budget.jan and assign to it the description "2002 Budget for Proj 1", you would enter:

```
dsmc archive -des="2003 Budget for Proj 1" /home/plan/  
proj1/budget.jan
```

Note:

1. The maximum length of a description is 254 characters.
2. Enclose the value in quotation marks (" ") if the option value that you enter contains a blank space.

Use the description option with the following commands:

- **archive**
- **delete archive**
- **query archive**
- **query backupset**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—DEscription =— —*description*—————►►

Parameters

description

Assigns a description to the file you are archiving. If you do not specify a description with the **archive** command, the default is Archive Date:x, where x is the current system date. Note that the date is always 10 characters long. If your date format uses a two digit year, there are two blank spaces at the end of the date. For example, a default description using a four-digit year might be "Archive Date: 2002/05/03", and the same default with a two-digit year might be "Archive Date: 02/05/03 " (note the two spaces at the end). When retrieving files using the two-digit year description, you can enter the -description option string in either of the following ways:

```
-description="ArchiveDate: 02/05/03 "  
or  
-description="ArchiveDate: 02/05/03"
```

If you use the **archive** command to archive more than one file, the description you enter applies to each file. For example, to archive a group of files and assign the same description, *Project X*, to each file, you would enter:

```
dsmc archive -description="Project X" "/Users/van/Documents/*.x"  
dsmc archive -description="Project X" "/home/allproj/*"
```

You can then use the description to retrieve all of the files.

Examples

Command line:

```
dsmc archive "/Users/van/Documents/*.prj" -des="2003 Budget for Proj 1"  
dsmc archive "/home/foo/*.prj" -des="2003 Budget for Proj 1"  
dsmc query backupset -loc=server -descr="My Laptop"
```

Detail

Use the `detail` option to display management class, file space, backup, archive information, and additional information, depending on the command with which it is used.

Use the `detail` option with the **query mgmtclass** command to display detailed information about each management class in your active policy set. If you do not use the `detail` option, only the management class name and a brief description are displayed on the screen. If you specify the `detail` option, information about attributes in each copy group contained in each management class is displayed on the screen. A management class can contain a backup copy group, an archive copy group, both, or neither.

A Unicode-enabled file space might not display correctly if the server cannot display the Unicode name. In this case, use the file space identifier (fsID) of the file space to identify these file spaces on the server. Use the `detail` option with the **delete filespace** and **query filespace** commands to determine the fsID of a file space. The fsID also appears in the file information dialog in the backup-archive client and web client GUIs.

Use the `detail` option with the **query backup** and **query archive** commands to display these attributes of the file that you specify:

- Last modification date
- Last access date
- Compression
- Encryption type
- Client-side data deduplication
- Whether the HSM client migrated or premigrated the file

Use the `detail` with the **query vm** command to display the following statistics:

- The average number of Tivoli Storage Manager objects that are needed to describe a single megablock, across all megablocks in a backup.
- The average number of Tivoli Storage Manager objects that are needed to describe a single megablock, for all megablocks in a filespace.
- The ratio of the amount of data, reported by Change Block Tracking, versus the amount of data that was actually backed up, in a specific backup.
- The ratio of the amount of data, reported by Change Block Tracking, versus the amount of data that was actually backed up, for all backups in this filespace.
- The number of backups that were created since the last full backup was created from the production disks.

The values returned on **query vm** can help you fine tune the heuristics (see the `Mbobjrefreshthresh` and `Mbpctrefreshthresh` options) to fine tune the values trigger for megablock refreshes.

Use the `detail` option with the following commands:

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

- `query vm`

Supported Clients

This option is valid for all clients. This option is not set in the client options file; use it by adding it to the command line when you enter any of the commands that support it. The Tivoli Storage Manager API does not support this option.

Syntax

►—DETail—►

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc query mgmtclass -detail
dsmc query filespace -detail
dsmc query backup file1 -detail

dsmc query vm -detail
```

Diffsnapshot

The `diffsnapshot` option controls whether the backup-archive client creates the differential snapshot when it runs a snapshot difference incremental backup.

If the differential snapshot is not created by the client, the latest snapshot found on the volume is used as the differential snapshot and as the source for the backup operation.

The default value is to create the differential snapshot. This option is ignored the first time that the `snappdiff` option is used. The first time the `snappdiff` option is used on a volume, a snapshot must be created and used as the source for a full incremental backup. Snapshots that are created by the backup-archive client are deleted by the client after the next snapshot difference incremental backup is complete.

Snapshots can be created with the Network Appliance FilerView tool. Use the `latest` parameter if you want the client to use the most recent snapshot that was created with this or any other method. Snapshots that are created by methods outside of Tivoli Storage Manager are never deleted by the client.

Supported Clients

This option is valid for the following clients:

- AIX 64-bit clients
- Linux x86_64 clients

Syntax



Parameters

create

Specifies that you want to create a new, persistent, snapshot to use as the source snapshot. This value is the default.

latest

Specifies that you want to use the latest snapshot that is found on the file server as the source snapshot.

Examples

Command line:

Perform a snapshot difference incremental backup of an NFS mounted file system `/vol/vol1` hosted on the file server `homestore.example.com`, where `/net/home1` is the mount point of `/vol/vol1`.

```
incremental -snapdiff -diffsnapshot=latest /net/home1
```

The `-diffsnapshot` option value of `latest` means that the operation uses the latest snapshot (the active snapshot).

Related concepts:

“Snapshot differential backup with an HTTPS connection” on page 167

Related tasks:

“Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 106

Related reference:

“Snapdiff” on page 519

“Snapdiffhttps” on page 524

“Createnewbase” on page 342

Diffsnapshotname

The `diffsnapshotname` option allows you to specify which differential snapshot, on the target filer volume, to use during a snapshot differential backup. This option is only specified if you also specify `diffsnapshot=latest`.

If this option is not specified, `diffsnapshot=latest` selects the most recent existing snapshot on the filer volume and uses it as the differential snapshot.

Supported Clients

This option can be used with supported x86_64 Linux and AIX clients.

Options File

This option can be specified in the client options file or on the command line.

Syntax

►—DIFFSNAPSHOTName— —*snapshot_name*—►

Parameters

snapshot_name

Specifies the name of an existing snapshot to use as the differential snapshot.

You can also use a pattern with wildcard characters to select a snapshot.

Wildcards can be either of the following characters:

* An asterisk (*) matches any character.

? A question mark (?) matches a single character.

The most recent snapshot that matches the wildcard pattern is selected as the differential snapshot.

Examples

Options file:

```
diffsnapshotname volume_base_snap
diffsnapshotname nightly.?
```

Command line:

```
dsmc incr \\DRFiler\UserDataVol_Mirror_Share -snapdiff
-useexistingbase -basenameshotname="nightly.?"
-diffsnapshot=latest -diffsnapshotname="nightly.?"
```

Related information

Basesnapshotname

Useexistingbase

Dirmc

The **dirmc** option specifies the management class you want to use for directories.

If you do not specify this option to associate a management class with directories, the client program uses the management class in the active policy set of your policy domain with the longest retention period. Select a management class for individual directories that retains directories at least as long as it retains the files associated with them.

If you specify a management class with this option, all directories specified in a backup operation are bound to that management class.

The **dirmc** option specifies the management class of directories that you back up and it does not affect archived directories. Use the **archmc** option with the **archive** command to specify the available management class for your policy domain to which you want to bind your archived directories and files. If you do not use the **archmc** option, the server binds archived directories to the default management class. If the default management class has no archive copy group, the server binds archived directories to the management class with the shortest retention period.

Important: Only extended attributes and ACLs are stored in storage pools. The directory information, other than extended attributes and ACLs, remains in the database. On Windows systems, directories occupy storage pool space.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

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

Syntax

►►—DIRMc— *—mgmtclassname—*—————►►

Parameters

mgmtclassname

Specifies the name of the management class that you want to associate with directories. The client uses the management class name that you specify for all of the directories that you back up. If you do not specify this option, the client associates the management class with the longest retention period with directories.

Examples

Options file:

`dirm managdir`

Command line

Does not apply.

Related information

If you want to back up specific files to a management class see “Assign a management class to files” on page 281 for more information.

Dirsonly

The `dirsonly` option processes directories *only*. The client does not process files.

Use the `dirsonly` option with the following commands:

- **archive**
- **incremental**
- **query archive**
- **query backup**
- **restore**
- **restore backupset**
- **retrieve**
- **selective**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—Dirsonly—►►

Parameters

There are no parameters for this option.

Examples

Command line:

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

Command line:

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

Disablenqr

The `disablenqr` option specifies whether the backup-archive client can use the no-query restore method for restoring files and directories from the server.

If you set the `disablenqr` option to `no` (the default), the client can use the no-query restore process.

If you set the `disablenqr` option to `yes`, the client can use only the standard restore process (also known as "classic restore").

Note: There is no option or value to specify that the client can use only the no-query restore method.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

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

Syntax

►►—DISABLENQR—

No
Yes

—►►

Parameters

No Specifies that the client can use the no-query restore method. This is the default.

Yes

Specifies that the client uses only the standard restore method. The no-query restore method is not allowed.

Examples

Options file:

disablenqr yes

Command line

-disablenqr=yes

Diskbuffsize

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

Optimal backup, archive, or HSM migration client performance can usually be achieved if the value for this option is equal to or smaller than the amount of file read ahead provided by the client file system. A larger buffer requires more memory and it might not improve performance.

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

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza.

Syntax

►►—DISKBuffsize— *size* —————►◄

Parameters

size

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

For AIX: If `enablelanfree no` is set, the default setting for `diskbuffsize` is 256.

Examples

Options file:

diskbuffsize 64

Command line:

Does not apply.

Diskcachelocation

The `diskcachelocation` option specifies the location where the disk cache database is created if the option `memoryefficientbackup=diskcachemethod` is set during an incremental backup.

You can specify the `diskcachelocation` option in your option file, or with the `include.fs` option. If the `diskcachelocation` option appears in the option file, its value is used for all file systems not represented by an `include.fs` option containing the `diskcachelocation` option.

The disk cache is a temporary file which is deleted after the **incremental** command is run. Use this option to select one of the following:

1. A location that has more free disk space if, when you are using `memoryefficientbackup=diskcachemethod`, you get the message that the disk cache file cannot be created because you do not have enough disk space.
2. A location on a different physical volume to reduce contention for the disk access mechanism, and therefore improve performance.

Important: For performance reasons, do not use a remote drive for `diskcachelocation`.

The actual amount of disk space required for the disk cache file created by disk cache incremental backups depends on the number of files and directories included in the backup and on the average length of the files and directories to be backed up. For UNIX and Linux, estimate 1 byte per character in the path name. For Mac OS X, estimate 4 bytes per character in the path name. For example, if there are 1 000 000 files and directories to be backed up and the average path length is 200 characters, then the database occupies approximately 200 MB for UNIX and Linux, and 800 MB for Mac OS X clients. Another way to estimate for planning purposes is to multiply the number of files and directories by the length of the longest path to establish a maximum database size.

A second disk cache file is created for the list of migrated files when backing up an HSM managed file system. The combined disk cache files, created by disk cache incremental backups and HSM managed file system backups, can require above 400 MB of disk space for each million files being backed up. The disk cache file can become very large. Large file support must be enabled on the file system that is being used for the disk cache file.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

Place this option in the `dsm.sys` file within a server stanza.

Syntax

►►—DISKCACHELocation— —*path*—◄◄

Parameters

path

Specifies the location where the disk cache database is created if `memoryefficientbackup=diskcachemethod`. The default location is to create the disk cache file in the root of the file space being processed.

Examples

Options file:

```
diskcachelocation /home  
diskcachelocation /Volumes/hfs2
```

Command line:

Does not apply.

See “Include options” on page 421 for more information about `include.fs`.

Domain

The `domain` option specifies what you want to include for incremental backup.

Domain objects are backed up only if you start the **incremental** command without a file specification.

The backup-archive client uses the domain value in the following situations to determine which file systems to process during an incremental backup:

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

You can define the domain option in the following locations:

- In an options file.
- On the command line, when entered with a client command.
- In a client option set, which is defined on the server with the **define clientopt** command.
- As an option on a scheduled command, which is defined on the server with the **define schedule** command.

If any of these sources contain a domain definition, the client backs up that domain. If more than one source specifies a domain, the client backs up all specified domains. The same domain object can be defined more than once, but the effect is the same as defining it only once. If you do not specify a domain, the client backs up the default domain, as described in the `all-local` parameter.

You can exclude objects from the domain by specifying the exclusion operator (-) before the object. If any domain definition excludes an object, that object is excluded from the domain, even if another definition includes the object. You cannot use the domain exclusion operator (-) in front of any domain keyword that begins with `all-`.

If a domain statement excludes one or more objects and no domain statement includes any objects, the result is an empty domain (nothing is backed up). You must specify the objects to include in the domain if any domain statements exclude objects.

Example 1: This example uses one domain statement to back up all local file systems except for /fs1:

```
domain all-local -/fs1
```

Example 2: This example uses multiple domain statements to back up all local file systems except for /fs1:

```
domain all-local domain -/fs1
```

Example 3: This example excludes /fs1 during a backup operation. If no other domain statement is used, the result is an empty domain. Nothing is backed up.

```
domain -/fs1
```

If you start the incremental command with a file specification, the client ignores any domain definitions and backs up only the file specification.

You can include a virtual mount point in your client domain.

Important: If you are running GPFS for AIX or GPFS for Linux x86_64 in a multinode cluster, and all nodes share a mounted GPFS file system, the client processes this file system as a local file system. The client backs up the file system on each node during an incremental backup. To avoid this situation, you can do one of the following tasks:

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

Automounted file systems

When you perform a backup with the domain option set to all-local, files that are handled by automounter and loopback file systems are not backed up.

If you back up a file system with the domain option set to all-local, any subdirectories that are mount points for an automounted file system (AutoFS) are excluded from a backup operation. Any files that exist on the server for the automounted subdirectory are expired.

When you perform a backup with the domain option set to all-lofs, all explicit loopback file systems (LOFS) are backed up and all automounted file systems are excluded. For loop devices and local file systems that are handled by automounter, set the domain option to all-auto-lofs.

Use the automount option with the domain parameters, all-auto-nfs, and all-auto-lofs to specify one or more automounted file systems to be mounted and added into the domain. If you specify the automount option, automounted file systems are remounted if they go offline during the execution of the **incremental** command.

Virtual mount points cannot be used with automounted file systems.

Important: On some Linux distributions, automounted file system mount points or maps of file system type (AutoFS) might not be listed in the current mount table. As a result, the automounted files systems, which are unmounted during backup or archive processing, might be incorrectly processed and stored as part of a wrong domain (for example, as part of domain `all-local`, `all-nfs`, or `all-lofs`, depending on the actual file system type). Therefore, in such Linux distribution environments, you must specify the appropriate automount option setting to correctly process your domain option setting at all points in time.

For HP-UX, the domain option is enhanced with the keywords `all-`, `auto-lofs`, and `all-auto-nfs` to support automounted file systems. To use this enhancement, you must use the automounter, AutoFS. To activate AutoFS, set the `autofs` parameter to 1 in the `/etc/rc.config.d/nfsconf` file. Changing this parameter requires a restart. If you encounter problems with NFS automounted file systems, install patches PHCO_24777 and PHNE_26388 (or later patches). For more information, refer to the HP-UX documentation.

For Mac OS X, automounted file systems are not supported. If an automounted file system is part of a domain statement, the backup fails and no files in the automounted file system are processed. Back up and restore the automounted file system from the host system. Do not back up or restore the automounted file system over a network connection.


Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the options file, `dsm.opt` or `dsm.sys`. In the `dsm.sys` file, you must place this option within a server stanza. You can set this option on the **Backup** tab, **Domain for Backup** section of the Preferences editor.

Syntax



Parameters

all-local

Backs up all local file systems except LOFS file systems and LOFS through automounter. This parameter is the default. The `/tmp` directory is not included.

domain

Defines the file systems to include in your default client domain.

When you use *domain* with the **incremental** command, it processes these file systems in addition to those file systems you specify in your default client domain.

-domain

Defines the file systems to exclude in your default client domain.

all-lofs

Backs up all loopback file systems, except those file systems that are handled by automounter. This parameter is not supported on Mac OS X.

all-nfs

Backs up all network file systems, except those file systems that are handled by automounter. This parameter is not supported on Mac OS X.

all-auto-nfs

Backs up all network file systems (but not local file systems) which are handled by automounter. This parameter is not supported on Mac OS X.

all-auto-lofs

Backs up all loop devices and local file systems that are handled through automounter. This parameter is not supported on Mac OS X.

object

Specifies the domain objects to include in the domain.

An object name must be enclosed in quotation marks if the name includes any spaces.

-object

Specifies the domain objects to exclude from the domain.

An object name must be enclosed in quotation marks if the name includes any spaces.

Examples

Options file:

An options file can contain more than one domain statement. However, each of the domain statements is an example of a single statement in an options file.

```
domain all-local
domain all-local -/Volumes/volume2
domain all-local '-/Volumes/Macintosh HD'
domain /tst /datasave /joe
"domain all-local"
domain ALL-LOCAL -/home
domain ALL-NFS -/mount/nfs1
```

A single domain statement can list one or more objects for the domain. You can use more than one domain statement. The following two examples from two options files yield the same domain result:

Example 1

```
...
domain fs1
domain all-local
domain -fs3
...
```

Example 2

```
...  
domain all-local fs1 -fs3  
...
```

Command line:

```
-domain="/ /Volumes/volume2"  
-domain="all-local -/Volumes/volume2"  
-domain="/fs1 /fs2"  
-domain=/tmp  
-domain="ALL-LOCAL -/home"
```

Domain definition interaction

Domain can be defined in several sources, and the result is a summation of all domain definitions. As an example of the interaction of domain definitions, consider how domain definitions from several sources yield different backup results. In the table, *FS* followed by a number (for example, FS1) is a file system. This table shows only commands that are entered on the command line. For scheduled commands, the command-line column is not relevant, and options from the scheduled command must be considered.

Table 72. Interaction of domain definitions from several sources

Options file	Command line	Client option set	Objects backed up using the incremental command
domain FS1	incremental -domain=FS2	domain FS3	FS1 FS2 FS3
domain FS1	incremental	domain FS3	FS1 FS3
	incremental -domain=FS2		FS2
	incremental -domain=FS2	domain FS3	FS2 FS3
	incremental	domain FS3	FS3
	incremental		all-local
domain all-local	incremental	domain FS3	all-local + FS3
domain all-local domain -FS1	incremental		all-local, but not FS1
domain -FS1	incremental		none
domain FS1 FS3	incremental	domain -FS3	FS1
domain all-local	incremental	domain -FS3	all-local, but not FS3
	incremental FS1 -domain=all-local		FS1
	incremental FS1	domain all-local	FS1
domain -FS1	incremental FS1		FS1

Related information

For information about defining a virtual mount point, see “Virtualmountpoint” on page 566.

For information about specifying one or more automounted file systems to be mounted and added into the domain, see “Automount” on page 328.

Domain.image

The `domain.image` option specifies what you want to include in your client domain for an image backup.

If you do not specify a file system with the **backup image** command, the file systems you specify with the `domain.image` option are backed up.

When you specify a file system with the **backup image** command, the `domain.image` option is ignored.

If you do not use the `domain.image` option to specify file systems in your client options file, and you do not specify a file system with the **backup image** command, a message is issued and no backup occurs.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86_64, Linux on POWER, and Solaris. The server can also define this option. The Tivoli Storage Manager API does not support this option.

This option is valid for all supported Windows clients. The server can also define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option in the **Backup > Domain for Backup** box in the Preferences editor.

Syntax



Parameters

domain

Defines the file systems or raw logical volumes to include in your default client image domain.

Examples

Options file:

`domain.image /fs1 /fs2`

Command line:

Does not apply.

Domain.nas

The `domain.nas` option specifies the volumes to include in your NAS image backups.

You can specify `all-nas` to include all the mounted file systems on the NAS file server, except those you exclude with the `exclude.fs.nas` option.

The backup-archive client uses your domain for NAS image backups when you run a **backup nas** command and you do not specify which volumes to process.

When you use this option in your client system options file (dsm.sys), the `domain.nas` option defines your default domain for NAS image backups. When you perform a NAS file system image backup using the **backup nas** command, the client adds the volumes that you specify on the command line to the volumes defined in your dsm.sys file. For example, if you enter `domain.nas nas1/vol/vol0 nas1/vol/vol1` in your dsm.sys file and you enter `dsmc backup nas -nasnodename=nas1 /vol/vol2` on the command line, the client 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, the client backs up all mounted volumes on the NAS file server. When performing a backup, if you use a file specification and set the `domain.nas` option to `all-nas` in the `dsm.sys` file, `all-nas` takes precedence.

Supported Clients

This option is only valid for AIX and Solaris clients. The server can also define this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza.

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, the client 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.

Domain.vcd

The `domain.vcd` option specifies the vCloud vApps to include when the backup-archive client processes a **Backup VAPP** command where the domain is not specified on the command line.

The `domain.vcd` option works with the `vcldhost` option. The `vcldhost` option identifies the vCloud Director server that manages the vApps that you want to protect. The `domain.vcd` option narrows the focus of the **Backup VAPP** command to a subset of the vApps in an organization and an organization virtual data center.

All vApps that are included on this `domain.vcd` option are processed by a **backup vapp** command that does not specify which vApps to process. To selectively backup vApps, include the vApps that you want to process on the vApp specification parameters on the **backup vapp** command. For more information about selectively backing up vApps, see “**Backup VAPP**” on page 628.

The **restore vapp** and **query vapp** commands do not globally process all vApps defined on this `domain.vcd` statement. Instead, you specify the vApps that you want to process on the vApp specification parameters on the on the command line. For more information about including vApps in a restore or query operation, see “**Restore VAPP**” on page 712 or “**Query VAPP**” on page 687.


Supported Clients

This option is valid for Linux clients that are configured to perform vCloud Director vApp backups.

Options file

Set this option in the client options file (`dsm.opt`), or on the command line.

Syntax



Parameters

ORG=organization_name

Specifies the name of the organization that the organization virtual data center belongs to.

ORGVDC=org_vdc_name

Specifies the name of the organization virtual data center, within the organization, that contains the vApps to protect.

VAPP=vapp_name

Specifies one or more vApps to protect. If you omit vApp names, all vApps in the specified organization and organization virtual data center are processed. To specify multiple vApp names, separate the names in a comma-separated list.

You must specify these keyword parameters in the following order: ORG=, ORGVDC=, VAPP=. The following syntax rules also apply:

- To specify more than one organization virtual data center name, separate the organization virtual data center names with commas and do not insert a space character after the commas.

```
"org=organization_name,orgvdc=org_vdc_name1,org_vdc_name2"
```

- To specify more than one vApp name, separate the vApp names with commas and do not insert a space character after the commas.

```
"org=organization_name_name,orgvdc=org_vdc_name,  
vapp=vapp_name1,vapp_name2,vapp_name3"
```

- To specify more than one domain, separate the domain parameters with semicolons and do not insert a space character after the semicolons.

```
"org=organization_name1,orgvdc=org_vdc_name1;  
organization=organization_name2,orgvdc=org_vdc_name2"
```

Examples

The following are examples showing vApp domains specified on the `domain.vcd` option in the client options file.

To include vApps from Org1 and Org1_vdc, specify the following option:

```
"org=Org1,orgvdc=Org1_vdc"
```

To include all vApps in multiple organizations and organization virtual data centers, specify each domain and separate them with a semicolon.

```
"org=Haifa,orgvdc=vdc1;  
org=Mainz,orgvdc=vdc2;org=SanJose,orgvdc=vdc4"
```

For information about using domain options in a vApp specification on the command line, see “**Backup VAPP**” on page 628, “**Query VAPP**” on page 687, and “**Restore VAPP**” on page 712

Domain.vmfull

The `domain.vmfull` option specifies the virtual machines to include in your full virtual machine image backup operations.

For VMware virtual machine backups, the `domain.vmfull` option works with the `vmchost` option. The `vmchost` option identifies the vCenter server or ESX server that contains the virtual machines that you want to protect. The `domain.vmfull` parameters are used to narrow the focus of an operation to a subset of the virtual machines that are running on the system that is identified by `vmchost`.

You can specify which virtual machines are to be processed by using any of the following techniques:

- Use the `VM=` option and specify the name of a virtual machine.
- Provide a comma-separated list of virtual machine names.
- Use wildcard syntax to process virtual machines that match the name pattern.

- Use one of the following domain-level parameters:

```
all-vm
all-windows
vmhost
vmfolder
vmhostcluster
vmdatastore
```

When you use domain-level parameters, virtual machines that are created in the domain are automatically included when the next backup occurs. For example, if you use the `vmfolder` parameter to back up all virtual machines included in a folder, any new virtual machines that get added to that folder are included in the next backup. The same is true of pattern-matched names that are included in a wildcard match.

The virtual machines that are specified on the `domain.vmfull` option are processed only when the **backup vm** command is entered without specifying a virtual machine or a list of virtual machines on the command line.

Supported Clients

This option can be used with supported x86_64 Linux clients.


The server can also define this option.

Options file

Set this option in the client options, by using the command line, or by using the **VM Backup** tab of the Preferences editor.

Restriction: The `vmname:vmrk=vmrk_label` parameters cannot be set in the Preferences Editor. Include this setting in the options file, or on the command line when you run a **backup vm** command.

Syntax for VMware virtual machines



Syntax rules: Multiple keywords must be separated by a semicolon. There cannot be any spaces after the semicolons. Multiple virtual machine or domain names must be separated by commas, with no space characters. For examples, see `vm=vmname`.

Parameters

vmname

Specifies the virtual machine name that you want to process. The name is the virtual machine display name. You can supply a list of virtual machine host names, by separating the names with commas (vm1,vm2,vm5).

vm=vmname

The vm= keyword specifies that the next set of values is a list of virtual machine names. The vm= keyword is the default and is not required.

In this example, vm= is not specified and commas are used to separate the machine names.

```
domain.vmfull my_vm1,my_vm2
```

If you specify multiple keywords, such as vm= and vmfolder=, the values that the keywords refer to must be separated by semicolons, with no intervening space characters:

```
domain.vmfull vm=my_vm1;vm=my_vm2
domain.vmfull vm=my_vm1;vmfolder=folder1;vmfolder=folder2
```

Wildcard characters can be used to select virtual machine names that match a pattern. An asterisk (*) matches any sequence of characters. A question mark (?) matches any single character, for example:

- Exclude all files that have “test” in the host name: -vm=*test*
- Include all virtual machines with names such as: “test20”, “test25”, “test29”, “test2A”: vm=test2?

You can exclude a virtual machine from a backup operation by specifying the exclude operator (-) before the vm= keyword. For example, -vm is used to exclude a particular machine, or machines, from a domain level backup, such as, ALL-Windows, ALL-VM, and VMFolder. If “vm1” is the name of a virtual machine in a folder that is named “accountingDept”, you can back up all of the virtual machines in the folder, but prevent the virtual machine “vm1” from being backed up. Set the following option:

```
domain.vmfull VMFolder=accountingDept;-vm=vm1
```

You cannot use the exclude operator (-) to exclude a domain, such as ALL-VM, ALL-Windows, or VMFolder. The exclude operator works only at the virtual machine name level.

vmname:vmdk=vmdk_label

The :vmdk= keyword applies only to VMware virtual machines and its use requires a Tivoli Storage Manager for Virtual Environments: Data Protection for VMware license.

This option is typically used to exclude disks (see the :-vmdk syntax) from being backed up. You can also include virtual machine disks by using the INCLUDE.VMDISK option or exclude virtual machine disks by using the EXCLUDE.VMDISK option.

Detailed guidance for excluding disks when you back up VMware virtual machines is provided in the Tivoli Storage Manager for Virtual Environments Data Protection for VMware product documentation on IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS8TDQ/welcome>.

The virtual disks within a virtual machine have disk labels that uniquely identify each virtual disk. You use the :vmdk= keyword to specify the labels of

the virtual disks that you want to be included in a **Backup VM** operation. If you do not specify `:vmdk=` and a disk label, all virtual disks in the virtual machine are backed up.

Assume that there is a virtual machine named "my_vm_example". This virtual machine has four disks (labeled Hard Disk 1, Hard Disk 2, Hard Disk 3, Hard Disk 4). To include only Hard Disk 2 and Hard Disk 3 in a backup, add the `:vmdk=` keyword and disk label for those disks. Quotation marks are necessary around the parameters because the disk labels contain space characters. For example:

```
domain.vmfull "my_vm_example:vmdk=Hard Disk 2:vmdk=Hard Disk 3"
```

This next example backs up Hard Disk 1 and Hard Disk 2 on VM1, and Hard Disk 3 and Hard Disk 4 on VM2. A comma is used to separate the virtual machine information.

```
domain.vmfull "vm1:vmdk=Hard Disk 1:vmdk=Hard Disk 2",  
              "vm2:vmdk=Hard Disk 3:vmdk=Hard Disk 4"
```

Similar to the `-vm=` keyword, you can also use the exclusion operator (`-`) with `:vmdk=` to exclude disks from a backup operation.

To back up a virtual machine (vm1) and exclude disks 3 and 4, use the following syntax:

```
domain.vmfull "vm1:-vmdk=Hard Disk 3:-vmdk=Hard Disk 4"
```

To back up two virtual machines, vm1 and vm2, and exclude the first two disks on each machine, use the following syntax:

```
domain.vmfull "vm1 :-vmdk=Hard Disk 1:-vmdk=Hard Disk 2",  
              "vm2:-vmdk=Hard Disk 1:-vmdk=Hard Disk 2"
```

You can include one or more disks on a `domain.vmfull` statement. You can exclude one or more disks on a `domain.vmfull` statement. You can mix include and exclude disks on the same statement. For example, the following statement is valid:

```
domain.vmfull  
  "vm1:vmdk=Hard Disk 1:-vmdk=Hard Disk 2:vmdk=Hard Disk 3:vmdk:Hard Disk 4"
```

If an include statement is present, all other disks in the virtual machine are excluded from a backup operation, unless the other disks are also specified in an include statement. For example, the following statement excludes all hard disks on vm1, except for Hard Disk 1:

```
domain.vmfull "vm1:vmdk=Hard Disk 1"
```

Both of the following exclude Hard Disk 4 from a backup of vm1:

```
domain.vmfull "vm1:vmdk=Hard Disk 1:vmdk=Hard Disk 2:vmdk=Hard Disk 3"  
domain.vmfull "vm1:-vmdk=Hard Disk 4"
```

all-vm

For VMware virtual machines. This option processes all virtual machines that are defined to the Virtual Center or to the ESX server that is specified on the `vmchost` option.

all-windows

For VMware virtual machines. This option processes all virtual machines that are defined to the Virtual Center or to the ESX server that is specified on the `vmchost` option. The virtual machines must also have a guest operating system type of Windows.

vmhost=hostname

For VMware virtual machines. This option processes all virtual machines that

are defined to the Virtual Center or to the ESX server that is specified on the `vmhost` option. The host name that you specify must match the fully qualified host name or IP address, as it is specified in the vCenter server Hosts and Clusters view.

All virtual machines that are added to this host are automatically included in backup and restore processing. To be included, the virtual machines must also be running on the ESX server that is specified by the host name; they cannot be powered off.

This parameter can include multiple ESX servers that are separated by commas. When the Virtual Center contains multiple ESX servers, this option does not determine the ESX server from which a snapshot is taken. The ESX server from which a snapshot is taken is determined by the VMware VirtualCenter web service.

When you connect directly to an ESXi or ESX host, the `vmhost` option applies only if the **vmhost** is the server that you connect to. If it is not, a warning level message is sent to the console and is recorded in the client `dsmerror.log` file; it is also recorded as a server event message.

If the `vmenabletemplatebackups` option is set to yes, and VM templates are part of the domain, they are included in the backup.

Restriction: VMware templates for virtual machines cannot be backed up when they are in an ESX or ESXi host because ESX and ESXi hosts do not support templates.

`vmfolder=foldername`

For VMware virtual machines. This option processes all virtual machines that are defined to the Virtual Center or to the ESX server that is specified on the `vmhost` option. The virtual machines must also exist in the VMware folder that is specified by the folder name. Folder name can include multiple VMware folders that are separated by commas.

`vmhostcluster=hostclustername`

For VMware virtual machines. This option processes all virtual machines that are defined to the Virtual Center or to the ESX server that is specified on the `vmhost` option. The virtual machines must also be running on the ESX host cluster that is specified by the host cluster name. To include more than one host cluster name, separate the cluster names with commas:
`VMHOSTCLUSTER=cluster1,cluster2.`

If the `vmenabletemplatebackups` option is set to yes, and VM templates are part of the domain, they are included in the backup. A VMware host cluster is not available if you connect directly to an ESXi or ESX host. If you connect directly to an ESXi/ESX host and a domain is processed that includes a host cluster, a warning level message is sent to the console and is recorded in the client `dsmerror.log` file; it is also recorded as a server event message.

`vmdatastore=datastorename`

For VMware virtual machines. This option processes all virtual machines that are defined to the Virtual Center or to the ESX server that is specified on the `vmhost` option. The configured datastore location for a virtual machine must match the datastore name that is specified by `datastorename`. The datastore name can include multiple datastores that are separated by commas:
`VMDATASTORE=datastore1,datastore2`

Virtual machines can have their disk (vmdk files) on more than one datastore; but there is only one default datastore location. This default datastore location

is defined in the virtual machine configuration and is always where the virtual machine configuration file (.vmx file) is located. When a machine is selected for backup by using a domain keyword, the virtual machine configuration file, and all of the virtual machine's disks are included in the backup, including the disks that are on a different datastore than the one specified as the domain.

Examples

Options file:

Include all virtual machines in full VM backup operations.

```
domain.vmfull all-vm
```

Include all virtual machines in full VM backup operations, except for the ones that have a name suffix of _test.

```
domain.vmfull all-vm;-vm=*_test
```

Include all virtual machines that have Windows as the operating system, in full VM backup operations.

```
domain.vmfull all-windows
```

Include all virtual machines in cluster servers 1, 2, and 3 in full VM backup operations.

```
domain.vmfull vmhostcluster=cluster1,cluster2,cluster3
```

Include all virtual machine data in datastore1 in full VM backup operations.

```
domain.vmfull vmdatastore=datastore1
```

Include all virtual machines in full VM backup operations, but exclude virtual machines testvm1 and testvm2.

```
domain.vmfull all-vm;-VM=testvm1,testvm2
```

Include the virtual machines that are defined in the VM folders that are named lab1 and lab2 in full VM backup operations.

```
domain.vmfull vmfolder=lab1,lab2
```

Include all virtual machines on the ESX hosts named "brovar", "doomzoo", and "kepler" in full VM backup operations.

```
domain.vmfull vmhost=brovar.example.com,  
doomzoo.example.com,kepler.example.com
```

Related reference:

"Exclude.vmdisk" on page 394

"Include.vmdisk" on page 431

Dontload

x86_64 Linux clients can use the dontload option to suppress specific plug-in libraries from being loaded when the backup-archive client is started.

The TIVsm_BAhdw.x86_64 package provided in Linux x86_64 distributions contains software that is required to support snapshot incremental backups for NetAPP and N-Series file servers. When this package is installed on a Linux x86_64 system that is used to perform backup-archive client operations for a virtual machine, the files in this package cause all VMware backup-archive operations to fail. When these failures occur, the following message is displayed:

```
ANS8811E
```

VMware operations cannot be run when the hardware plug-in product TIVsm-BAhbw is installed and loaded. Either uninstall the hardware product TIVsm-BAhbw, or set the option DONTLOAD PIHDW in the options file to prevent the hardware plug-in from being loaded.

Use this option to prevent the plug-in library from being loaded into RAM when the client is started. Alternatively, you can uninstall the TIVsm_BAhbw package if it is not needed for snapshot operations.

Supported Clients

This option is only valid for Linux x86_64 clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza.

Syntax

►►—DONTLoad—PIHDW—◄◄

Parameters

PIHDW

Specifies that the hardware plug-in (TIVsm-BAhbw) is not loaded into RAM when the client is started. Use this option on backup-archive clients that have the hardware plug-in installed, to prevent the plug-in from causing failures when performing backup-archive operations on VMware virtual machines. There is no default for the dontload option.

To determine whether the plug-in is installed, enter the following command and examine the output.

```
rpm -q -a | grep TIV
```

If the output contains a package starting with “TIVsm-BAhbw” (followed by a version string), the hardware plug-in package is installed.

Examples

Options file:

```
DONTLoad PIHDW
```

Command line:

Does not apply. Do not use this option on the command line.

Related reference:

“Backup VM” on page 634

“Restore VM” on page 715

Dynamic image

Use the dynamicimage option with the **backup image** command or the **include.image** option to specify that you want to perform a dynamic image backup.

Supported Clients

This option is valid for AIX, Solaris, HP-UX and all Linux clients. The Tivoli Storage Manager API does not support this option.

Options File

Place the `include.image` statement containing the `dynamicimage` value in the server stanza in your system-options file, `dsm.sys`. You can also set this option using the Preferences editor.

Syntax

►►—DYNAMICImage— —value—————►►

Parameters

value

Specifies one of the following values:

yes

Use this option only if the volume cannot be unmounted and remounted as read-only. The client backs up the volume as is without remounting it as read-only. Corruption of the backup can occur if applications write to the volume while the backup is in progress. In this case, run `fsck` after a restore and manually mount the file system in order to regain access to the volume. This option is valid for AIX, Solaris, HP-UX and all Linux clients.

Note: This option is not allowed for AIX JFS2 file systems.

no Use this option if you do not want to perform a dynamic image backup. This is the default. The default behavior depends on the platform and file system type. For platforms and file systems that support snapshot, namely AIX JFS2 file systems and LINUX LVM file systems, the default is snapshot-based image backup. For all other UNIX platforms and file systems, the default is static image backup.

Examples

Options file:

```
include.image /kalafs1 dynamicimage=yes
```

Command line on backup image:

```
dynamicimage=yes
```

Efsdecrypt

The `efsdecrypt` option allows you to control whether or not files encrypted by an AIX Encrypted File System (EFS) are read in encrypted or decrypted format.

The `efsdecrypt` option default is `no`, which is to back up the encrypted or raw data. If you specify `yes`, the files are backed up as clear text, which means that they are backed up as normal files, as if the files existed in unencrypted form on the file system.

Important: Whenever you run a backup that includes any files encrypted on an EFS, you must ensure that you use the correct specification of the `efsdecrypt` option. If the `efsdecrypt` option value changes between two incremental backups,

all encrypted files on EFS file systems are backed up again, even if they have not changed since the last backup. For example, if you are running an incremental backup of encrypted files that were previously backed up as "raw," then ensure that `efsdecrypt` is specified as `no`. If you change `efsdecrypt` to `yes`, all the files are backed up again in clear text even if they are unchanged, so ensure that you use this option carefully.

Note: This is a global option that is applied to the complete backup. Two separate invocations of the client are required to back up some encrypted files as raw data and others as clear text.

Supported Clients

This option is valid for AIX clients.

Options File

Place this option in the `dsm.sys` file or the client user-options file (`dsm.opt`). In the `dsm.sys` file, you must place this option within a server stanza.

Syntax



Parameters

No Encrypted files are read in encrypted or raw data format, and Tivoli Storage Manager encryption and compression is forced off. This is the default.

Yes

Encrypted files are read in decrypted or clear text format.

Examples

Options file:

```
EFSDecrypt yes
```

Command line:

```
-EFSDecrypt=no
```

Enablearchiveretentionprotection

The `enablearchiveretentionprotection` option allows the client to connect to a Tivoli Storage Manager data retention server. This ensures that archive objects will not be deleted from the server until policy-based retention requirements for that object have been satisfied.

This option is ignored if the client connects to a server that is not retention protection enabled. If the option is `no` (the default) and an attempt is made to connect to a data retention server, the connection is refused.

The data retention server is specially configured for this task, so normal backup or restore processing is rejected by the server. When the client is connected to a data retention server, the following commands will not be available. If you attempt to use these commands, a message is displayed indicating that they are not valid with this server.

- **incremental**
- **backup** (all subcommands)
- **selective**
- **restore** (all subcommands except **restore backupset -location=file** or **-location=tape**)

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

- **restart restore**
- **delete backup**
- **delete group**
- **expire**
- All queries *except*:
 - **query access**
 - **query archive**
 - **query filespace**
 - **query inclexcl**
 - **query managementclass**
 - **query node**
 - **query options**
 - **query schedule**
 - **query session**
 - **query systeminfo**
 - **query tracestatus**


Supported Clients

This option is valid for all clients.

Options File

This option is valid only in the `dsm.sys` file *within* a server stanza and is not valid in a client option set from the server. It is not valid on any command line.

Syntax



Parameters

No The data retention server connection is refused. This is the default.

Yes

The client connects to a data retention server.

Enablededupcache

Use the `enablededupcache` option to specify whether you want to use a cache during client-side data deduplication. Using a local cache can reduce network traffic between the Tivoli Storage Manager server and the client.

When you perform a backup or archive operation with the data deduplication cache enabled, the specification of data extents that are backed up or archived are saved to the cache database. The next time you run a backup or archive, the client

queries the data deduplication cache and identifies the extents of data that have been previously saved to the server. Data extents that are identical to data extents on the server are not resent to the server.

If the server and the cache are not synchronized, the cache is removed and a new one is created.

Only one process can access the distributed data deduplication cache at a time. Concurrent backup instances on a workstation, that use the same server and storage pool, must either use unique node names or unique cache specifications. In this way, all the instances can use a local cache and optimize the client-side data deduplication.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API also supports this option.

Options File

Place this option in the system-options file (dsm.sys) within a server stanza. You can set this option on the **Deduplication > Enable Deduplication Cache** check box of the Preferences editor. The option can be set in the client option set on the Tivoli Storage Manager server.

Syntax



Parameters

Yes

Specifies that you want to enable data deduplication cache. If data deduplication is not enabled, this setting is not valid. Yes is the default for the backup-archive client. No is the default for the Tivoli Storage Manager API.

No Specifies that you do not want to enable data deduplication cache.

Examples

Options file:

enablededupcache no

Command line:

```
-enabledupcache=no
```

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

Related reference:

"Deduplication" on page 349

"Dedupcachepath" on page 348

"Dedupcachesize" on page 348

Enableinstrumentation

By default, instrumentation data is automatically collected by the backup-archive client and Tivoli Storage Manager API to identify performance bottlenecks during backup and restore processing. To disable or later enable instrumentation, use the `enableinstrumentation` option.

With this option enabled, you do not have to wait for a customer service representative to direct you to collect performance data when a problem occurs. Instead, the data can be collected whenever you run a backup or restore operation. This feature can be helpful because you do not have to re-create the problem just to collect performance data. The information is already collected by the client.

This option replaces the `-TESTFLAG=instrument:detail`, `-TESTFLAG=instrument:API`, and `-TESTFLAG=instrument:detail/API` options that are used in previous versions of the client and API.

For each process, the following types of performance instrumentation data are collected:

- The activity names for each thread (such as File I/O, Data Verb, Compression, and Transaction), the average elapsed time per activity, and the frequency of the activity.
- The total activity time of each thread.
- The command that was issued and the options that were used.
- The summary of the backup, restore, or query command.

By default, the performance data is stored in the instrumentation log file (`dsminstr.log`) in the directory that is specified by the `DSM_LOG` environment variable (or the `DSMI_LOG` environment variable for API-dependent products such as Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server and Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server). If you did not set the `DSM_LOG` environment variable, the instrumentation log file is stored in the current directory (the directory where you issued the `dsmc` command).

You can optionally change the name and location of the instrumentation log file by using the `instrlogname` option. You can also control the size of the log file by specifying the `instrlogmax` option.

Performance data is not collected for the backup-archive client GUI or web client GUI.

Performance data is collected for the following products when the `enableinstrumentation` option is specified in the client options file:

- Scheduled file-level backup operations with the backup-archive client
- Tivoli Storage Manager for Virtual Environments: Data Protection for VMware backups
- Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V backups
- Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server backups
- Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server backups

Performance data is also collected during archive and retrieve processing.

Supported Clients


This option is valid for all clients and the Tivoli Storage Manager API.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza.

Tip: This option is enabled by default, so typically, you do not need to place this option in the client options file unless you need to disable the option.

Syntax



Parameters

Yes

Specifies that you want to collect performance data during backup and restore operations. The default value is Yes, which means that performance data is collected even if you do not specify this option.

By default, the performance data is stored in the instrumentation log file (dsminstr.log) in the directory that is specified by the DSM_LOG environment variable. If you did not set the DSM_LOG environment variable, the instrumentation log file is stored in the current directory (the directory where you issued the **dsmc** command). If the file does not exist, the client creates the file and adds performance data to the file.

No Specifies that you do not want to collect performance data during backup and restore operations. If the instrumentation log exists, no more data is added to the file.

Examples

Options file:

```
enableinstrumentation yes
```

Command line:

```
dsmc sel /home/mydir/* -subdir=yes -enableinstrumentation=yes
```

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

Related tasks:

➡ Collecting client instrumentation data

➡ Collecting API instrumentation data

Related reference:

“Instrlogmax” on page 438

“Instrlogname” on page 439

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 backup-archive client and the SAN-attached storage device.

To support LAN-free data movement you must install and configure the IBM Tivoli Storage Manager for SAN storage agent on the client workstation.

Note:

1. If you place the `enablelanfree` option in the client option file (`dsm.opt`), but zero (0) bytes were transferred through the SAN during an operation, ensure that you bind the data to a LAN-free enabled management class.
2. To restore backup sets in a SAN environment, see “**Restore Backupset**” on page 698 for more information.
3. When a LAN-free path is enabled, the SAN Storage Agent settings override the client `tcpserveraddress`, `tcpport`, and `ssl` options. This override action occurs to ensure that both the client and the Storage Agent use the same server communication options.

Supported Clients

This option is valid for AIX, HP-UX, Linux x86_64, Linux on POWER, and Solaris clients.

Options File

Place this option in the `dsm.sys` file within a server stanza. You can also set this option by selecting the **Enable Lanfree** check box on the **General** tab in 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.

Related information

To specify a communication protocol between the backup-archive client and storage agent, see “Lanfreemethod” on page 441.

Encryptiontype

Use the encryptiontype option to specify the algorithm for data encryption.

The encryptiontype affects only backup and archive operations. The data that you include is stored in encrypted form, and encryption does not affect the amount of data that is sent or received. During restore and retrieve operations the encrypted data is decrypted with the proper encryption algorithm, regardless of the setting for this option.


Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza. 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 is the default.

AES256

AES 256-bit data encryption. AES 256-bit data encryption provides the highest level of data encryption available in backup and archive operations.

Restriction:

Files that are backed up using AES 256-bit encryption cannot be restored with an older level client.

DES56

DES 56-bit data encryption.

Examples

Options file:

encryptiontype aes128

Command line:

Does not apply.

Encryptkey

The backup-archive client supports the option to encrypt files that are being backed up or archived to the Tivoli Storage Manager server. This option is enabled with the `include.encrypt` option.

All files matching the pattern on the `include.encrypt` specification are encrypted before the data is sent to the server. There are three options for managing the key used to encrypt the files (prompt, save, and generate). All three options can be used with either the backup-archive client or the Tivoli Storage Manager API.

The encryption key password is case-sensitive and can be up to 64 characters in length.

The following characters can be included in the encryption key password:

A-Z	Any letter, A through Z, uppercase or lowercase. You cannot specify national language characters.
0-9	Any number, 0 through 9
+	Plus
.	Period
_	Underscore
-	Hyphen
&	Ampersand

Note:

1. The 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 API returns an error message.

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

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. 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 backup-archive client password file. A prompt is issued for an initial encryption key password, and after the initial prompt, the saved encryption key password in the password file is used for the backups and archives of files matching the `include.encrypt` specification. The key is retrieved from the password file on restore and retrieve operations.

The password can be up to 64 bytes in length.

When the `save` option is specified for an API application (does not apply to Mac OS X), the initial key password must be provided by the application using the API in the `dsmInitEx` function call. The API itself does not issue a prompt to the user but relies on the application to prompt the user as necessary.

This parameter is the default.

Note: The following restrictions apply:

- This option can only be used when `passwordaccess generate` is also specified.
- The root user or an authorized user must specify the initial encryption key password.

prompt

The management of the encryption key password is provided by the user. The user is prompted for the encryption key password when the client begins a backup or archive. A prompt for the same password is issued when restoring or retrieving the encrypted file.

This password can be up to 64 bytes in length.

When the `prompt` option is specified for an API application (does not apply to Mac OS X), the key password must be provided by the application using the API in the `dsmInitEx` function call. The API itself does not issue a prompt to the user but relies on the application to prompt the user as necessary.

generate

An encryption key password is dynamically generated when the 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 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. The default name for the error log is `dsmerror.log`.

Log wrapping is controlled by the `errorlogmax` option. If `errorlogmax` is set to zero (0), the size of the log is unlimited; logged entries never “wrap” and begin

overwriting earlier logged entries. If `errorlogmax` is not set to zero, the newest log entries overwrite the oldest log entries after the log file reaches its maximum size.

Log pruning is controlled by the `errorlogretention` option. Pruned logs do not wrap. Instead, log entries that are older than the number of days specified by the `errorlogretention` option are removed from the log file.

If you change from log wrapping (`errorlogmax` option) to log pruning (`errorlogretention` option), all existing log entries are retained and the log is pruned using the new `errorlogretention` criteria. Pruned log entries are saved in a file called `dsmerlog.pru`.

If you change from using log pruning (`errorlogretention` option) to using log wrapping (`errorlogmax` option), all records in the existing log are copied to the `dsmerlog.pru` log file, the existing log is emptied, and logging begins using the new log wrapping criteria.

If you simply 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.

If neither `errorlogmax` nor `errorlogretention` is specified, the error log can grow without any limit on its size. You must manually manage the log contents to prevent the log from depleting disk resources. When the log has been created with neither option specified, if you later issue a command and specify the `errorlogretention` option, the log is pruned using the retention value specified. When the log has been created with neither option specified, if you later issue a command and specify the `errorlogmax` option, the existing log is treated as if it was a pruned log. That is, the content of the `dsmerror.log` file is copied to a file called `dsmerlog.pru` and new log entries are created in `dsmerror.log` and the log is wrapped when it reaches its maximum size.

Note: If you specify a non-zero value for `errorlogmax` (which enables log wrapping), you cannot use the `errorlogretention` option to create pruned logs. Logs can be pruned or wrapped, but not both.

Logs created with the `errorlogmax` option contain a log header record that contains information similar to this example record:

```
LOGHEADERREC 661 104857600 IBM Tivoli Storage Manager 6.3.0.52 11080B Tue Aug 9 06:46:53 2011
```

Note that the dates and time stamps in the `LOGHEADERREC` text are not translated or formatted using the settings specified on the `dateformat` or `timeformat` options.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza.

You can also set this option on the **Client preferences** tab in the GUI, by selecting **Enable error log file wrapping** and by specifying a non-zero **maximum size** for the log file. To prevent log file wrapping, set the **maximum size** to zero. When the maximum wrapping is set to zero, clearing or setting the **Enable error log file**

wrapping option has no effect; log wrapping does not occur if the **maximum size** is set to zero.

Syntax

►►—ERRORLOGMAX— —*size*—►►

Parameters

size

Specifies the maximum size, in megabytes, for the log file. The range of values is 0 to 2047; the default is 0, which disables log file wrapping and allows the log file to grow indefinitely.

Examples

Options file:

errorlogmax 2000

Command line:

-errorlogmax=2000

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

Errorlogname

This option specifies the fully qualified path and file name of the file that contains the error messages.

The value for this option overrides the DSM_LOG environment variable. The dsmwebcl.log and dsmsched.log files are created in the same directory as the error log file you specify with the errorlogname option.

For Mac OS X, the default location is one of the following:

~/Library/Logs/tivoli/tsm/
/Library/Logs/tivoli/tsm/

The dsmerror.log cannot be a symbolic link.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza. You can set this option on the **General** tab, **Select Error Log** button of the Preferences editor.

Syntax

►►—ERRORLOGName— —*filespec*—►►

Parameters

filespec

The fully qualified path and file name in which to store error log information. If any part of the path you specify does not exist, the client attempts to create it.

The dsmerror.log file cannot be a symbolic link.

Examples

Options file:

```
errorlogname /tmp/tsmerror.log
```

Command line:

```
-errorlogname=/tmp/tsmerror.log
```

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

Errorlogretention

The errorlogretention option specifies how many days to maintain error log entries before pruning, and whether to save the pruned entries in other files.

The error log is pruned when the first error is written to the log after a client 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 log pruning (errorlogretention option) to log wrapping (errorlogmax option), all records in the existing log are copied to the dsmerlog.pru log file, the existing log is emptied, and logging begins using the new log wrapping criteria.

If you change from log wrapping (errorlogmax option) to log pruning (errorlogretention option), all existing log entries are retained and the log is pruned using the new errorlogretention criteria. Pruned log entries are saved in a file called dsmerlog.pru.

If neither errorlogmax nor errorlogretention is specified, the error log can grow without any limit on its size. You must manually manage the log contents to prevent the log from depleting disk resources. When the log has been created with neither option specified, if you later issue a command and specify the errorlogretention option, the log is pruned using the retention value specified. When the log has been created with neither option specified, if you later issue a command and specify the errorlogmax option, the existing log is treated as if it was a pruned log. That is, the content of the dsmerror.log file is copied to a file called dsmerlog.pru and new log entries are created in dsmerror.log and the log is wrapped when it reaches its maximum size.

Note: If you specify errorlogretention option to create pruned logs, you cannot specify the errorlogmax option. Logs can be pruned or wrapped, but not both.

Supported Clients


This option is valid for all clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza.

You can also set this option on the **Client preferences** tab in the GUI, by selecting **Prune old entries** and by specifying a value for **Prune entries older than**. Selecting the **Save pruned entries** option saves the pruned log entries in the dsmerlog.pru log file.

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 dsmerror.log file.

Examples

Options file:

Prune log entries from the dsmerror.log file that are older than 365 days and save the pruned entries in dsmerlog.pru.errorlogretention 365 S

Command line:

-errorlogr=365,S

Options file:

Prune log entries from the dsmerror.log file that are older than 365 days and do not save the pruned entries.errorlogretention 365 D

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

Exclude options

Use the exclude options to exclude objects from backup, image, or archive services.

For example, you might want to exclude this type of information:

- All temporary files
- Any local caches of network files
- All files that contain compiled object code that you can easily reproduce using other methods
- Your operating system files

You can exclude specific files from encryption processing during a backup.

Note:

1. With the exception of `exclude.fs`, when you exclude a file that was previously included, existing backup versions become inactive during the next incremental backup.
2. The server can define exclude options with the `incl excl` option.

Exclude any system files or images that could corrupt the operating system when recovered. Also exclude the directory containing the Tivoli Storage Manager client files.

Use wildcard characters to exclude a broad range of files.

To exclude an entire directory called `/any/test`, enter the following:

```
exclude.dir /any/test
```

To exclude subdirectories that begin with `test` under the `/any` directory, enter the following:

```
exclude.dir /any/test*
```

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set these options on the **Include-Exclude** tab, **Define Include-Exclude Options** section of the Preferences editor.

Syntax

►► *options* — *pattern* —————►►

exclude, exclude.backup, exclude.file, exclude.file.backup

Use these options to exclude a file or group of files from backup services and space management services (if the HSM client is installed). The `exclude.backup` option only excludes files from normal backup, but not from HSM.

exclude.archive

Excludes a file or a group of files that match the pattern from archive services *only*.

exclude.attribute.symlink

Excludes a file or a group of files that are symbolic links or aliases (aliases apply to Mac OS X) from backup processing only.

Note: For Mac OS X aliases are excluded.

exclude.compression

Excludes files from compression processing if the `compression` option is set to `yes`. This option applies to backups and archives.

exclude.dedup

Excludes files from client-side data deduplication. To control a client-side data deduplication operation, specify `ieobjtype` as the value of the `exclude.dedup` option.

Valid `ieobjtype` parameters are:

- File
- Image

The default is `File`.

exclude.dir

Excludes a directory, its files, and all its subdirectories and their files from backup processing. For example, the statement `exclude.dir /test/dan/data1` excludes the `/test/dan/data1` directory, its files, and all its subdirectories and their files.

If you exclude a directory that was previously included, the server expires existing backup versions of the files and directories beneath it during the next incremental backup. Use this option to exclude a portion of your data that has no underlying files to back up.

Note: Avoid performing a selective backup, or a partial incremental backup, of an individual file within an excluded directory. The next time that you perform an incremental backup, any files backed up in this manner is expired.

exclude.encrypt

Excludes the specified files from encryption processing. This option does not affect whether files are excluded from backup or archive processing, only whether they are excluded from encryption processing.

exclude.fs

Excludes file systems that match the specified pattern from backup, incremental image backup, and archive operations. If files from the excluded file systems were ever backed up, then management class rebinding and deleted file expiration does not occur. However, existing backup versions remain on the server subject to associated management class settings. The files that were previously archived from the excluded file system remain on the server as archive copies.

The `exclude.fs` option does NOT prevent the backup or archive of any virtual mount points that are subdirectories of the excluded file system.

Use `exclude.image` to exclude file systems from full image backup operations.

exclude.fs.nas

Excludes file systems on the NAS file server from an image backup when used with the **backup nas** command. The NAS node name must be prefixed to the file system name, for example: `netappsj1/vol/vol1`. To apply the exclude to all NAS nodes, replace the NAS node name with a wildcard, for example: `*/vol/vol1`. The **backup nas** command ignores all other exclude statements including `exclude.fs` and `exclude.dir` statements. This option is valid for AIX and Solaris clients *only*.

exclude.image

Excludes mounted file systems and raw logical volumes that match the specified pattern from full image backup operations. This option is valid for

AIX, HP-UX, all Linux clients, and Solaris only. Use `exclude.fs` to exclude file systems from incremental image backup operations.

Restriction: This option does not apply to Mac OS X.

Parameters

pattern

Specifies the file or group of files that you want to exclude.

Note: For NAS file systems: You must prefix the NAS node name to the file specification to specify the file server to which the exclude statement applies. If you do not specify a NAS node name, the file system identified refers to the NAS node name specified in the client system-options file (`dsm.sys`) or on the command line.

If the pattern begins with a single or double quote or contains any embedded blanks or equal signs, you must surround the value in either single (') or double (") quotation marks. The opening and closing quotation marks must be the same type of quotation marks.

For the `exclude.image` option, the pattern is the name of a mounted file system or raw logical volume.

Examples

Options file:

```
exclude /unix/
exclude ../../core
exclude /home/jones/proj1/*
exclude.archive ../../core
exclude.backup /home/jones/proj1/devplan/
exclude.dir /home/jones/tmp
exclude.backup /users/home1/file1
exclude.image /usr/*/*
exclude.encrypt /users/home2/file1
exclude.compression /home/gordon/proj1/*
exclude.fs.nas netappsj/vol/vol0
exclude.attribute.symlink ../../*
exclude.dedup /Users/Administrator/Documents/Important/../../*
```

Command line:

Does not apply.

Related information

See “System files to exclude” on page 119 for a list of files that you should always exclude.

“Incl excl” on page 420

See “Include and exclude groups of files with wildcard characters” on page 121 for a list of wildcard characters that you can use. Then, if necessary, use the `include` option to make exceptions.

Controlling symbolic link and alias processing

The backup-archive client treats symbolic links and aliases (aliases apply to Mac OS X only) as actual files and backs them up. However, the file referenced by the symbolic link is not backed up. In some cases symbolic links can be easily recreated and need not be backed up.

In addition, backing up these symbolic links can increase backup processing time and occupy a substantial amount of space on the Tivoli Storage Manager server. You can use the `exclude.attribute.symlink` option to exclude a file or a group of files that are symbolic links from backup processing. If necessary, you can use the `include.attribute.symlink` option to include symbolic links within broad group of excluded files for backup processing.

For example, to exclude all symbolic links from backup processing, except those that exist under the `/home/spike` directory, enter these statements in your `dsm.sys` file:

```
exclude.attribute.symlink  /.../*
include.attribute.symlink  /home/spike/.../*
```

Related reference:

“Include options” on page 421

Controlling compression processing

This topic lists some items to consider if you want to exclude specific files or groups of files from compression processing during a backup or archive operation.

- Remember that the backup-archive client 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, the client does not perform compression processing.

If you set the compression option to yes and no `exclude.compression` statements exist, the client considers all files for compression processing.

- The client processes `exclude.fs`, `exclude.dir`, and other include-exclude statements first. The client then considers any `exclude.compression` statements. For example, consider the following include-exclude list:

```
exclude /home/jones/proj1/*.*
exclude.compression /home/jones/proj1/file.txt
include /home/jones/proj1/file.txt
```

The client examines the statements (reading from bottom to top) and determines that the `/home/jones/proj1/file.txt` file is a candidate for backup, but is not a candidate for compression processing.

- Include-exclude compression processing is valid only for backup and archive processing. The `exclude.compression` option does not affect whether files are excluded from backup or archive processing, only whether they are excluded from compression processing.

Related reference:

“Compression” on page 339

Processing NAS file systems

Use the `exclude.fs.nas` option to exclude file systems from NAS image backup processing.

Note: The `exclude.fs.nas` option does not apply to a snapshot difference incremental backup.

A NAS file system specification uses the following conventions:

- NAS nodes represent a unique node type. The NAS node name uniquely identifies a NAS file server and its data to the backup-archive client. 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
```

Virtual machine exclude options

Virtual machine include and exclude options influence the behavior of backup and restore operations for virtual machines. These options are processed before any command-line options are processed, so that options on the command line can override options specified on any of the virtual machine include options or virtual machine exclude options. See the individual option descriptions for information about the options.

Related reference:

“Exclude.vmdisk”

Exclude.vmdisk:

The EXCLUDE.VMDISK option excludes a VMware virtual machine disk from backup operations.

The EXCLUDE.VMDISK option specifies the label of a virtual machine's disk to be excluded from a **Backup VM** operation. If you exclude a disk on the **Backup VM** command, the command-line parameters override any EXCLUDE.VMDISK statements in the options file.

This option is available only if you are using the Tivoli Storage Manager for Virtual Environments licensed product. For additional information about this option, see the Tivoli Storage Manager for Virtual Environments product documentation on IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS8TDQ/welcome>.

Supported clients

This option can be used with supported x86_64 Linux clients.

Options file

Set this option in the client options file. Command line parameters override statements in the options file.

Syntax

►►—EXCLUDE.VMDISK—*vmname*—*vmdk_label*—◀◀

Parameters

vmname

Specifies the name of the virtual machine that contains a disk that you want to exclude from a **Backup VM** operation. The name is the virtual machine display

name. You can specify only one virtual machine name on each EXCLUDE.VMDISK statement. Specify additional EXCLUDE.VMDISK statements for each virtual machine disk to exclude.

The virtual machine name can contain an asterisk (*), to match any character string, and question mark (?) to match any one character. Surround the VM name with quotation marks (" ") if the VM name contains space characters.

Tip: If the virtual machine name contains special characters, such as bracket characters ([or]), the virtual machine name might not be correctly matched. If a virtual machine name uses special characters in the name, you might need to use the question mark character (?) to match the special characters in the VM name

For example, to exclude Hard Disk 1 in the backup of a virtual machine named "Windows VM3 [2008R2]", use this syntax in the options file: EXCLUDE.VMDISK "Windows VM3 ?2008R2?" "Hard Disk 1"

vmdk_label

Specifies the disk label of the disk that you want to exclude. Wildcard characters are not allowed. Use the **Backup VM** command with the -preview option to determine the disk labels of disks in a given virtual machine. See "**Backup VM**" on page 634 for the syntax.

Do not exclude disks on virtual machines that you are protecting with the INCLUDE.VMTSMVSS option, if the disks contain application data.

Examples

Options file

Assume that a virtual machine named vm1 contains four disks, labeled Hard Disk 1, Hard Disk 2, Hard Disk 3, and Hard Disk 4. To exclude disk 2 from **Backup VM** operations, specify the following in the options file:

```
EXCLUDE.VMDISK "vm1" "Hard Disk 2"
```

Exclude disks 2 and 3 from **Backup VM** operations:

```
EXCLUDE.VMDISK "vm1" "Hard Disk 2"  
EXCLUDE.VMDISK "vm1" "Hard Disk 3"
```

Command line

The command line examples show the use of the exclusion operator (-) before the vmdk= keyword, to indicate that the disk is to be excluded.

Exclude a single disk:

```
dsmc backup vm "vm1:-vmdk=Hard Disk 1"
```

Exclude disk 2 and disk 3:

```
dsmc backup vm "vm1:-vmdk=Hard Disk 2:-vmdk=Hard Disk 3"
```

Exclude disk 1 and disk 2 on vm1:

```
dsmc backup vm "vm1:-vmdk=Hard Disk 1:-vmdk=Hard Disk 2"
```

Related reference:

"**Backup VM**" on page 634

"**Restore VM**" on page 715

"Domain.vmfull" on page 369

"Include.vmdisk" on page 431

"INCLUDE.VMTSMVSS" on page 435

Fbbranch

Use the fbbranch option with the **backup fastback** or **archive fastback** commands.

The fbbranch option specifies the branch ID of the remote FastBack server to back up or archive. The fbbranch option is only required when the backup-archive client is installed on a FastBack DR Hub or when a dedicated proxy is connecting to a replicated DR Hub repository. Do not specify the fbbranch option when the backup-archive client is installed on the FastBack server.

Supported Clients

This option is valid for Linux x86_64 clients.

Options File

None. You can specify this option only on the command line. The server can also define or override this option.

Syntax

►►—FBBBranch=—*branch_ID*—————◄◄

Parameters

branch_ID

Specifies the FastBack server branch ID. The value is part of the disaster recovery configuration of the FastBack server.

Examples

Command line:

-FBBBranch=oracle

On a backup-archive client that is installed on a FastBack DR Hub:

```
dsmc backup fastback -fbpolicyname=policy1 -fbserver=myFbServer  
-fbbranch=oracle
```

Command line:

On a backup-archive client that is connecting to a repository on a remote FastBack DR Hub:

```
dsmc backup fastback -fbpolicyname=policy1 -fbserver=server1  
-Fbreposlocation=\\myDrHub.company.com\\REP  
-fbbranch=oracle
```

If the fbbranch option is specified on a backup-archive client workstation that is installed on a FastBack server, the fbbranch option is ignored.

Fbclientname

Use the fbclientname option with the **backup fastback** or **archive fastback** commands.

The fbclientname option is the name of one or more comma-separated FastBack clients to back up or archive from the backup proxy. The values for the fbclientname option are invalid if more than one policy is specified in the fbpolicyname option.

You cannot include spaces in the `fbclientname` option values.

If you do not specify any values for the `fbvolumename` option, all the volumes from all the FastBack clients in the policy that is specified are backed up. If you specify multiple FastBack clients in the `fbclientname` option, you cannot specify values for the `fbvolumename` option.

Supported Clients

This option is valid for Linux x86_64 clients.

Options File

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

Syntax



Parameters

client_name

Specifies the name of one or more FastBack clients. You can specify up to 10 FastBack client names.

Important:

When specifying the **archive fastback** or **backup fastback** command:

1. At least one FBpolicyName is always required.
2. You can specify up to 10 values for FBPolicyName, if no values are specified for both FBClientName and FBVolumeName.
3. When you specify a FBClientName value, there must be only one value for FBPolicyName.
4. You can specify up to 10 values for FBClientName if only one PolicyName is specified, and no values for FBVolumeName are specified.
5. When you specify the FBVolumeName option, you can have only one FBPolicy, and only one FBClientName specified.
6. You can specify multiple FBVolumeNames if condition 5 is satisfied.
7. You must always specify the FBReposLocation option for Linux.

Examples

Command line:

```

dsmc backup fastback -fbpolicyname=Policy1
                    -fbclientname=fbclient1,fbclient2
                    -fbserver=myFbServer
                    -fbreposlocation=/mnt/FBLocation

```

Backs up all volumes for FastBack clients `fbclient1` and `fbclient2` that are found in policy `Policy1`.

Fbpolicyname

Use the `fbpolicyname` option with the **backup fastback** or **archive fastback** commands.

The `fbpolicyname` option is the name of one or more comma-separated Tivoli Storage Manager FastBack policies that you want to back up or archive from the backup proxy. You must specify at least one policy name. Specify multiple policy names using a comma-delimited list of policies. There is no default value.

If one or more FB policy names contain spaces, you must specify them within quotation marks. Here is an example: "FB Policy NAME1, FBPolicy Name 2".

If you do not specify any values for the `fbclientname` and `fbvolumename` options, all the volumes from all the FastBack clients in the policies that are specified are backed up. If you specify multiple policies in the `fbpolicyname` option, you cannot specify values for the `fbclientname` and `fbvolumename` options.

If a policy specification contains both Windows and Linux FastBack clients, only the Linux volumes will be backed up or archived to the Tivoli Storage Manager server by the Linux backup-archive client.

At least one snapshot should exist in the Tivoli Storage Manager FastBack repository for the FastBack policies being archived or backed up prior to issuing the **dsmc** command

Supported Clients

This option is valid for Linux x86_64 clients.

Options File

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

Syntax



Parameters

policy_name

Specifies the name of the FastBack policies. You can specify up to 10 FastBack policy names.

Important:

When specifying the **archive fastback** or **backup fastback** command:

1. At least one FBpolicyName is always required.
2. You can specify up to 10 values for FBPolicyName, if no values are specified for both FBClientName and FBVolumeName.
3. When you specify a FBClientName value, there must be only one value for FBPolicyName.

4. You can specify up to 10 values for FBClientName if only one PolicyName is specified, and no values for FBVolumeName are specified.
5. When you specify the FBVolumeName option, you can have only one FBPolicy, and only one FBClientName specified. You must specify exactly one FBClientName. It cannot be omitted.
6. You can specify multiple FBVolumeNames if condition 5 is satisfied.
7. You must always specify the FBReposLocation option for Linux.

Examples

Command line:

```
dsmc backup fastback -fbpolicyname=Policy1,Policy2,Policy3
-fbserver=myFbServer
-fbreposlocation=\\myFbServer.company.com\REP
```

Backs up all volumes for all FastBack clients found in policies Policy1, Policy2 and Policy3.

To specify policies with spaces, enclose them in double quotation marks, for example:

```
-fbpolicyname="Policy 1,Policy2,Policy3"
```

Fbreposlocation

Use the `fbreposlocation` option with the **backup fastback** or **archive fastback** commands.

The `fbreposlocation` option specifies the location of the Tivoli Storage Manager FastBack repository for the backup-archive client proxy to connect to issue Tivoli Storage Manager FastBack shell commands necessary to mount appropriate snapshots.

This option is required on Linux systems. There is no default location.

If you specify the `fbreposlocation` option for a snapshot on the FastBack server, use the `server_name@WORKGROUP` format.

There are two ways to specify the FastBack repository location on the FastBack DR Hub.

- Specify the complete repository location via the option `-fbreposlocation=\\DR_Hub\rep_server`. When using this format, `DR_Hub` is the FastBack DR Hub machine name and `rep_server` is the name of the replicated FastBack server repository on the DR Hub.
- Specify the repository location using a combination of the `-fbreposlocation=` and `-fbbranch` options. When using this format, specify the DR Hub repository location via the option `-fbreposlocation=DR_Hub@WORKGROUP`, and specify the name of the replicated FastBack server repository on the DR Hub using the `-fbbranch` option.

If you use the format `-fbr=\\<fbserver>\REP`, specify two backslashes before `<fbserver>` and one backslash before `REP` when using the backup-archive client in interactive mode. If you are using this format as a Linux command `dsmc backup fastback -fbr=\\\\<fbserver>\\REP`, you must specify four backslashes before `<fbserver>` and two backslashes before `REP`. This is because the Linux shell interprets a backslash as an escape character; the first backslash is treated as an escape character for the following backslash.

Supported Clients

This option is valid for Linux x86_64 clients.

Options File

None. You can specify this option only on the command line. The server can also define or override this option.

Syntax

►►—FBReposlocation—*repository_location*—————►►

Parameters

repository_location

Specifies the Tivoli Storage Manager FastBack repository location.

Examples

Command line:

```
dsmc backup fastback -fbpolicyname=Policy1  
-fbclientname=fbclient1,fbclient2 -fbserver=myFbDrHub  
-fbreposlocation=\\myFbDrHub\\rep_myFbServer
```

Note: Because Linux is supported only as a dedicated proxy configuration, a repository location is always required on Linux.

Command line:

```
dsmc backup fastback -fbpolicyname=Policy1  
-fbclientname=fbclient1,fbclient2 -fbserver=myFbDrHub  
-fbreposlocation=myFbDrHub -fbbranch=rep_myFbServer
```

Note: Because Linux is supported only as a dedicated proxy configuration, a repository location is always required on Linux.

Fbserver

Use the fbserver option with the **backup fastback** or **archive fastback** commands.

The fbserver option specifies the short host name of the Tivoli Storage Manager FastBack server workstation that owns the repository specified by the fbreposlocation option. For a DR Hub, the fbserver option specifies the short name of the FastBack server workstation whose branch repository the backup-archive client is connecting to.

The fbserver option is a key to retrieving the necessary user credentials required to connect to the FastBack server repository or the DR Hub server repository for mount processing.

Supported Clients

This option is valid for Linux x86_64 clients.

Options File

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

Syntax

►► — -FBServer — *server_name* —►►

Parameters

server_name

Specifies the short hostname of the machine on which a FastBack server is installed.

Examples

Command line:

The backup-archive client is installed on a Linux proxy client machine. Use this command to archive all FastBack volumes for all Linux FastBack clients that are defined for FastBack policy1:

```
dsmc archive fastback -fbpolicyname=Policy1
-fbserver=myfbserver
-fbreposlocation=myfbserver@WORKGROUP
```

The repository location is required. If you do not provide the repository location, the command will fail.

The FastBack server name, -myfbserver, is the short host name of the FastBack server where the repository is located.

Command line:

The repository, rep_server1, is located on a FastBack disaster hub, myFbDrHub.

```
dsmc archive fastback -fbpolicyname="Policy 1"
-fbserver=myFbDrHub
-fbreposlocation=\\myFbDrHub\\rep_server1
```

The FastBack server name, -myFbDrHub is the short host name of the FastBack disaster recovery hub server where the repository is located

The -fbreposlocation specifies the location of the repository. The repository location is required. If you do not provide the repository location, the command fails.

-fbserver should point to the short host name of the FastBack DR hub in this case.

Command line:

Archive all volumes protected by FastBack policy named policy1 from the FastBack server named basil:

```
dsmc archive fastback -Fbpolicyname=policy1
-FBServer=basil -ARCHMC="my_tsm_mgmt_class"
-fbreposlocation=basil@WORKGROUP
```

Fbvolumename

Use the fbvolumename option with the **backup fastback** or **archive fastback** commands.

The fbvolumename option is the name of one or more comma-separated Tivoli Storage Manager FastBack volumes to back up or archive from the backup proxy. Values for the fbvolumename option are not valid if more than one FastBack client is specified in the fbclientname option.

If you specify multiple FastBack clients in the `fbclientname` option, you cannot specify values for the `fbvolumename` option.

Supported Clients

This option is valid for Linux x86_64 clients.

Options File

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

Syntax



Parameters

volume_name

Specifies the name of the Tivoli Storage Manager FastBack volumes. You can specify up to 10 FastBack volume names.

Important:

When specifying the **archive fastback** or **backup fastback** command:

1. At least one FBpolicyName is always required.
2. You can specify up to 10 values for FBPolicyName, if no values are specified for both FBClientName and FBVolumeName.
3. When you specify a FBClientName value, there must be only one value for FBPolicyName.
4. You can specify up to 10 values for FBClientName if only one PolicyName is specified, and no values for FBVolumeName are specified.
5. When you specify the FBVolumeName option, you can have only one FBPolicy, and only one FBClientName specified. You must specify exactly one FBClientName. It cannot be omitted.
6. You can specify multiple FBVolumeNames if condition 5 is satisfied.
7. You must specify the FBReposLocation option.

Examples

Command line:

```
dsmc backup fastback -fbpolicyname=Policy1 -fbclientname=client1
                    -fbvolumename=data1,data2 -fbserver=myFbDrHub
                    -fbreposlocation=\\myFbDrHub\\rep_server1
```

Backs up volumes data1 and data2 from FastBack client client1 found in policy Policy1.

Filelist

Use the `filelist` option to process a list of files.

You can use the `filelist` option with the following commands:

- **archive**

- **backup group**
- **delete archive**
- **delete backup**
- **expire**
- **incremental**
- **query archive**
- **query backup**
- **restore**
- **retrieve**
- **selective**

The backup-archive client opens the file you specify with this option and processes the list of files within according to the specific command. Except for the **restore** and **retrieve** commands, when you use the `filelist` option, the client 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-qualified or a relative path to a file or directory. Note that if you include a directory in a `filelist` entry, the directory is backed up, but the contents of the directory are not.
- Each path must be specified on a single line. A line can contain only one path.
- Paths must not contain control characters, such as 0x18 (CTRL-X), 0x19 (CTRL-Y) and 0x0A (newline).
- By default, paths must not contain wildcard characters. Do not include asterisk (*) or question marks (?) in a path. This restriction can be overridden if you enable the option named `wildcardsareliteral`. For more information about that option, see “`Wildcardsareliteral`” on page 600.
- The `filelist` can be an MBCS file or a Unicode file with all Unicode entries. For Mac OS X, the `filelist` can be encoded in the current operating system language or UTF-16.
- If it is set, the client option called `quotesareliteral` allows quotation marks in a file specification to be interpreted literally, as quotation marks and not as delimiters. For more information about that option, see “`Quotesareliteral`” on page 487. If `quotesareliteral` and `wildcardsareliteral` are not set, quotation mark and wildcard processing works as described in the following list:
 - If a path or file name contains a space, enclose the entire path in quotation marks (") or single quotation marks ('). For example "C:\My Documents\spreadsheet.xls" or 'C:\My documents\spreadsheet.xls'.
 - If a path contains one or more single quotation marks ('), enclose the entire entry in quotation marks ("). If a path contains one or more quotation marks, enclose the entire path in single quotation marks. File list processing does not support paths that include a mix of quotation marks and single quotation marks.

The following examples illustrate the correct and incorrect use of quotation marks and single quotation marks in paths.

This path example contains a single quotation mark, so the path must be enclosed in quotation marks:

```
"/home/gatzby/mydir/gatzby's_report.out"
```

This path example contains quotation marks, so it must be enclosed in single quotation marks:

```
'/home/gatzby/mydir/"top10".out'
```

This path example contains a space character, so it must be enclosed in either quotation marks or single quotation marks:

```
"/home/gatzby/mydir/top 10.out"
```

or

```
'/home/gatzby/mydir/top 10.out'
```

This path example is not supported for filelist processing because it contains unmatched delimiters (" and '):

```
/home/gatzby/mydir/andy's_"top 10" report.out
```

These paths are not supported for filelist processing because they contain wildcard characters:

```
/home/gatzby*  
/home/*/20??.txt
```

- Any Tivoli Storage Manager filelist entry that does not comply with these rules is ignored.

The following are examples of valid paths in a filelist:

```
/home/dir/file1  
/usr/tivoli/file2  
/usr/avi/dir1  
/fs1/dir2/file3  
"/fs2/Ha Ha Ha/file.txt"  
"/fs3/file.txt"
```

To override standard processing of quotation marks and wildcard characters, see “Quotesareliteral” on page 487 and “Wildcardsareliteral” on page 600.

You can use the `filelist` option during an open file support operation. In this case, the client 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 is processed and not the files within the directory.

If the file name (the `filelistspec`) you specify with the `filelist` option does not exist, the command fails. The client skips any entries in the filelist that are not valid files or directories. The client logs errors and processing continues to the next entry.

Use file specifications with the **restore** and **retrieve** commands to denote the destination for the restored filelist entries. For example, in the following **restore** command, the file specification `/usr/record/` represents the restore destination for all entries in the filelist.

```
restore -filelist=/home/dir/file3 /usr/record/
```

However, in the following **selective** command, the file specification `/usr/record/` is ignored.

```
selective -filelist=/home/dir/file3 /usr/record/
```

If you specify a directory in a filelist for the **delete archive** or **delete backup** command, the directory is not deleted. filelists that you use with the **delete archive** or **delete backup** command should not include directories.

The entries in the list are processed in the order they appear in the filelist. For optimal processing performance, pre-sort the filelist by file space name and path.

Note: The client might back up a directory twice if the following conditions exist:

- The filelist contains an entry for the directory
- The filelist contains one or more entries for files within that directory
- No backup of the directory exists

For example, your filelist includes the entries `/home/dir/file1` and `/home/dir`. If the `/dir` directory does not exist on the server, the `/home/dir` directory is sent to the server a second time.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—FILEList =— —*filelistspec*—►►

Parameters

filelistspec

Specifies the location and name of the file that contains the list of files to process with the command.

Note: When you specify the `filelist` option on the command line, the `subdir` option is ignored.

Examples

Command line:

```
sel -filelist=/home/avi/filelist.txt
```

Related information

“Quotesareliteral” on page 487

“Wildcardsareliteral” on page 600

Filename

Use the `filename` option with the **query systeminfo** command to specify a file name in which to store information.

You can store information gathered from one or more of the following items:

- DSMOPTFILE - The contents of the `dsm.opt` file.
- DSMSYSFILE - The contents of the `dsm.sys` file.
- ENV - Environment variables.
- ERRORLOG - The Tivoli Storage Manager error log file.
- FILE - Attributes for the file name that you specify.
- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- OPTIONS - Compiled options.

- OSINFO - Name and version of the client operating system (includes ULIMIT information for UNIX and Linux).
- POLICY - Policy set dump.
- SCHEDLOG - The contents of the schedule log (usually dsmsched.log).
- CLUSTER - AIX cluster information.

Note: The **query systeminfo** command is intended primarily as an aid for IBM support to assist in diagnosing problems, although users who are familiar with the concepts addressed by this information might also find it useful. If you use the console option, no special formatting of the output is performed to accommodate screen height or width. Therefore, the console output might be difficult to read due to length and line-wrapping. In this case, use the **filename** option with the **query systeminfo** command to allow the output to be written to a file that can subsequently be submitted to IBM support.

Supported Clients

This option is valid for all clients.

Syntax

►►—FILENAME =— —*outputfilename*—————◄◄

Parameters

outputfilename

Specifies a file name in which to store the information. If you do not specify a file name, by default the information is stored in the dsminfo.txt file.

Examples

Command line:

```
query systeminfo dsmoptfile errorlog -filename=tsminfo.txt
```

Related information

“Console” on page 341

Filesonly

The **filesonly** option restricts backup, restore, retrieve, or query processing to files *only*.

You cannot restore or retrieve directories from the Tivoli Storage Manager server when using the **filesonly** option with the **restore** or **retrieve** commands. However, directories with default attributes are created, if required, as placeholders for files that you restore or retrieve.

You can also use the **filesonly** option with the following commands:

- **archive**
- **incremental**
- **query archive**
- **query backup**
- **restore**
- **restore backupset**
- **restore group**

- **retrieve**
- **selective**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►—FILESOnly—►

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc incremental -filesonly
```

Followsymbolic

During a backup operation, the `followsymbolic` option specifies whether you want to use a symbolic link as a virtual mount point. During a restore or retrieve operation, the `followsymbolic` option specifies how the backup-archive client restores a directory whose name matches a symbolic link on the restore target file system.

For backup operations, the `followsymbolic` option can influence the `virtualmountpoint` option setting. If you use the `virtualmountpoint` option to specify a symbolic link as a virtual mount point, you must also set the `followsymbolic` option.

During restore and retrieve operations, `followsymbolic` can influence how the client handles a symbolic link on the file system. Set `followsymbolic` only when the client attempts to restore a directory whose name matches a symbolic link on the restore target file system.

If you specify `followsymbolic=no` (the default), the client does not restore the contents of the directory, but returns this error message:

```
ANS4029E Error processing 'filespace name path-name file-name':
unable to build a directory path; a file exists with the same name
as a directory.
```

If you specify `followsymbolic=yes`, the client restores the contents of the directory to the target of the symbolic link.

For example, assume the client backed up a file with this path:
`/fs1/dir1/subdir1/file1`. Assume also that a symbolic link `/fs1/dir1`, that exists on the restore target file system, links to the directory `/fs88/dir88/subdir88`. Restore the file with the command:

```
restore /fs1/dir1/subdir1/file1
```

If you specify `followsymbolic=no`, the client does not restore the file, but returns the preceding error message. If you specify `followsymbolic=yes`, the client restores file1 to the `/fs88/dir88/subdir88/subdir1/file1` directory.

If you restore a symbolic link (not a directory) whose name matches a symbolic link on the restore target file system, the client restores the symbolic link.

If a symbolic link is used as a virtual mount point, the path to the link target must be specified by using an absolute file path.

Use this option with the **restore** and **retrieve** commands, or in the client user-options file (`dsm.opt`).

Supported Clients

This option is valid for all UNIX clients except Mac OS X.

Options File

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

Syntax



Parameters

No Do not back up a virtual mount point that is a symbolic link. Do not restore a directory if the restore target file system contains a symbolic link with matching name. This is the default.

Yes

Restore the contents of a directory to the target of a symbolic link.

Examples

Options file:

```
followsymbolic Yes
```

Command line:

```
-fol=Yes
```

Related information

During archive, the “Archsymlinkasfile” on page 319 option determines how the client handles symbolic links.

For operating systems other than Mac OS X, see “Back up symbolic links” on page 214 for more information about how the backup-archive client handles symbolic links.

Forcefailover

The `forcefailover` option enables the client to immediately fail over to the secondary server.

You can use the `forcefailover` option to immediately connect to the secondary server, even if the primary server is still online. For example, you can use this option to verify that the backup-archive client is failing over to the expected secondary server.

Do not edit this option during normal operations.


Supported Clients

This option is valid for all clients.

Options File

Place this option in the client-system options file (`dsm.sys`).

Syntax



Parameters

Yes

Specifies that the client immediately connects to the secondary server.

No Specifies that the client fails over to the secondary server during the next logon if the primary server is unavailable. This value is the default.

Examples

Options file:

```
FORCEFAILOVER yes
```

Command line:

```
-FORCEFAILOVER=yes
```

Related concepts:

“Automated client failover configuration and use” on page 87

Related tasks:

“Configuring the client for automated failover” on page 90

Fromdate

Use the `fromdate` option with the `fromtime` option to specify a date and time from which you want to search for backups or archives during a restore, retrieve, or query operation.

Files that were backed up or archived before this date and time are not included, although older directories might be included, if necessary, to restore or retrieve the files.

Use the `fromdate` option with the following commands:

- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**

- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►► FROMDate =— —*date*—————►►

Parameters

date

Specifies the date from which you want to search for backup copies or archived files. Enter the date in the format you selected with the `dateformat` option.

When you include `dateformat` with a command, it must precede the `fromdate`, `pitdate`, and `todate` options.

Examples

Command line:

```
dsmc query backup -fromdate=12/11/2003 "/Users/van/Documents/*"
```

Command line:

```
dsmc query backup -fromdate=12/11/2003 /home/dilbert/*
```

Fromnode

The `fromnode` option permits one node to perform commands for another node. A user on another node must use the **set access** command to permit you to query, restore, or retrieve files for the other node.

Use the `fromnode` option with the following commands:

- **query archive**
- **query backup**
- **query filespace**
- **query image**
- **query mgmtclass**
- **restore**
- **restore group**
- **restore image**
- **retrieve**

Supported Clients

This option is valid for all clients.

Syntax

►► FROMNode =— —*node*—————►►

Parameters

node

Specifies the node name on a workstation or a file server whose backup copies or archived files you want to access.

Examples

Command line:

```
dsmc query archive -fromnode=bob -subdir=yes  "/Users/van/Documents/
*"
```

Command line:

```
dsmc query archive -fromnode=bob -subdir=yes  "/home/jones/*"
```

Fromowner

The *fromowner* option specifies an alternate owner from which to restore backup versions or archived files or images. The owner must give access to another to use the files or images.

For example, to restore files from the */home/devel/proja* directory belonging to *usermike* on system *puma*, and place the restored files in a directory you own named */home/id/proja*, enter the following command:

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

Note: Archiving image restores does not apply to Mac OS X operating systems.

Non-root users can specify *fromowner=root* to access files owned by the root user if the root user has granted them access.

Note: If you specify the *fromowner* option without the *fromnode* option, the active user must be on the same node as the *fromowner* user.

Use the *fromowner* option with the following commands:

- **query archive**
- **query backup**
- **query group**
- **query image**
- **restore**
- **restore image**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

►►—FROMOwner =— —owner—►►

Parameters

owner

Name of an alternate owner.

Examples

Command line:

```
dsmc query archive "/home/id/proja/*" -fromowner=mark
```

Fromtime

Use the **fromtime** option with the **fromdate** option to specify a beginning time from which you want to search for backups or archives during a restore, retrieve, or query operation.

The backup-archive client ignores this option if you do not specify the **fromdate** option.

Use the **fromtime** option with the following commands:

- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►FROMTime =— —*time*—————►►

Parameters

time

Specifies a beginning time on a specific date from which you want to search for backed up or archived files. If you do not specify a time, the time defaults to 00:00:00. Specify the time in the format you selected with the **timeformat** option.

When you include the **timeformat** option in a command, it must precede the **fromtime**, **pittime**, and **totime** options.

Examples

Command line:

```
dsmc q b -timeformat=4 -fromt=11:59AM -fromd=06/30/2003 -tot=11:59PM  
-tod=06/30/2003 /home/*
```

Groupname

Use the **groupname** option with the **backup group** command to specify the name for a group. You can only perform operations on new groups or the current active version of the group.

Supported Clients

This option is valid for all UNIX and Linux clients except Mac OS X.

Syntax

►► —GROUPName = — —name ————— ►►

Parameters

name

Specifies the name of the group which contains the files backed up using the `filelist` option. Directory delimiters are not allowed in the group name since the group name is not a file specification, but a name field.

Examples

Command line:

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

Groups (deprecated)

This option is deprecated.

Guitreeviewafterbackup

The `guitreeviewafterbackup` option specifies whether the client returns to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client user-options file (`dsm.opt`) or the `dsm.sys` file. You can set this option on the **General** tab, **Return to tree window after function completed** check box of the Preferences editor.

Syntax

►► —GUITREEViewafterbackup *No*
Yes ————— ►►

Parameters

No Returns you to the Tivoli Storage Manager main window after a successful operation completes. This is the default.

Yes

Returns you to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.

Examples

Options file:

```
guitreeviewafterbackup yes
```

Command line:

Does not apply.

Host

The **host** option specifies the target ESX server location where the new virtual machine is created during a VMware restore operation.

Use this option on **restore vm** commands to specify the ESX host server to restore the data to.

Example

Restore the virtual machine to the ESX server named `vmesxbld1`.

```
restore vm -host=vmesxbld1.us.acme.com
```

Httpport

The **httpport** option specifies a TCP/IP port address for the web client.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **Web Client** tab, in the **HTTP Port** field of the Preferences editor.

Syntax

►►—HTTPport— *port_address*—————►►

Parameters

port_address

Specifies the TCP/IP port address that is used to communicate with the web client. The range of values is 1000 through 32767; the default is 1581.

Examples

Options file:

```
httpport 1502
```

Command line:

Does not apply.

Hsmreparsetag

The **hsmreparsetag** option specifies a unique reparse tag that is created by an HSM product installed on your system.

Many HSM products use reparse points to retrieve or recall migrated files. After a file is migrated, a small stub file, with the same name as the original file, is left on the file system. The stub file is a reparse point that triggers a recall of the original

file when a user or application accesses the stub file. The reparse point includes a unique identifier called a *reparse tag* to identify which HSM product migrated the file.

If the Tivoli Storage Manager backup-archive client does not recognize the reparse tag in a stub file, the Backup-Archive Client causes the HSM product to recall the original file. You can prevent files from being recalled if you specify the reparse tag with the `hsmreparsetag` option.

The backup-archive client recognizes the reparse tag of HSM products from the following companies:

- International Business Machines Corp.
- Wisdata System Co. Ltd.
- BridgeHead Software Ltd.
- CommVault Systems, Inc.
- Data Storage Group, Inc.
- Enigma Data Solutions, Ltd.
- Enterprise Data Solutions, Inc.
- Global 360
- GRAU DATA AG
- Hermes Software GmbH
- Hewlett Packard Company
- International Communication Products Engineering GmbH
- KOM Networks
- Memory-Tech Corporation
- Moonwalk Universal
- Pointsoft Australia Pty. Ltd.
- Symantec Corporation

If the HSM product you use is not in the preceding list, use the `hsmreparsetag` option to specify the reparse tag. Ask your HSM product vendor for the reparse tag used by the product.

Supported clients

This option is valid for all Windows clients.

Option file

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

Syntax

►►—HSMREPARSETAG—*reparse_tag_value*—————►►

Parameters

reparse_tag_value

A decimal (base 10) or hexadecimal (base 16) value that specifies the reparse tag.

Options file:

hsmreparsetag 22

```
hsmreparsetag 0x16
```

Does not apply.

Use the `ieobjtype` option to specify an object type for a client-side data deduplication operation within include-exclude statements.

The `ieobjtype` option is an additional parameter to the `include.dedup` or `exclude.dedup` options.

This option is valid for all clients. The Tivoli Storage Manager API also supports this option.

Place this option in the system-options file (dsm.sys). You can set this option on the **Include/Exclude** tab of the Preferences editor. The option can be set in the client option set on the Tivoli Storage Manager server.

```

graph LR
    IEObjtype --> File
    IEObjtype --> Image

```

File

Specifies that you want to include files for, or exclude files from, client-side data deduplication processing. File is the default.

Image

Specifies that you want to include images for, or exclude images from, client-side data deduplication processing.

Options file:

```
exclude.dedup /home/*/* ieobjtype=image
```

Does not apply.

"Exclude options" on page 389

"Include options" on page 421

Ifnewer

The `ifnewer` option replaces an existing file with the latest backup version only if the backup version is newer than the existing file.

Only active backups are considered unless you also use the `inactive` or `latest` options.

Note: Directory entries are replaced with the latest backup version, whether the backup version is older or newer than the existing version.

Use the `ifnewer` option with the following commands:

- `restore`
- `restore backupset`
- `restore group`
- `retrieve`

Note: This option is ignored if the `replace` option is set to *No*.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—IFNewer—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/Users/grover/Documents/*" -sub=y -rep=y -ifnewer
dsmc restore "/home/grover/*" -sub=y -rep=y -ifnewer
```

Imagegapsize

Use the `imagegapsize` option with the **backup image** command, in the options file, or with the `include.image` option to specify the minimum size of empty regions on a volume that you want to skip during image backup.

Use this option for LAN-based and LAN-free image backup.

For example, if you specify a gap size of 10, this means that an empty region on the disk that is larger than 10 KB in size is not backed up. Gaps that are exactly 10 KB are backed up. Empty regions that are exactly 10 KB and that are smaller than 10 KB is backed up, even though they do not contain data. However, an empty region that is smaller than 10 KB is backed up, even though it does not contain data. A smaller image gap size means less data needs to be transferred, but with potentially decreased throughput. A larger image gap size results in more data being transferred, but with potentially better throughput.

Supported Clients

This option is valid for AIX, Linux, and JFS2 clients only. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the server stanza of the client systems options file (`dsm.sys`), or in the `include.image` statement in the `dsm.sys` file.

Syntax

►►—IMAGEGapsize— —size—►►

Parameters

size

Specifies the minimum size of empty regions in an AIX JFS2 file system that should be skipped during an image backup. You can specify *k* (kilobytes) *m* (megabytes) or *g* (gigabytes) qualifiers with the value. Without a qualifier, the value is interpreted in kilobytes. Valid values are 0 through 4294967295 KB. If you specify a value of 0, all blocks, including unused blocks at the end of the volume, is backed up. If you specify any value other than 0, unused blocks at the end of the volume are not backed up. For LAN-based and LAN-free image backup the default value is 32 KB. This option is applicable to both static and snapshot-based image backup.

Note: This option is valid for AIX JFS2 file systems. If you specify an `imagegapsize` that is greater than 0 for a file system other than AIX JFS2, you get a warning message.

Examples

Options file:

Add the following to the server stanza in the `dsm.sys` file: `imagegapsize 1m`

Include-exclude list example: `include.image /kalafsl imagegapsize=-128k`

Command line:

`-imagegapsize=64k`

Imagetofile

Use the `imagetofile` option with the **restore image** command to specify that you want to restore the source image to a file.

You might need to restore the image to a file if bad sectors are present on the target volume, or if you want to manipulate the image data. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume.

Linux supports mounting an image file as a logical volume, so you can get access to file data within the image. The following are some examples:

- The file system `/usr` has been backed up by Tivoli Storage Manager. The following command restores the file system image to the file `/home/usr.img`:

```
# dsmc restore image /usr /home/usr.img -imagetofile
```


- To mount the image file at the /mnt/usr directory, the following mount command can be executed:

```
# mount /home/usr.img /mnt/usr -o loop=/dev/loop0
```

Now the image contents are available from /mnt/usr as if a regular file system was mounted at that directory.

Supported Clients

This option is valid for AIX, HP-UX, all Linux, and Solaris clients. The Tivoli Storage Manager API does not support this option.

Syntax

▶▶—IMAGETOfile—————▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore image /usr /home/usr.img -imagetofile
```

Inactive

Use the inactive option to display both active and inactive objects.

You can use the inactive option with the following commands:

- **delete group**
- **query backup**
- **query group**
- **query image**
- **query nas**
- **restore**
- **restore group**
- **restore image**
- **restore nas**

Important: When using the inactive option during a restore operation, also use the pick or some other filtering option because, unlike the latest option, all versions are restored in an indeterminate order. This option is implicit when pitdate is used.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

▶▶—INActive—————▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/Users/zoe/Documents/*" -inactive -pick
```

Command line:

```
dsmc restore "/home/zoe/*" -inactive -pick
```

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, such as `/etc`.

When processing occurs, the include-exclude statements within the include-exclude file are placed in the list position occupied by the `incl excl` option, in the same order, and processed accordingly.

If you have the HSM client installed on your workstation, you can use an include-exclude options file to exclude files from backup and space management, from backup only or from space management only.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

Place this option in the `dsm.sys` file *within* a server stanza. You can set this option on the **Include-Exclude** tab of the Preferences editor.

Syntax

►►—INCLExcl— *filespec*—————►►

Parameters

filespec

Specifies the path and file name of *one* include-exclude options file.

Examples

Options file:

```
INCLExcl /Users/user1/Documents/backup.excl
incl excl /usr/dsm/backup.excl
incl excl /etc/incl excl.def
```

Command line:

Does not apply.

Related information

For more information about creating an include-exclude options file, see “Creating an include-exclude list” on page 115.

Considerations for Unicode-enabled clients

An include-exclude file can be in Unicode or non-Unicode format.

If the codeset used to create an include-exclude list file does not match the codeset used on the client computer, characters in the file that cannot be mapped by the client’s codeset to a displayable character cannot be processed when backups are performed.

Using Unicode encoding for files containing include-exclude lists eliminates the unmapped character problem, so you no longer need to use wildcard characters as substitutes for the unrecognized characters.

Mac users: Create an include-exclude file in Unicode format by performing the following steps:

1. Open TextEdit. Click **Format > Make PlainText**.
2. Enter your include and exclude statements.
3. Click **File** and then click **Save As**.
4. From **PlainText Encoding**, select **Unicode (UTF-8)** or **Unicode (UTF-16)**, specify the file and target directory, and then save the file. Do not add the .txt extension.
5. Place an `incl excl` option specifying the include-exclude file you just created in your `dsm.sys` file.
6. Restart the backup-archive client.

Include options

The include options specify objects that you want to include for backup and archive services.

The include options specify any of the following:

- Objects within a broad group of excluded objects that you want to include for backup and archive services.
- Objects within a broad group of excluded objects that you want to include for backup, archive, image, and space management services.
- Files that are included for backup or archive processing that you want to include for encryption processing.
- Files that are included for backup or archive processing that you also want to include for compression processing.
- Objects to which you want to assign a specific management class.
- A management class to assign to all objects to which you do not explicitly assign a management class.
- File spaces to which you want to assign memory-efficient backup processing
- File spaces where you want to use the `diskcachelocation` option to cause specific file systems to use different, specific locations for their disk cache.

If you do not assign a specific management class to objects, the default management class in the active policy set of your policy domain is used. Use the **query mgmtclass** command to display information about the management classes available in your active policy set.

Remember: The backup-archive client compares the files it processes against the patterns specified in the include-exclude statements, reading from the bottom to the top of the options file.

Note:

1. The `exclude.fs` and `exclude.dir` statements override all include statements that match the pattern.
2. The server can also define these options with the `incl excl` option.

Supported Clients

This option is valid for all clients. The server can also define `include.fs.nas`.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set these options on the **Include-Exclude** tab in the Preferences editor.

Syntax

►► *options* — *pattern* — — *optional_parameter* — ►►

include, include.backup, include.file

Use these options to include files or assign management classes for backup processing.

The `include` option affects archive and backup processing. If you want to assign different management classes for archive and backup processing, always specify `include.archive` and `include.backup` with their own management classes. In this example, the `archmc` management class is assigned when an archive operation is performed. The management class is assigned when an archive operation is performed because `include.backup` is used only for backup processing, and not for archive processing.

```
include.archive /home/test/* archmc
include.backup /home/test/*
```

include.archive

Includes files or assigns management classes for archive processing.

include.attribute.symlink

Includes a file or a group of files that are symbolic links or aliases, within a broad group of excluded files for backup processing only.

Note: For Mac OS X, aliases are included.

include.compression

Includes files for compression processing if you set the `compression` option to `yes`. This option applies to backups and archives.

include.dedup

Includes files for client-side data deduplication. To control a client-side data

deduplication operation, specify `ieobjtype` as the value of the `include.dedup` option. By default, all data deduplication-eligible objects are included for client-side data deduplication.

Valid `ieobjtype` parameters are:

File
Image

The default is File.

include.encrypt

Includes the specified files for encryption processing. By default, the client does not perform encryption processing.

Notes:

1. The `include.encrypt` option is the only way to enable encryption on the backup-archive client. If no `include.encrypt` statements are used, encryption does not occur.
2. Encryption is not compatible with client-side deduplication. Files that are included for encryption are not deduplicated by client-side deduplication.
3. Encryption is not compatible with VMware virtual machine backups that use the incremental forever backup modes (`MODE=IFIncremental` and `MODE=IFFull`). If the client is configured for encryption, you cannot use incremental forever backup. However, you can use the full or incremental backup modes (`MODE=Full` and `MODE=Incremental`).
4. Encryption is not compatible with the Tivoli Storage Manager for Virtual Environments Data Protection for VMware Recovery Agent. If the client is configured for encryption, you can use the client to restore backups that were created with the full or incremental backup modes (`MODE=Full` and `MODE=Incremental`). However, you cannot use the Recover Agent to restore the encrypted backups.

include.fs

For AIX JFS2 file systems: Use the `snapshotcachesize` option in the `dsm.sys` file or with the `include.fs` option, to specify an appropriate snapshot size so that all old data blocks can be stored while the snapshot-based file backup or archive occurs.

To control how the client processes your file space for incremental backup, you can specify these additional options in your `dsm.sys` file, as values of the `include.fs` option: `diskcachelocation` and `memoryefficientbackup`.

Each of the `include.fs`, `memoryefficientbackup` and `diskcachelocation` options must be on the same line in the options file.

```
include.fs /home
    memoryefficientbackup=diskcachemethod
    diskcachelocation=/usr
include.fs /usr
    memoryefficientbackup=diskcachemethod
    diskcachelocation=/home
include.fs /Volumes/hfs3
    memoryefficientbackup=diskcachemethod
    diskcachelocation=/Volumes/hfs2
AIX JFS2 filesystems only: include.fs
    /kalafsl snapshotproviderfs=JFS2
```

If these options appear both in the options file and an `include.fs` option, the `include.fs` values are used for the specified file space in place of any values in an option file or on the command line.

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 the client saves Table of Contents (TOC) information during a NAS file system image backup, using the `toc` option with the `include.fs.nas` option in your `dsm.sys` file. This option is only valid for AIX and Solaris clients.

include.image

Includes a file space or logical volume, or assigns a management class when used with the **backup image** command. The **backup image** command ignores all other include options.

For Linux x86_64 clients: Use the `snapshotcachesize` option in these situations:

- With the **backup image** command
- In the `dsm.sys` file
- With the `include.image` option

Using the `snapshotcachesize` option in these situations lets you specify an appropriate snapshot size, so that all old data blocks can be stored while the image backup occurs.

A snapshot size of 100 percent ensures a valid snapshot.

For AIX JFS2 file systems: Use the `snapshotcachesize` option in these situations:

- With the **backup image** command
- In the `dsm.sys` file
- With the `include.image` option

Using the `snapshotcachesize` option in these situations lets you specify an appropriate snapshot size, so that all old data blocks can be stored while the image backup occurs.

This option is valid for AIX, HP-UX, Solaris, and all Linux clients only.

Parameters

pattern

Specifies the objects to include for backup or archive processing or to assign a specific management class.

Note: For NAS file systems: You must prefix the NAS node name to the file specification to specify the file server to which the include statement applies. If you do not specify a NAS node name, the file system identified refers to the NAS node name specified in the client system-options file (`dsm.sys`) or on the command line.

If the pattern begins with a single or double quotation mark, or contains any embedded blanks or equal signs, you must surround the value in either single (') or double (") quotation marks. The opening and closing quotation marks must be the same type of quotation marks.

For the `include.image` option, the pattern is the name of a mounted file system or raw logical volume.

optional_parameter

management_class_name

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. To associate a management class with a backup group on an include statement, use the following syntax:

```
include virtual_filespace_name\group_name management_class_name
```

where:

virtual_filespace_name

Specifies the name of the Tivoli Storage Manager server virtual filesystem that you associated with the group, on the **Backup Group** command.

group_name

Is the name of the group that you created when you ran the **Backup Group** command.

management_class_name

Is the name of the management class to associate with the files in the group.

For example, a group named MyGroup is stored in a virtual file space called MyVirtualFileSpace. To associate a management class, named TEST, with the group, use the following syntax:

```
include MyVirtualFileSpace/MyGroup TEST
```

Table 73. Other optional parameters

optional_parameter	Use with option
ieobjtype "Ieobjtype" on page 416	include.dedup
memoryefficientbackup "Memoryefficientbackup" on page 452	include.fs
diskcachelocation "Diskcachelocation" on page 359	include.fs
dynamicimage "Dynamicimage" on page 375	include.image
postsnapshotcmd "Postsnapshotcmd" on page 476	include.image
presnapshotcmd "Presnapshotcmd" on page 483	include.image
snapshotcachesize "Snapshotcachesize" on page 525	include.image
snapshotproviderfs "Snapshotproviderfs" on page 526	include.image
snapshotproviderimage "Snapshotproviderimage" on page 527	include.image

Examples

Options file:

```
include /home/proj/text/devel.*
include /home/proj/text/* textfiles
include * managall
include /WAS_ND_NDNODE mgmtclass
include /WAS_APPNODE mgmtclass
include.image /home
include.archive /home/proj/text/
* myarchiveclass
```

```

include.backup /home/proj/text/
* mybackupclass
include.compression /home/proj/text/
devel.*
include.encrypt /home/proj/gordon/*
include.fs.nas netappsj/vol/vol0
homemgmtclass

include.dedup /Users/Administrator/Documents/Important/.../*
AIX only:

include.image /home
MGMTCLASSNAME
snapshotproviderimage=JFS2
snapshotcachesize=40
include.image /home
snapshotproviderimage=NONE
include.fs /kalafs1
snapshotproviderfs=JFS2
Linux only:

include.image /home
snapshotproviderimage=LINUX_LVM
include.image /myfs1 dynamicimage=yes
include.image /home MGMTCLASSNAME
snapshotproviderimage=NONE
include.image /myfs1 dynamicimage=yes
include.attribute.symlink /home/spike/.../*
include.fs /usr
memoryefficientbackup=diskcachemethod

```

Command line:

Does not apply.

Related reference:

“Snapshotcachesize” on page 525

“Toc” on page 552

Related information:

-  [mmbackup command: Tivoli Storage Manager requirements](#)
-  [Guidance for integrating IBM Spectrum Scale AFM with IBM Spectrum Protect](#)
-  [Using IBM Spectrum Protect include and exclude options with IBM Spectrum Scale mmbackup command](#)

Controlling symbolic link and alias processing

Tivoli Storage Manager treats symbolic links and aliases (aliases apply to Mac OS X only) as actual files and backs them up. However, the file referenced by the symbolic link is not backed up.

In some cases symbolic links and aliases can be easily recreated and need not be backed up. In addition, backing up these symbolic links or aliases can increase backup processing time and occupy a substantial amount of space on the Tivoli Storage Manager server.

You can use the `exclude.attribute.symlink` option to exclude a file or a group of files that are symbolic links or aliases from backup processing. If necessary, you can use the `include.attribute.symlink` option to include symbolic links or aliases within broad group of excluded files for backup processing. For example, to exclude all symbolic links or aliases from backup processing, except those that exist under the `/home/spike` directory, enter these statements in your `dsm.sys` file:


```
exclude.attribute.symlink /.../*
include.attribute.symlink /home/spike/.../*
```

Related reference:

“Exclude options” on page 389

Compression and encryption backup processing

This topic lists some items to consider if you want to include specific files or groups of files for compression and encryption processing during a backup or archive operation.

- You must set the compression option to *yes* to enable compression processing. If you do not specify the compression option or you set the compression option to *no*, the backup-archive client does not perform compression processing.
- The client processes `exclude.fs`, `exclude.dir`, and other include-exclude statements first. The client then considers any `include.compression` and `include.encrypt` statements. For example, consider the following include-exclude list:

```
exclude /home/jones/proj1/file.txt
include.compression /home/jones/proj1/file.txt
include.encrypt /home/jones/proj1/file.txt
```

The client examines the `exclude /home/jones/proj1/file.txt` statement first and determines that `/home/jones/proj1/file.txt` is excluded from backup processing and is, therefore, not a candidate for compression and encryption processing.

- Include-exclude compression and encryption processing is valid for backup and archive processing *only*.

Related reference:

“Compression” on page 339

Processing NAS file systems

Use the `include.fs.nas` option to bind a management class to NAS file systems and to control whether Table of Contents information is saved for the file system backup.

Note: The `include.fs.nas` option does not apply to incremental snapshot difference incremental backup.

A NAS file system specification uses the following conventions:

- NAS nodes represent a new node type. The NAS node name uniquely identifies a NAS file server and its data to the backup-archive client. 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 operating system, NAS file system specifications use the forward slash (/) separator, as in this example: `/vol1/vol10`.

Use the following syntax:

➤—*pattern*— *mgmtclassname*— *toc=value*—➤

Where:

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

For more information, see “Toc” on page 552.

Example 1: To assign a management class to the /vol/vol1 file system of a NAS node that is called netappsj, specify the following include statement:

```
include.fs.nas netappsj/vol/vol1 nasMgmtClass toc=yes
```

Example 2: To assign the same management class to all paths that are subordinate to the /vol/ file system on a NAS node called netappsj (for example, /vol/vol1, /vol/vol2, and /vol/vol3), specify the following include statement:

```
include.fs.nas netappsj/vol/* nasMgmtClass toc=yes
```

Virtual machine include options

Virtual machine include and exclude options influence the behavior of backup and restore operations for virtual machines. These options are processed before any command-line options are processed, so that options on the command line can override options specified on any of the virtual machine include options or virtual machine exclude options. See the individual option descriptions for information about the options.

Related reference:

“Include.vmdisk” on page 431

“INCLUDE.VMTSMVSS” on page 435

“INCLUDE.VMSNAPSHOTATTEMPTS” on page 433

Include.vapp:

When you back up a VMware vCloud vApp, the backed up data is bound to a management class that is specified by the vappmc option. You can use the include.vapp option to override the management class that is specified by vappmc and bind the backed up data for a vApp to a different management class.

Supported Clients

This option can be used with supported Linux clients that are configured to back up VMware vCloud vApps.

Options File

Set this option in the client options file.

Syntax

►►—INCLUDE.VAPP— *vapp_specification*— *mgmtclassname*—►►

Parameters

vapp_specification

Required parameter. Specifies the organization, organization virtual data center, and vApp to bind to the specified management class when you backup a vApp.

The vApp specification must be in the following form:

org_name::orgvdc_name::vapp_name

where:

org_name::

Specifies the name of the organization that contains the vApp. The organization name must be followed by two colon (:) characters.

orgvdc_name::

Specifies the name of the organization virtual data center that contains the vApp. The organization virtual data center name must be followed by two colon (:) characters.

vapp_name

Specifies the name of a vApp.

Only one vApp can be specified on each `include.vapp` statement. However, you can specify as many `include.vapp` statements as needed to bind each vApp to a specific management class.

You can include wildcards in the vApp specification to match more than one organization, organization virtual data center, or vApp. An asterisk (*) matches any character string. A question mark (?) matches a single character. If any part of the vApp specification contains a space character, enclose the entire specification in double quotation marks (").

mgmtclassname

Optional parameter. Specifies the management class to use when the specified vApp is backed up. If this parameter is not specified, the management class defaults to the global vApp management class that is specified by the `vappmc` option.

Examples

Assume that the following management classes exist and are active on the Tivoli Storage Manager server:

- MCFORTESTVAPPS
- MCFORPRODVAPPS
- MCUNIQUEVAPP

The following `include.vapp` statement in the client options file binds all vApp backups to the management class called MCFORTESTVAPPS. Wildcards are used in the vApp specification to target all vApps in all organizations and all organization virtual data centers:

```
include.vapp *:::*::* vm vmtest* MCFORTESTVAPPS
```

The following `include.vapp` statement in the client options file binds all vApps in ORG1 and ORG1VDC to the management class called MCFORPRODVAPPS:

```
include.vapp "ORG1::ORGVDC1::*" MCFORPRODVAPPS
```

The following `include.vapp` statement in the client options file binds a vApp that is named VAPP1 to a management class that is named MCUNIQUEVAPP:

```
include.vapp MY_ORG::MY_VDC::VAPP1 MCUNIQUEVAPP
```

Related information

“Vappmc” on page 561

Include.vm:

For Hyper-V or VMware operations, this option overrides the management class that is specified on the `vmmc` option.

The management class specified on the `vmmc` option applies to all Hyper-V and VMware backups. You can use the `include.vm` option to override that management class, for one or more virtual machines. The `include.vm` option does not override or affect the management class that is specified by the `vmctlmc` option. The `vmctlmc` option binds backed-up virtual machine control files to a specific management class.

Supported Clients

This option can be used with supported Linux clients that are configured to back up VMware virtual machines.

Options File

Set this option in the client options file.

Syntax

►►—INCLUDE.VM— *vmname*— *mgmtclassname*—►

Parameters

vmname

Required parameter. Specifies the name of a virtual machine that you want to bind to the specified management class. The name is the virtual machine display name.

Only one virtual machine can be specified on each `include.vm` statement. However, you can specify as many `include.vm` statements as needed to bind each virtual machine to a specific management class.

You can include wildcards in the virtual machine name. An asterisk (*) matches any character string. A question mark (?) matches a single character. If the virtual machine name contains a space character, enclose the name in double quotation marks (").

Tip: If the virtual machine name contains special characters, type the question mark wildcard in place of the special characters when you specify the virtual machine name.

mgmtclassname

Optional parameter. Specifies the management class to use when the specified

virtual machine is backed up. If this parameter is not specified, the management class defaults to the global virtual machine management class that is specified by the `vmmc` option.

Examples

Assume that the following management classes exist and are active on the Tivoli Storage Manager server:

- MCFORTESTVMS
- MCFORPRODVMS
- MCUNIQUEVM

Example 1

The following `include.vm` statement in the client options file binds all virtual machines that have names that begin with VMTEST to the management class called MCFORTESTVMS:

```
include.vm vmtest* MCFORTESTVMS
```

Example 2

The following `include.vm` statement in the client options file binds a virtual machine that is named WHOPPER VM1 [PRODUCTION] to the management class called MCFORPRODVMS:

```
include.vm "WHOPPER VM1 ?PRODUCTION?" MCFORPRODVMS
```

The virtual machine name must be enclosed in quotation marks because it contains space characters. Also, the question mark wildcard is used to match the special characters in the virtual machine name.

Example 3

The following `include.vm` statement in the client options file binds a virtual machine that is named VM1 to a management class that is named MCUNIQUEVM:

```
include.vm VM1 MCUNIQUEVM
```

Related information

“Vmmc” on page 584

Include.vmdisk:

The `INCLUDE.VMDISK` option includes a VMware virtual machine disk in backup operation. If you do not specify one or more disk labels, all disks in the virtual machine are backed up.

This option is available only if you are using the Tivoli Storage Manager for Virtual Environments licensed product. For additional information about this option, see the Tivoli Storage Manager for Virtual Environments product documentation on IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS8TDQ/welcome>.

The `INCLUDE.VMDISK` option specifies the label of a virtual machine's disk to be included in a **Backup VM** operation. If you include a disk on the **Backup VM** command, the command-line parameters override any `INCLUDE.VMDISK` statements in the options file.

Supported clients

This option can be used with supported x86_64 Linux clients.

Options file

Set this option in the client options file. Command line parameters override statements in the options file.

Syntax

►►—INCLUDE.VMDISK—*vmname*——*vmdk_label*——►►

Parameters

vmname

Specifies the name of the virtual machine that contains a disk that you want to include in a **Backup VM** operation. The name is the virtual machine display name. You can specify only one virtual machine name on each INCLUDE.VMDISK statement. Specify additional INCLUDE.VMDISK statements for each virtual machine disk to include.

The virtual machine name can contain an asterisk (*), to match any character string, and question mark (?) to match any one character. Surround the VM name with quotation marks (" ") if the VM name contains space characters.

Tip: If the virtual machine name contains special characters, such as bracket characters ([or]), the virtual machine name might not be correctly matched. If a virtual machine name uses special characters in the name, you might need to use the question mark character (?) to match the special characters in the VM name

For example, to include Hard Disk 1 in the backup of a virtual machine named "Windows VM3 [2008R2]", use this syntax in the options file: INCLUDE.VMDISK "Windows VM3 ?2008R2?" "Hard Disk 1"

vmdk_label

Specifies the disk label of the disk that you want to include. Wildcard characters are not allowed. Use the **Backup VM** command with the -preview option to determine the disk labels of disks in a given virtual machine. See "Backup VM" on page 634 for the syntax.

Examples

Options file

Assume that a virtual machine named vm1 contains four disks, labeled Hard Disk 1, Hard Disk 2, Hard Disk 3, and Hard Disk 4. To include only disk 2 in a **Backup VM** operations, specify the following in the options file:

```
INCLUDE.VMDISK "vm1" "Hard Disk 2"
```

Include disks 2 and 3 in **Backup VM** operations:

```
INCLUDE.VMDISK "vm1" "Hard Disk 2"  
INCLUDE.VMDISK "vm1" "Hard Disk 3"
```

Command line

Include a single disk when backing up vm1:

```
dsmc backup vm "vm1:vmdk=Hard Disk 1"
```

Include disk 2 and disk 3 on vm1:

```
dsmc backup vm "vm1:vmdk=Hard Disk 2:vmdk=Hard Disk 3"
```

Related reference:

“Backup VM” on page 634

“Restore VM” on page 715

“Domain.vmfull” on page 369

“Exclude.vmdisk” on page 394

INCLUDE.VMSNAPSHOTATTEMPTS:

For VMware backup operations, this option determines the total number of snapshot attempts to try for a VMware virtual machine that fails during backup processing due to snapshot failure.

Supported Clients

This option can be used with supported Linux clients that are configured to back up VMware virtual machines.

Options File

Set this option in the client options file.

Syntax

```
►—INCLUDE.VMSNAPSHOTATTEMPTS—vmname—num_with_quiescing—————►  
►—num_without_quiescing—————►
```

Parameters

vmname

A required positional parameter that specifies the name of the virtual machine to attempt the total number of snapshots for, if a backup attempt fails due to snapshot failure. The name is the virtual machine display name.

Only one virtual machine can be specified on each INCLUDE.VMSNAPSHOTATTEMPTS statement. However, to configure the total snapshot attempts for other virtual machines, you can use the following methods:

- Specify as many INCLUDE.VMSNAPSHOTATTEMPTS statements as you need to, to have failed snapshots reattempted, for each virtual machine that you want this option to apply to.
- Use wildcard characters in **vmname** to select virtual machine names that match the wildcard pattern. An asterisk (*) matches any character string. A question mark (?) matches a single character. If the virtual machine name contains a space character, enclose the name in double quotation marks (").

Tip: If the virtual machine name contains special characters, type the question mark wildcard (?) in place of the special characters when you specify the virtual machine name.

num_with_quiescing

This is a positional parameter that specifies the following action:

- For Windows virtual machines with Tivoli Storage Manager application protection enabled, *num_with_quiescing* specifies the number of times to attempt the snapshot with Tivoli Storage Manager VSS quiescing. VSS quiescing applies only to Windows virtual machines.
- For Windows and Linux virtual machines without Tivoli Storage Manager application protection enabled, *num_with_quiescing* specifies the number of times to attempt the snapshot with VMware Tools file system quiescing.

The maximum value that you can specify is ten (10). The default value is two (2). The minimum value that you can specify is zero (0).

num_without_quiescing

This is a positional parameter that specifies the number of times to attempt the snapshot with VMware Tools file system quiescing and application (VSS) quiescing disabled after the specified number of attempts with VSS quiescing (*num_with_quiescing*) completes. For example, you can specify this parameter for a virtual machine that is already protected by an IBM Data Protection agent installed in a guest virtual machine.

The maximum value that you can specify is ten (10). The minimum value that you can specify is zero (0), which is the default value.

Important: When this parameter is applied to a virtual machine backup, the backup is considered crash-consistent. As a result, operating system, file system, or application consistency is not guaranteed. An `include.vmsnapshotattempts 0 0` entry is not valid. Backup operations require at least one snapshot.

Examples

Example 1

The following `INCLUDE.VMSNAPSHOTATTEMPTS` statement in the client options file tries two total snapshot attempts (with VSS quiescing) for virtual machine `VM_a`:

```
INCLUDE.VMSNAPSHOTATTEMPTS VM_a 2 0
```

Example 2

The following `INCLUDE.VMSNAPSHOTATTEMPTS` statement in the client options file tries three total snapshot attempts for virtual machines that match the `vmServer_Dept*` string:

- First and second attempts to take a snapshot with VSS quiescing.
- Third snapshot attempt without VSS quiescing.

```
INCLUDE.VMSNAPSHOTATTEMPTS vmServer_Dept* 2 1
```

Example 3

The following `INCLUDE.VMSNAPSHOTATTEMPTS` statement in the client options file tries one total snapshot attempt (with VSS quiescing) for virtual machines that match the `vmDB_Dept*` string:

```
INCLUDE.VMSNAPSHOTATTEMPTS vmDB_Dept* 1 0
```

Example 4

The following `INCLUDE.VMSNAPSHOTATTEMPTS` statement in the client options file tries two total snapshot attempts (with VSS quiescing) for all virtual machines:

```
INCLUDE.VMSNAPSHOTATTEMPTS * 2 0
```


Example 5

In this example, the virtual machine DB15 has an IBM Data Protection agent installed in a guest virtual machine and does not need an application-consistent snapshot. The following `INCLUDE.VMSNAPSHOTATTEMPTS` statement in the client options file tries one total snapshot attempt (without VSS quiescing) for virtual machine DB15:

```
INCLUDE.VMSNAPSHOTATTEMPTS DB15 0 1
```

INCLUDE.VMTSMVSS:

The `INCLUDE.VMTSMVSS` option notifies virtual machine applications that a backup is about to occur. This option allows the application to truncate transaction logs and commit transactions so that the application can resume from a consistent state when the backup completes. An optional parameter can be specified to suppress truncation of the transaction logs. To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

When a virtual machine is included by this option, Tivoli Storage Manager provides application protection. That is, the client freezes and thaws the VSS writers and, optionally, truncates the application logs. If a virtual machine is not protected by this option, application protection is provided by VMware, and VMware freezes and thaws the VSS writers, but application logs are not truncated.

Supported clients

This option can be used with supported x86_64 Linux clients.

Options file

Set this option in the client options file. This option cannot be set by the preferences editor or on the command line.

Syntax

►► `INCLUDE.VMTSMVSS` *vmname* `—OPTions=KEEPSqllog` ►►

Parameters

vmname

Specifies the name of the virtual machine that contains the applications to quiesce. The name is the virtual machine display name. Specify one virtual machine per `INCLUDE.VMTSMVSS` statement. For example, to include a virtual machine named Windows VM3 [2008R2], use this syntax in the options file: `INCLUDE.VMTSMVSS "Windows VM3 [2008R2]"`.

To protect all virtual machines with this option, use an asterisk as a wildcard (`INCLUDE.VMTSMVSS *`). You can also use question marks to match any single character. For example, `INCLUDE.VMTSMVSS vm??` protects all virtual machines that have names that begin with `vm` and are followed by any two characters (`vm10`, `vm11`, `vm17`, and so on).

Tip: If the virtual machine name contains special characters, such as bracket characters ([or]), the virtual machine name might not be correctly matched. If a virtual machine name uses special characters in the name, you can use the question mark character (?) to match the special characters in the virtual

machine name.

There is no default value for this parameter. To enable application protection, you must include virtual machines to be protected on one or more `INCLUDE.VMTSMVSS` statements. Make sure that you do not exclude a disk on a virtual machine (by using the `EXCLUDE.VMDISK` option) if the disk contains application data that you want protected.

OPTions=KEEPSqllog

If the `OPTions KEEPSqllog` parameter is specified on an `INCLUDE.VMTSMVSS` statement, the parameter prevents SQL server logs from being truncated when a backup-archive client that is installed on a data mover node backs up a virtual machine that is running a SQL server. Specifying this parameter allows the SQL server administrator to manually manage (backup, and possibly truncate) the SQL server logs, so that they can be preserved and be used to restore SQL transactions to a specific checkpoint, after the virtual machine is restored.

When this option is specified, the SQL log is not truncated and the following message is displayed and logged on the server:

```
ANS4179I IBM Tivoli Storage Manager application protection
did not truncate the Microsoft SQL Server logs on VM 'VM'.
```

You can remove the `OPTIONS=KEEPSQLLOG` option to enable truncation of the SQL logs when a backup completes.

Note: The client does not back up the SQL log files. The SQL administrator must back up the log files so that they can be applied after the database is restored.

Examples

Options file

Configure application protection for a virtual machine that is named `vm_example`:

```
INCLUDE.VMTSMVSS vm_example
```

Configure application protection for `vm11`, `vm12`, and `vm15`:

```
INCLUDE.VMTSMVSS vm11
INCLUDE.VMTSMVSS vm12
INCLUDE.VMTSMVSS vm15 options=keepsqlllog
```

Command line

Not applicable; this option cannot be specified on the command line.

Related reference:

“`Vmtimeout`” on page 598

“`Exclude.vmdisk`” on page 394

“`Include.vmdisk`” on page 431

Incrbydate

Use the `incrbydate` option with the **incremental** command to back up new and changed files with a modification date later than the last incremental backup stored at the server, unless you exclude the file from backup.

Important: Files that are modified or created after their respective directory was processed by the backup-archive 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 run a regular incremental backup periodically, without specifying the `incrbydate` option.

An incremental-by-date updates the date and time of the last incremental at the server. If you perform an incremental-by-date on only part of a file system, the date of the last full incremental is not updated and the next incremental-by-date backs up these files again.

Important:

The last incremental backup time refers to the server time and the file modification time refers to the client time. If the client and server time are not synchronized, or the client and server are in different time zones, this affects incremental-by-date backup with `mode=incremental`.

The last incremental backup time refers to the server time and the file modification time refers to the client time. If the client and server time are not synchronized, or the client and server are in different time zones, this affects incremental-by-date backup and image backup with `mode=incremental`.

Both full incremental backups and incrementals-by-date backups backup 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 backup, an incremental-by-date backup does not maintain current server storage of all your workstation files for the following reasons:

- It does not expire backup versions of files that are deleted from the workstation.
- It does not rebind backup versions to a new management class if the management class has changed.
- It does not back up files with attributes that have changed, such as Access control list (ACL) data, unless the modification dates and times have also changed.
- It ignores the copy group frequency attribute of management classes.

Tip: If you have limited time during the week to perform backups, but extra time on weekends, you can maintain current server storage of your workstation files by performing an incremental backup with the `incrbydate` option on weekdays and a full incremental backup on weekends.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—INCRbydate—►►

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc incremental -incrbydate
```

Incremental

Use the incremental option with the **restore image** command to ensure that any changes that were made to the base image are also applied to the restored image.

If you also use the deletefiles option, changes include the deletion of files and directories that were in the original image but later deleted from the workstation.

Note: Using the incremental option with the **restore image** command to perform a dynamic image backup is not supported.

Supported Clients

This option is valid only for AIX, HP-UX, Linux x86_64, Linux on POWER, and Solaris. The Tivoli Storage Manager API does not support this option.

Syntax

►—INCRemental—►

Examples

Command line:

```
res i "/home/devel/projecta/*" -incremental
```

Instrlogmax

The instrlogmax option specifies the maximum size of the instrumentation log (dsminstr.log), in MB. Performance data for the client is collected in the dsminstr.log file during backup or restore processing when the enableinstrumentation option is set to *yes*.

If you change the value of the instrlogmax 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.

Supported Clients

This option is valid for all clients and the Tivoli Storage Manager API.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza.

Syntax

►—INSTRLOGMAX— *size*—►

Parameters

size

Specifies the maximum size, in MB, for the instrumentation log file. The range of values is 0 - 2047. The default value is 25.

When the size of the `dsminstr.log` file exceeds the maximum size, the log file is renamed to `dsminstr.log.bak`. Subsequent instrumentation data continues to be saved to the `dsminstr.log` file.

If you specify 0, the log file grows indefinitely.

Examples

Options file:

```
instrlogmax 100
```

Command line:

```
dsmc sel /home/mydir/* -subdir=yes -enableinstrumentation=yes  
-instrlogmax=100
```

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

Related reference:

“Enableinstrumentation” on page 380

“Instrlogname”

Instrlogname

The `instrlogname` option specifies the path and file name where you want to store performance information that the backup-archive client collects.

When you use the `enableinstrumentation yes` option to collect performance data during backup and restore operations, the client automatically stores the information in a log file.

By default, the performance data is stored in the instrumentation log file (`dsminstr.log`) in the directory that is specified by the `DSM_LOG` environment variable (or the `DSMI_LOG` environment variable for the API-dependent products Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server and Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server). If you did not set the `DSM_LOG` environment variable, the instrumentation log file is stored in the current directory (the directory where you issued the `dsmc` command).

Use this option only when you want to change the file name and location of the instrumentation log.

If you want to control the size of the log file, use the `instrlogmax` option.

Supported Clients

This option is valid for all clients and the Tivoli Storage Manager API.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza.

Important: Set the DSM_LOG environment variable to name a directory where the log is to be placed. The directory that is specified must have permissions that allow write-access from the account under which the client is run. The root directory is not a valid value for DSM_LOG.

Syntax

►—INSTRLOGNAME— *filespec*—►

Parameters

filespec

Specifies the path and file name where you want to store performance information during backup or restore processing. If any part of the path that you specify does not exist, the client attempts to create it.

If you specify a file name only, the file is stored in the directory that is specified by the DSM_LOG environment variable. If you did not set the DSM_LOG environment variable, the instrumentation log file is stored in the current directory (the directory where you issued the **dsmc** command). The instrumentation log file cannot be a symbolic link.

For Mac OS X, if you specify a file name only, the file is stored in your default folder. The default directories are:

```
~/Library/Logs/tivoli/tsm  
/Library/Logs/tivoli/tsm
```

This instrumentation log file name replaces the previous instrumentation log file name `dsminstr.report.pXXX` that was created by the `TESTFLAG=instrument:detail` or `instrument:API` option.

Examples

Options file:

For AIX, HP-UX, Linux, and Oracle Solaris clients:

```
instrlogname /home/user1/mydir/mydsminstr.log
```

For Mac OS X clients:

```
instrlogname /Users/user1/Library/Logs/mydsminstr.log
```

Command line:

For AIX, HP-UX, Linux, and Oracle Solaris clients:

```
dsmc sel /home/user1/mydir/* -subdir=yes -instrlogname=/usr/log/mydsminstr.log
```

For Mac OS X clients:

```
dsmc sel /Users/user1/mydir/* -subdir=yes -instrlogname=/Users/user1/Library/Logs/mydsminstr.log
```

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

Related reference:

“Enableinstrumentation” on page 380

“Instrlogmax” on page 438

Lanfreecommmethod

The `lanfreecommmethod` option specifies the communications protocol between the Tivoli Storage Manager client and Storage Agent. This enables processing between the client and the SAN-attached storage device.

If you are using LAN failover, you must have `lanfreecommmethod` in the `dsm.sys` file within a server stanza.

For AIX, HP-UX, Linux and Solaris, use the `lanfreeshmport` option to specify the shared memory port number where the Storage Agent is listening.

Supported Clients

This option is only valid for AIX, HP-UX, Linux, and Solaris clients.

Options File

Place this option in the `dsm.sys` file within a server stanza.

Syntax

►►—LANFREECommmethod— *commmethod* —►►

Parameters

commmethod

Specifies the supported protocol for the backup-archive client:

TCPip

The Transmission Control Protocol/Internet Protocol (TCP/IP) communication method.

Use the `lanfreetcpport` option to specify the TCP/IP port number where the Storage Agent is listening. The TCP/IP communication method is the default for non-root users on all supported platforms.

V6Tcpip

Indicates that either TCP/IP v4 or v6 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.

SHAREmem

Use the shared memory communication method when the client and Storage Agent are running on the same system. Shared memory provides better performance than the TCP/IP protocol. This is the default communication method for AIX, HP-UX, Linux, and Solaris root users. When specifying this communication method on AIX, the backup-archive client user can be logged in as root or non-root, as long as the Storage Agent is running as root. If the Storage Agent is not running as root, the user ID running the backup-archive client must match the user ID running the Storage Agent.

Examples

Options file:

```
lanfreecommmethod tcp
```

Use only TCP/IP v4

lanfreecommmethod V6Tcpip

Use both TCP/IP v4 or v6, depending on how the system is configured and the results of a domain name service lookup.

Command line:

-lanfreec=tcp

-lanfreec=V6Tcpip

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

Related information

"Lanfreeshmport"

"Lanfreecpport" on page 443

Lanfreeshmport

Use the lanfreeshmport option when lanfreecommmethod=SHAREdmem is specified for communication between the backup-archive client and the storage agent. This enables processing between the client and the SAN-attached storage device.

Supported Clients

This option is valid for AIX, HP-UX, Linux, and Oracle Solaris clients only.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza.

Syntax

►►—LANFREEShmport— *port_address*—————►►

Parameters

port_address

Specifies the number that is used to connect to the storage agent. The range of values is 1 through 32767.

For Windows clients, the default is 1.

For all clients except Windows clients, the default is 1510.

Examples

Options file:

lanfrees 1520

Command line:

-lanfrees=1520

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

Related information

“Lanfreecommmethod” on page 441

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 backup-archive client and Storage Agent. Do not specify the `lanfreetcpport` option if you want to use the `NAMEDpipes` communication method for LAN-free communication.

Supported Clients

This option is only valid for AIX, HP-UX, Linux x86_64, Linux on POWER, and Solaris clients.

Options File

Place this option in the `dsm.sys` file within a server stanza.

Syntax

►►—LANFREETCPPort— *port_address*—◄◄

Parameters

port_address

Specifies the TCP/IP port number where the Storage Agent is listening. The range of values is 1 through 32767; the default is 1500.

Note: The client `lanfreetcpport` value must match Storage Agent `tcpport` value for communications with the Storage Agent (virtual server). The client `tcpport` value must match the server `tcpport` value for communications with the actual server.

Examples

Options file:

```
lanfreetcpp 1520
```

Command line:

```
-lanfreetcpp=1520
```

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

Related information

“Lanfreecommmethod” on page 441

Lanfreessl

Use the `lanfreessl` option to enable Secure Sockets Layer (SSL), to provide secure client and Storage Agent communications. This option is deprecated if you are connecting to a Tivoli Storage Manager server V7.1.8 and later V7 levels, and V8.1.2 and later levels.


Supported Clients

This option is supported on all clients, except for Mac OS X clients.

Options File

Place this option in the client options file. You cannot set this option in the GUI or on the command line.

Syntax



Parameters

No Specifies that the backup-archive client does not use SSL when communicating with the Storage Agent. No is the default.

Yes

Specifies that the backup-archive client enables SSL when communicating with the Storage Agent. To enable SSL, specify `lanfreessl=yes` and change the value of the `lanfreetcppport` option. Changing the value of the `lanfreetcppport` option is necessary because the Tivoli Storage Manager Storage Agent is typically set up to listen for SSL connections on a separate port.

Examples

Options file:

```
lanfreessl yes
lanfreessl no
```

Command line:

Not applicable. You cannot set this option on the command line.

Lanfreetcpserveraddress

The `lanfreetcpserveraddress` option specifies the TCP/IP address for a Tivoli Storage Manager Storage Agent.

Use this option when you specify `lanfreecommmethod=TCPIP` or `V6TCPIP` for communication between the backup-archive 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 only for AIX, HP-UX, Linux x86_64, Linux on POWER, and Solaris clients.

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.example.com
```

```
LANFREETCPServeraddress 192.0.2.1
```

Command line:

Does not apply.

Latest

Use the latest option to restore the most recent backup version of a file, even if the backup is inactive.

You can use the latest option with the following commands:

- **restore**
- **restore group**

If you are performing a point-in-time restore (using the pitdate option), it is not necessary to specify latest since this option is implicit when pitdate is used.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—LATest—►►

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/Users/devel/projecta/*" -latest
```

Command line:

```
dsmc restore "/home/devel/projecta/*" -latest
```

Localbackupset

The `localbackupset` option specifies whether the backup-archive client GUI bypasses initial logon with the Tivoli Storage Manager server to restore a local backup set on a standalone workstation.

If you set the `localbackupset` option to `yes`, the GUI does not attempt initial logon with the server. In this case, the GUI only enables the restore functionality.

If you set the `localbackupset` option to `no` (the default), the GUI attempts initial logon with the server and enables all GUI functions.

Note: The `restore backupset` command supports restore of local backup sets on a standalone workstation without using the `localbackupset` option.


Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza.

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 **dsmc** command-line client. (Does not apply to Mac OS X)

Related information

“Restore Backupset” on page 698

Makesparsefile

Use the `makesparsefile` option with the **restore** or **retrieve** commands to specify how sparse files are recreated.

Sparse files do not have disk space allocated for every block in the whole address space, leading to holes within the file. The backup-archive client detects sparse files during a backup operation and marks them as sparse on the Tivoli Storage Manager server. Holes are detected by their content, which is always zeros.

If you set the `makesparsefile` option to `yes` (default), holes within the file are not written to disk so no additional disk space is allocated during a restore.

If you set the `makesparsefile` option to `no`, holes are not recreated, leading to disk blocks allocated for the whole address space. This might result in a larger amount of used disk space. Ensure that you have enough disk space to restore all data.

On some UNIX and Linux systems, it might be necessary to back up system specific files as non-sparse files. Use the `makesparsefile` option for files where the existence of physical disk blocks is required, such as `ufsboot` on Solaris, which is executed during boot time. The boot file loader of the operating system accesses physical disk blocks directly and does not support sparse files.


Supported Clients

This option is valid for all UNIX and Linux clients except Mac OS X.

Options File

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

Syntax



Parameters

Yes

Specifies that holes within the file are not written so that no additional disk space is allocated during a restore. This is the default.

No

Specifies that holes are not recreated leading to disk blocks allocated for the whole address space.

Examples

Options file:

```
makesparsefile no
```

Command line:

```
-makesparsefile=no
```

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

Managedservices

The managedservices option specifies whether the Tivoli Storage Manager client acceptor service manages the scheduler, the web client, or both.

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

The client acceptor daemon serves as an external timer for the scheduler. When the scheduler is started, it queries the server for the next scheduled event. The event is either executed immediately or the scheduler exits. The client acceptor daemon restarts the scheduler when it is time to execute the scheduled event.

Note:

1. If you set the schedmode option to prompt, the server prompts the client acceptor daemon when it is time to run the schedule. The scheduler connects to and disconnects from the server when the client acceptor daemon is first started.
The dsmc schedule command cannot be used when both schedmode prompt and commethod V6Tcpip are specified.
2. For Mac OS X, if you do not specify the managedservices option, the client acceptor daemon manages both the scheduler program and the web client, by default.
3. Set the passwordaccess option to generate in your dsm.sys file and generate a password, so Tivoli Storage Manager can manage your password automatically.

Using the client acceptor daemon to manage the scheduler service can provide the following benefits:

- Memory retention problems that can occur when using traditional methods of running the scheduler are resolved. Using the client acceptor daemon to manage the scheduler requires very little memory between scheduled operations.
- The client acceptor daemon can manage both the scheduler program and the web client, reducing the number of background processes on your workstation.
- To use the web client, you must specify this option in the client system-options file.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza. You can set this option on the **Web Client** tab of the Preferences editor.

Syntax

►►—MANAGEDServices—mode—————◄◄

Parameters

mode

Specifies whether the client acceptor daemon manages the scheduler, the web client, or both.

webclient

Specifies that the client acceptor daemon manages the web client. This is the default for UNIX and Linux. Both *webclient* and *schedule* are the defaults for Mac OS X.

schedule

Specifies that the client acceptor daemon manages the scheduler. Both *webclient* and *schedule* are the defaults for Mac OS X.

none

For Mac OS X, specifies that the client acceptor daemon not manage the web client or schedules. Set *managedservices* to *none* to enable the **dsmc schedule** command.

Examples

Options file:

The following are examples of how you might specify the *managedservices* option in your client system-options file (*dsm.sys*).

Task Specify that the client acceptor daemon manages only the web client.

```
managedservices webclient
```

Task Specify that the client acceptor daemon manages only the scheduler.

```
managedservices schedule
```

Task Specify that the client acceptor daemon manages both the web client and the scheduler.

```
managedservices schedule webclient
```

Note: The order in which these values are specified is not important.

Task For Mac OS X, to use the **dsmc schedule** command, specify:

```
managedservices none
```

Command line:

Does not apply.

Related information

“Passwordaccess” on page 470

See “Configuring the scheduler” on page 61 for instructions to set up the client acceptor daemon to manage the scheduler.

“Sessioninitiation” on page 514

“Cadlistenonport” on page 332

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 administrator specifies a value for this option, that value overrides what you specify in the client options file after your client node successfully contacts the server.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **Scheduler** tab, in the **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.

Mbobjrefreshthresh

The `mbobjrefreshthresh` (megablock object refresh threshold) option is a number defining a threshold. When the number of Tivoli Storage Manager objects that are needed to describe any 128 MB megablock exceeds this value, the entire megablock is refreshed and the objects that were used to represent this area, in previous backups, are expired.

When you backup a virtual machine, the data is stored on the Tivoli Storage Manager server in 128 MB units, called *megablocks*. If an area on the production disk changes and a new incremental backup is performed, a new megablock is created to represent the changes that were made to the previously backed up data. Because a new megablock can be created with each incremental backup, eventually the megablocks can adversely affect the performance of the Tivoli Storage Manager database, and therefore, adversely affect the performance of most Tivoli Storage Manager operations.

Use this option when estimating Tivoli Storage Manager objects that represent production data for each virtual machine backup. For example, when the number of Tivoli Storage Manager objects exceed this value, the megablock is refreshed. This action means that the entire 128-MB block is copied to the server and is represented as a single Tivoli Storage Manager object. The minimum value is 2 and the maximum value is 8192. The default value is 50.

Supported clients

This option is valid for clients that act as data mover nodes that protect VMware virtual machines. To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

Options file

Specify this option in the client options file.

Syntax



Parameters

The minimum value you can specify is 2 megablocks, the largest value is 8192 megablocks; the default is 50 megablocks.

Examples

Set this option to trigger a megablock refresh when the number of objects needed to represent an updated megablock exceeds 20 objects:

```
MBOBJREFRESHTHRESH 20
```

Mbpctrefreshthresh

The `mbpctrefreshthresh` (megablock percentage refresh threshold) option is a number defining a threshold. When the percentage of Tivoli Storage Manager objects that are needed to describe any 128 MB megablock exceeds this value, the entire megablock is refreshed and the objects that were used to represent this area, in previous backups, are expired.

When you backup a virtual machine, data is stored on the Tivoli Storage Manager server in 128 MB units, called *megablocks*. If an area on the production disk changes and a new incremental backup is performed, a new megablock is created to represent the changes that were made to the previously backed up data. Because a

new megablock can be created with each incremental backup, eventually the megablocks can adversely affect the performance of the Tivoli Storage Manager database, and therefore, adversely affect the performance of most Tivoli Storage Manager operations.

Use this option when estimating the amount of additional data that is backed up for each virtual machine. For example, when a 128-MB block of a production disk changes more than the percentage specified, the entire 128-MB block is copied to the server. The block is represented as a single Tivoli Storage Manager object.

Supported clients

This option is valid for clients that act as data mover nodes that protect VMware virtual machines. To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

Options file

Specify this option in the client options file.

Syntax



Parameters

The minimum value you can specify is 1 percent, the largest value is 99 percent; the default is 50 percent.

Examples

Set this option to trigger a megablock refresh when 50 percent (or more) of the objects in a megablock on a production disk have changed:

```
MBPCTREFRESHTHRESHOLD 50
```

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 that the backup-archive client uses on a per-filespace basis.

Use `memoryefficientbackup=diskcachemethod` for any file space that has too many files for the client to complete the incremental backup with either the default setting, `memoryefficientbackup=no`, or with `memoryefficientbackup=yes`.

The actual amount of disk space required for the disk cache file created by disk cache incremental backups depends on the number of files and directories included in the backup and on the average path length of the files and directories to be backed up. For UNIX and Linux estimate 1 byte per character in the path name. For Mac OS X, estimate 4 bytes per character in the path name. For example, if there are 1 000 000 files and directories to be backed up and the average path length is 200 characters, then the database occupies approximately 200 MB for UNIX and Linux, and 800 MB for Mac OS X clients. Another way to estimate for planning purposes is to multiply the number of files and directories by the length of the longest path to establish a maximum database size.

A second disk cache file is created for the list of migrated files when backing up an HSM managed file system. The combined disk cache files, created by disk cache incremental backups and HSM managed file system backups, can require above 400 MB of disk space for each million files being backed up. The disk cache file can become very large. Large file support must be enabled on the file system that is being used for the disk cache file.


Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

This option is allowed in dsm.opt and within a server stanza in dsm.sys, but the value in dsm.opt is ignored if it also appears in dsm.sys. You can also place this option on the initial command line. In interactive mode, this option can be used with the **incremental** command. You can also set this option on the **Performance Tuning** tab in the Preferences editor, and selecting the **Use memory-saving algorithm** check box.

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

Related information

“Include options” on page 421

Mode

Use the mode option to specify the backup mode to use when performing specific backup operations.

The mode option has no effect on a when backing up a raw logical device.

You can use the mode option with the following backup commands:

backup image

To specify whether to perform a selective or incremental image backup of client file systems.

backup nas

To specify whether to perform a full or differential image backup of NAS file systems.

backup group

To specify whether to perform a full or differential group backup containing a list of files from one or more file space origins.

backup vm

For VMware virtual machines, this parameter specifies whether to perform a full, incremental, incremental-forever-full, or incremental-forever-incremental backup of VMware virtual machines.

Supported Clients

This option is valid on all supported clients, except Mac OS. The Tivoli Storage Manager API does not support this option.

Syntax

For image backups of client file systems



For image backup of NAS file systems



For group backups



For backing up VMware virtual machines



Parameters

Image backup parameters

selective

Specifies that you want to perform a full (selective) image backup. This is the default mode for image backups of client file systems.

incremental

Specifies that you want to back up only the data that has changed since the most recent image backup. If an image backup has not already been created, then the first backup is a full image backup (mode=selective), regardless of what mode option you specify.

NAS backup parameters

differential

This is the default for NAS objects. Specifies that you want to perform a NAS 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 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.

full

Specifies that you want to perform a full backup of NAS file systems.

Group backup parameters

full

Specifies that you want to perform a full backup of group objects. This is the default for group backups.

differential

Specifies that you want to perform a 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 GROUP server command.

A full image can be eligible for expiration based on versioning or retention (verexists retextra), but still be maintained on the 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 GROUP server command. The differential image backups that depend on an "expired" full image can be restored.

VMware virtual machine parameters

full

Specifies that you want to perform an image backup of all objects on a VMware virtual machine's disks. This is the default mode, and the only mode that you can use for VMware virtual machine backups, if you do not have a license for Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

IFFull

Specifies that you want to perform an incremental-forever-full backup of a virtual machine. An incremental-forever-full backup backs up all used blocks on a VMware virtual machine's disks. This mode requires a Tivoli Storage Manager for Virtual Environments: Data Protection for VMware license. By default, when the license for Tivoli Storage Manager for Virtual Environments: Data Protection for VMware is detected, the first backup of a VMware virtual machine is an incremental-forever-full (mode=iffull) backup, even if you specify mode=ifincremental (or let the mode option default). Subsequent backups default to mode=ifincremental.

You cannot use this backup mode to back up a virtual machine if the client is configured to encrypt the backup data.

For a description of the incremental-forever backup strategy for VMware virtual machines, see Backup and restore types.

IFIncremental

Specifies that you want to perform an incremental-forever-incremental backup of a virtual machine. An incremental-forever-incremental backup backs up only the disk blocks that have changed since the last backup. This mode requires a Tivoli Storage Manager for Virtual Environments: Data Protection for VMware license. With that license installed, this is the default backup mode for VMware virtual machine backups.

You cannot use this backup mode to back up a virtual machine if the client is configured to encrypt the backup data.

incremental

Specifies that you want to back up only new and changed data after the most recent image backup (full or incremental). If an image backup has not yet been created, then the first backup is a full image backup, regardless of the mode option. This mode requires a Tivoli Storage Manager for Virtual Environments: Data Protection for VMware license.

Examples

Task Perform a full VM incremental backup of a VMware virtual machine named vm1.

```
dsmc backup vm vm1 -mode=incremental  
-vmbackuptype=full
```

Task Perform a backup of a VMware virtual machine named `vm1`, using the incremental-forever-incremental mode to back up only the data that has changed since the last backup.

```
dsmc backup vm vm1 -mode=ifincremental  
-vmbackuptype=full
```

Task Perform the NAS image backup of the entire file system.

```
dsmc backup nas -mode=full -nasnodename=nas1  
/vol/vol0 /vol/vol1
```

Task Back up the `/home/test` file space using an image incremental backup that backs up only new and changed files after the last full image backup.

```
dsmc backup image /home/test -mode=incremental -snapshotproviderimage=none
```

Task Perform a full backup of all the files in filelist `/home/dir1/filelist1` to the virtual file space name `/virtfs` containing the group leader `/home/group1` file.

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

Related reference:

“Backup VM” on page 634

“Backup Group” on page 618

“Backup Image” on page 620

“Backup NAS” on page 626

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`, the backup-archive client 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`, the client does not monitor the current NAS image backup or restore operation and is available to process the next command.

Use this option with the **backup nas** or **restore nas** commands.

Supported Clients

This option is valid for AIX, Linux, and Solaris clients *only*.

Syntax

►► `MONitor` =  ►►

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
```

Myreplicationserver

The `myreplicationserver` option specifies which secondary server stanza that the client uses during a failover.

The secondary server stanza is identified by the `replservername` option and contains connection information about the secondary server.

This option is set by the Tivoli Storage Manager server administrator for the client node. During the normal (non-failover) logon process, the option is sent to the client and is saved in the client options file.

Do not edit this option during normal operations.

Edit this option only during situations such as the following ones:

- The primary server is offline and the information for the secondary server is not in the options file.
- The secondary server information is out-of-date or incorrect.

Any values that you edit are removed or updated the next time you log in to the primary server.

Supported Clients

This option is valid for all clients.

Options File

This option is placed within a server stanza in the `dsm.sys` file.

Syntax

►►—MYREPLICATIONServer—*repl_servername*—————►◄

Parameters

repl_servername

Specifies the name of the stanza for the secondary server to be used during a failover. This value is usually the name of the secondary server, not the host

name of the server. Also, the value of the repl_servername parameter is not case-sensitive, but the value must match the value that is specified for the REPLSERVERName option.

Examples

Options file:

```
MYREPLICATIONServer TargetReplicationServer1
```

Command line:

Does not apply.

Options file:

The following example demonstrates how to specify options for three different servers in the dsm.sys file, and how to reference the secondary server. Connection information for multiple secondary servers is presented in stanzas. Each stanza is identified by the **replservername** option and the name of the secondary server. The **servername** stanza must contain the **myreplicationserver** option, which points to the secondary server that is specified by the **replservername** stanza. Only one secondary server can be specified per **servername** stanza.

```
REPLSERVERNAME TargetReplicationServer1
REPLTCPSEVERADDRESS TargetReplicationServer1
REPLTCPPOINT 1505
REPLSSLPORT 1506
REPLSERVERGUID 91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00.00
```

```
REPLSERVERNAME TargetReplicationServer2
REPLTCPSEVERADDRESS TargetReplicationServer2
REPLTCPPOINT 1505
REPLSSLPORT 1506
REPLSERVERGUID 91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00.02
```

```
SErvername server_a
COMMMethod TCPip
TCPPOINT 1500
TCPSEVERADDRESS server_hostname1.example.com
PASSWORDACCESS prompt
MYREPLICATIONServer TargetReplicationServer1
```

```
SErvername server_b
COMMMethod TCPip
TCPPOINT 1500
TCPSEVERADDRESS server_hostname2.example.com
PASSWORDACCESS generate
INCLEXCL /adm/tsm/archive.excl
MYREPLICATIONServer TargetReplicationServer2
```

```
SErvername server_c
COMMMethod TCPip
TCPPOINT 1500
TCPSEVERADDRESS server_hostname3.example.com
PASSWORDACCESS generate
MYREPLICATIONServer TargetReplicationServer1
```

Related concepts:

“Automated client failover configuration and use” on page 87

Related tasks:

“Configuring the client for automated failover” on page 90

Nasnodename

The `nasnodename` option specifies the node name for the NAS file server when processing NAS file systems. The client prompts you for an administrator ID.

The node name identifies the NAS file server to the Tivoli Storage Manager server. The server must register the NAS file server.

You can specify this option on the command line or in the client system-options file (`dsm.sys`).

You can override the default value in the `dsm.sys` file by entering a different value on the command line. If you do not specify the `nasnodename` option in the `dsm.sys` file, you must specify this option on the command line when processing NAS file systems.

You can use the `nasnodename` option with the following commands:

- **backup nas**
- **delete filespace**
- **query backup**
- **query filespace**
- **restore nas**

You can use the **delete filespace** command to interactively delete NAS file spaces from server storage.

Use the `nasnodename` option to identify the NAS file server. Place the `nasnodename` option in your client system-options file (`dsm.sys`). The value in the client system-options file is the default, but this value can be overridden on the command line. If the `nasnodename` option is not specified in the client system-options file, you must specify this option on the command line when processing NAS file systems.

Use the `class` option to specify the class of the file space to delete. To display a list of file spaces belonging to a NAS node so that you can choose one to delete, use the `-class=nas` option.

To delete NAS file spaces using the web client, see the topic for backing up your data.

Supported Clients

This option is only valid for the AIX, Linux, and Solaris clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **General** tab of the Preferences editor.

Syntax

►►—NASNodename— —*nodename*—————►►

Parameters

nodename

Specifies the node name for the NAS file server.

Examples

Options file:

nasnodename nas2

Command line:

-nasnodename=nas2

Nfstimeout

The `nfstimeout` option specifies the number of seconds the client waits for a status system call on an NFS file system before it times out.

You can use this option to mitigate the default behavior of status calls on file systems. For example, if an NFS file system is stale, a status system call is timed out by NFS (soft mounted) or hang the process (hard mounted).

When the value of this option is changed to a value other than zero, a new thread is created by a caller thread to issue the status system call. The new thread is timed out by the caller thread and the operation can continue.

Note: On Solaris and HP-UX, the `nfstimeout` option can fail if the NFS mount is hard. If a hang occurs, deactivate the `nfstimeout` option and mount the NFS file system soft mounted, as follows:

```
mount -o soft,timeo=5,retry=5 machine:/filesystem /mountpoint
```

The parameters are defined as follows:

soft Generates a soft mount of the NFS file system. If an error occurs, the `stat()` function returns with an error. If the option `hard` is used, `stat()` does not return until the file system is available.

timeo=*n*

Sets the time out for a soft mount error to *n* tenths of a second.

retry=*n*

Set the internal retries and the mount retries to *n*, the default is 10000.

Supported Clients

This option is for all UNIX and Linux clients. The server can also define this option.

Options File

Place this option in the `dsm.sys` file within a server stanza *or* the client options file (`dsm.opt`).

Syntax

►►—NFSTIMEout— *number* —◄◄

Parameters

number

Specifies the number of seconds the client waits for a status system call on a file system before timing out. The range of values is 0 through 120; the default is 0 seconds.

Examples

Options file:

```
nfstimeout 10
```

Command line:

```
-nfstimeout=10
```

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

Related information

See “NFS hard and soft mounts” on page 217 for a discussion of how NFS hard and soft mounts are handled.

Nodename

Use the nodename option in your client options file to identify your workstation to the server. You can use different node names to identify multiple operating systems on your workstation.

When you use the nodename option, you are prompted for the password that is assigned to the node that you specify, if a password is required.

If you want to restore or retrieve files from the server while you are working from a different workstation, use the virtualnodename option. You can also use the asnodename option, if it is set up by the administrator.

When connecting to a server, the client must identify itself to the server. This login identification is determined in the following manner:

- In the absence of a nodename entry in the dsm.sys file, or a virtualnodename entry in the client user-options file (dsm.opt), or a virtual node name specified on a command line, the default login ID is the name that the **hostname** command returns.
- If a nodename entry exists in the dsm.sys file, the nodename entry overrides the name that the **hostname** command returns.
- If a virtualnodename entry exists in the client system-options file (dsm.sys), or a virtual node name is specified on a command line, it cannot be the same name as the name returned by the **hostname** command. When the server accepts the virtual node name, a password is required (if authentication is on), even if the passwordaccess option is generate. When a connection to the server is established, access is permitted to any file that is backed up using this login ID.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the `dsm.sys` file within a server stanza. You can set this option on the **General** tab, in the **Node Name** field of the Preferences editor.

Syntax

►►—NODename— —*nodename*—►►

Parameters

nodename

Specifies a 1 to 64 character node name for which you want to request Tivoli Storage Manager services. The default is the value returned with the **hostname** command.

Not specifying a node name permits the node name to default to the host name of the workstation

Examples

Options file:

`nodename cougar`

“Virtualnodename” on page 567

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:

- Non-default copy frequencies (other than 0) are not enforced on the Tivoli Storage Manager server.
- UNIX special file changes are not detected by the Journal daemon and are not, therefore, backed up.

For these reasons, you want to use the `nojournal` option periodically to perform a traditional full incremental backup.

Supported Clients

This option is valid for the AIX and Linux backup-archive client.

Syntax

►►—NOJournal—►►

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc incr /home -nojournal
```

Related concepts:

“Journal-based backup” on page 657

Noprompt

The **noprompt** option suppresses the confirmation prompt that is presented by the **delete group**, **delete archive**, **expire**, **restore image**, and **set event** commands.

- **delete archive**
- **delete backup**
- **delete group**
- **expire**
- **restore image**

Note: The **restore image** command does not apply to Mac OS X operating systems.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—NOPrompt—————◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc delete archive -noprompt "/Users/van/Documents/*"
```

Command line:

```
dsmc delete archive -noprompt "/home/project/*"
```

Nrtblepath

The **nrtblepath** option specifies the location of the node replication table on the client. The backup-archive client uses this table to store information about each backup or archive operation to the Tivoli Storage Manager server.

The server to which you back up your data must be at version 7.1 or newer and must replicate client node data to the secondary server.

When a failover occurs, the information that is on the secondary server might not be the most recent version if replication did not happen before the failover. The client can compare the information in the node replication table against the information that is on the secondary server to determine whether the backup on the server is the most recent backup version.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (`dsm.sys`).

This option can also be configured in the client option set on the Tivoli Storage Manager server.

Syntax

►—NRTABLEPath—*path*—►

Parameters

path

Specifies the location where the node replication table database is created. The default location is the backup-archive client installation directory.

For non-root users, you must specify a path that your user ID has write access to, such as a temporary directory. Most non-root users do not have access to the client installation directory.

Restriction: The node replication table cannot be created in the root directory (/). If you choose to specify a location for the node replication table, do not specify the root directory.

Example

Options file:

```
nrtablepath /Volumes/nrtbl
```

Command line:

Does not apply.

Related tasks:

“Determining the status of replicated client data” on page 92

“Configuring the client for automated failover” on page 90

Numberformat

The `numberformat` option specifies the format you want to use to display numbers.

The AIX, Solaris, and HP-UX clients support locales other than English that describe every user interface that varies with location or language.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time the client is called. Consult the documentation on your local system for details about setting up your locale definition.

Note: The `numberformat` option does not affect the web client. The web client uses the number format for the locale that the browser is running in. If the browser is not running in a supported locale, the web client uses the number format for US English.

You can use the `numberformat` option with the following commands:

- **delete archive**
- **delete backup**
- **expire**
- **query archive**
- **query backup**
- **query image**
- **query nas**
- **restore**
- **restore image**
- **restore nas**
- **retrieve**
- **set event**

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client user-options file (`dsm.opt`). You can set this option on the **Regional Settings** tab, **Number Format** field of the Preferences editor.

Syntax

►►—`NUMBERformat—` *—number—* ◀◀

Parameters

number

Displays numbers using any one of the following formats. Specify the number (0–6) that corresponds to the number format you want to use.

0 Use the locale-specified date format. This is the default (does not apply to Mac OS X).

1 1,000.00

This is the default for the following available translations:

- US English
- Japanese
- Chinese (Traditional)
- Chinese (Simplified)
- Korean

2 1,000,00

3 1 000,00

This is the default for the following available translations:

- French
- Czech
- Hungarian
- Polish
- Russian

4 1 000.00

5 1.000,00

This is the default for the following available translations:

- Brazilian Portuguese
- German
- Italian
- Spanish

6 1'000,00

For AIX, HP-UX, and Solaris: To define number formats, modify the following lines in the source file of your locale. Whatever format you select applies both to output and to input.

decimal_point

The character that separates the whole number from its fractional part.

thousands_sep

The character that separates the hundreds from the thousands from the millions.

grouping

The number of digits in each group that is separated by the `thousands_sep` character.

Examples

Options file:

num 4

Command line:

-numberformat=4

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the `dsm.opt` file unless overridden by the initial command line or by an option forced by the server.

Optfile

The `optfile` option specifies the client options file to use when you start a backup-archive client session.

Supported Clients

This option is valid for all clients.

Syntax

►►—OPTFILE =— —*file_name*—————►►

Parameters

file_name

Specifies an alternate client options file, if you use the fully qualified path name. If you specify only the file name, the client assumes the file name specified is located in the current working directory. The default is `dsm.opt`.

Restriction: Specify the full path when you use this option with the client acceptor daemon (dsmcad), because the client acceptor daemon changes its working directory to root ("/") after initialization.

Examples

Command line:

```
dsmc query session -optfile=myopts.opt
```

Client acceptor daemon:

```
dsmcad -optfile=/usr/tivoli/tsm/client/ba/bin/myopts.opt
```

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

Password

The password option specifies a Tivoli Storage Manager password.

If you do not specify this option and your administrator has set authentication to 0n, you are prompted for a password when you start a backup-archive client session.

Note:

1. If the server prompts for a password, the password is not displayed as you enter it. However, if you use the password option on the command line, your password is displayed as you enter it.
2. If the Tivoli Storage Manager server name changes or the backup-archive clients are directed to a different server, all clients must re-authenticate with the server because the stored encrypted password must be regenerated.

The password option is ignored when the passwordaccess option is set to generate.

Supported Clients

This option is valid for all clients.

Options File

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

Syntax

►►—PASsword— —password—►►

Parameters

password

Specifies the password you use to log on to the Tivoli Storage Manager server.

Passwords can be up to 63 character in length. Password constraints vary, depending on where the passwords are stored and managed, and depending on the version of the server that your client connects to.

If your Tivoli Storage Manager server is at version 6.3.3 or later, and if you use an LDAP directory server to authenticate passwords

Use any of the following characters to create a password:

```

a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9 . ! @ # $ % ^ & * _ - + = ` ( )
| { } [ ] : ; < > , ? / ~

```

Passwords are case-sensitive and are subject to more restrictions that can be imposed by LDAP policies.

If your Tivoli Storage Manager server is at version 6.3.3 or later, and if you do not use an LDAP directory server to authenticate passwords

Use any of the following characters to create a password:

```

a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9 . ! @ # $ % ^ & * _ - + = ` ( )
| { } [ ] : ; < > , ? / ~

```

Passwords are stored in the Tivoli Storage Manager server database and are not case-sensitive.

If your Tivoli Storage Manager server is earlier than version 6.3.3

Use any of the following characters to create a password:

```

a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9
_ - & + .

```

Passwords are stored in the Tivoli Storage Manager server database and are not case-sensitive.

Remember:

On the command line, enclose all parameters that contain one or more special characters in quotation marks. Without quotation marks, the special characters can be interpreted as shell escape characters, file redirection characters, or other characters that have significance to the operating system.

On AIX, HPUNIX, Linux, Mac, and Solaris systems:

Enclose the command parameters in single quotation marks (').

Command line example:

```

dsmc set password -type=vmguest 'Win 2012 SQL'
'tsm12dag\administrator' '7@#$$%^&7'

```

Quotation marks are not required when you type a password with special characters in an options file.

Examples

Options file:

```
password secretword
```

Command line:

```
-password=secretword
```

```
-password='my>pas$word'
```

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 one of the following methods:

- Set the password for your client node yourself and have the client prompt for it each time you request services.
- Let the client 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.
- If the server is not configured to require a password to log on to it, you can still be prompted to enter your node password when the backup-archive client establishes a connection with the server. This behavior occurs if this option, passwordaccess, is allowed to default or if you set it to passwordaccess prompt. The password that you supply in response to the prompt is used only to encrypt your login information; it is not used to log onto the server. In this configuration, you can avoid entering a password by setting this option to passwordaccess generate. Setting passwordaccess generate causes the client to create, store, and submit the password for you. When passwordaccess generate is set, the password option is ignored.

Setting the passwordaccess option to generate is required in the following situations:

- When using the HSM client.
- When using the web client.
- When performing NAS operations.
- When using Tivoli Continuous Data Protection for Files.


Supported Clients

This option is valid for all clients.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Authorization** tab, in the **Password Access** section of the Preferences editor.

Syntax



Parameters

prompt

You are prompted for your client node password each time a client connects to the server. This is the default.

To keep your client node password secure, enter commands without the password and wait for the client to prompt you for the password.

Each user must know the Tivoli Storage Manager password for your client node. Any user who knows the password for your client node can gain access to all backups and archives that originate from your client node. For example: If the user enters the node name and password for your client node from a different client node, the user becomes a virtual root user.

API applications must supply the password when a session is initiated. The application is responsible for obtaining the password (does not apply to Mac OS X).

generate

Encrypts and stores your password locally and generates a new password when the old password expires. The new password is randomly generated by the client. Password constraints vary, depending on where the passwords are stored and managed, and depending on the version of the server that your client connects to. Generated passwords are 63 characters in length and contain at least two of the following characters:

- upper case letters
- lower case letters
- numeric characters
- special characters

Additionally, the first and last character of a generated password is an alphabetic character, and they can be either upper or lower case. Generated passwords do not contain repeated characters.

A password prompt is displayed when registering a workstation with a server using open registration or if your administrator changes your password manually.

When logging in locally, users do not need to know the password for the client node. However, by using the nodename option at a remote node, users can access files they own and files to which another user grants access.

Examples

Options file:

passwordaccess generate

Command line:

Does not apply.

Related information

For information on where the password is stored, see “Passworddir.”

Passworddir

The passworddir option specifies the directory location in which to store an encrypted password file.

The default directory for AIX is /etc/security/adsm and for other UNIX and Linux platforms it is /etc/adsm. The default directory for Mac is /Library/Preferences/Tivoli Storage Manager. Regardless of where it is stored, the password file that is created by the client is always named TSM.sth. In turn three files comprise a password file. TSM.KDB stores the encrypted passwords. TSM.sth stores the random

encryption key that is used to encrypt passwords in the TSM.KDB file. This file is protected by the file system. TSM.IDX is an index file that is used to track the passwords in the TSM.KDB file.

Supported Clients

This option is valid for all UNIX clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza.

Syntax

►—PASSWORDDIR— *directoryname*—►

Parameters

directoryname

Specifies the path in which to store the encrypted password file. The name of the password file is TSM.sth. If any part of the specified path does not exist, Tivoli Storage Manager attempts to create it.

Examples

Options file:

```
passworddir "/Users/user1/Library/Preferences/Tivoli Storage Manager/"
passworddir /etc/security/tsm
```

Command line:

Does not apply.

Pick

The pick option creates a list of backup versions or archive copies that match the file specification you enter.

From the list, you can select the versions to process. Include the inactive option to view both active and inactive objects.

For images, if you do not specify a source file space and destination file space, the pick list contains all backed up images. In this case, the images selected from the pick list are restored to their original location. If you specify the source file space and the destination file space, you can select only one entry from the pick list.

Use the pick option with the following commands:

- **delete archive**
- **delete backup**
- **delete group**
- **expire**
- **restore**
- **restore group**
- **restore image**
- **restore nas**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—Pick—————►►

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore "/Users/van/Documents/*" -pick -inactive
```

Command line:

```
dsmc restore "/home/project/*" -pick -inactive
```

Pitdate

Use the pitdate option with the pittime option to establish a point in time to display or restore the latest version of your backups.

Files that were backed up *on or before* the date and time you specify, and which were not deleted *before* the date and time you specify, are processed. Backup versions that you create after this date and time are ignored.

Use the pitdate option with the following commands:

- **delete backup**
- **query backup**
- **query group**
- **query image**
- **restore**
- **restore group**
- **restore image**
- **restore nas**

When pitdate is used, the inactive and latest options are implicit.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—PITDate == —date—————►►

Parameters

date

Specifies the appropriate date. Enter the date in the format you selected with the `dateformat` option.

When you include `dateformat` with a command, it must precede the `fromdate`, `pitdate`, and `todate` options.

Examples

Command line:

```
dsmc restore "/Volumes/proj4/myproj/*" -sub=y -pitdate=08/01/2003  
-pittime=06:00:00
```

Command line:

```
dsmc restore "/fs1/*" -sub=y -pitdate=08/01/2003 -pittime=06:00:00
```

Pittime

Use the `pittime` option with the `pitdate` option to establish a point in time to display or restore the latest version of your backups.

Files that were backed up *on or before* the date and time you specify, and which were not deleted *before* the date and time you specify, are processed. Backup versions that you create after this date and time are ignored. This option is ignored if you do not specify `pitdate` option.

Use the `pittime` option with the following commands:

- **delete backup**
- **query backup**
- **query image**
- **restore**
- **restore image**
- **restore nas**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—PITtime =— —*time*—————►►

Parameters

time

Specifies a time on a specified date. If you do not specify a time, the time defaults to 23:59:59. Specify the time in the format you selected with the `timeformat` option.

When you include the `timeformat` option in a command, it must precede the `fromtime`, `pittime`, and `tottime` options.

Examples

Command line:

```
dsmc query backup -pitt=06:00:00 -pitd=08/01/2003  
"/Volumes/proj5/myproj/*"
```

Command line:

```
dsmc q b "/fs1/*" -pitt=06:00:00 -pitd=08/01/2003
```

Postschedulecmd/Postnschedulecmd

The `postschedulecmd/postnschedulecmd` option specifies a command that the client program processes after it runs a schedule.

If you want the client program to wait for the command to complete before it continues with other processing, use the `postschedulecmd` option. If you do not want to wait for the command to complete before the client continues with other processing, specify the `postnschedulecmd` option.

Return code handling and scheduled action behavior depends on both the option specified, and the type of operation that is scheduled:

- For scheduled operations where the scheduled action is something other than `COMMAND`:
If the `postschedulecmd` command does not complete with return code 0 (zero), the return code for the scheduled event is either 8, or the return code of the scheduled operation, whichever is greater. If you do not want the `postschedulecmd` command to be governed by this rule, you can create a script or batch file that starts the command and exits with return code 0. Then configure `postschedulecmd` to start the script or batch file.
- For scheduled operations where the scheduled action is `COMMAND`:
The return code from the command specified on the `postschedulecmd` option does not affect the return code that is reported to the server when the scheduled event completes. If you want the results of `postschedulecmd` operations to affect the return code of the scheduled event, include the `postschedulecmd` operations in the scheduled action command script instead of using the `postschedulecmd` option.
- If the scheduler action cannot be started, and the command specified on the `preschedulecmd` option completes with a return code of zero (0), the command specified by the `postschedulecmd` option is run.
- The return code from an operation specified on the `postnschedulecmd` option is not tracked, and does not influence the return code of the scheduled event.


Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **Scheduler** tab in the **Schedule Command** text box in the Preferences editor. The server can also define these options.

Syntax



Parameters

cmdstring

Specifies the command to process. You can enter a command to be run after a schedule with this option. Use only one postschedulecmd option.

If the command string contains blanks, enclose the command string in quotation marks. If you placed quotation marks within the command string, then enclose the entire command string in single quotation marks.

Use a blank, or null, string for *cmdstring* if you want to prevent any commands from running that the Tivoli Storage Manager server administrator uses for postschedulecmd or preschedulecmd. If you specify a blank or null string on either option, it prevents the administrator from using a command on both options.

If your administrator uses a blank or null string on the postschedulecmd option, you cannot run a post-schedule command.

For Mac OS X, if the postschedulecmd schedule command is a UNIX shell script, enter the name of the script to run. For example, if the file presched.sh is a UNIX shell script, enter this command:

```
postschedulecmd "/Volumes/La Pomme/Scripting/presched.sh"
```

Examples

Options file:

```
postschedulecmd "/Volumes/La Pomme/Scripting/postsched.sh"
```

Options file:

```
postschedulecmd "restart database"
```

The command string is a valid command for restarting your database.

Command line:

```
-postschedulecmd="/Volumes/La Pomme/Scripting/postsched.sh"
```

Command line:

```
-postschedulecmd="'restart database'"
```

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

Related concepts:

Chapter 8, "Client return codes," on page 273

Related reference:

 [DEFINE SCHEDULE command](#)

Postsnapshotcmd

The postsnapshotcmd option allows you to run operating system shell commands or scripts after the backup-archive client starts a snapshot during a snapshot-based backup operation.

AIX only: This option is only applicable to JFS2 snapshot-based file backup or archive and snapshot-based image backup. For a snapshot-based file backup or archive, use this option with the **backup** command, the `include.fs` option, or in the `dsm.sys` file.

Linux only: This option is only valid if the LVM is installed and configured on your system, allowing you to perform a snapshot-based image backup operation.

AIX and Linux only: For a snapshot-based image backup, use this option with the **backup image** command, the `include.image` option, or in the `dsm.sys` file.

If the `postsnapshotcmd` fails the operation continues, but appropriate warnings are logged.

Supported Clients

This option is valid for AIX clients and Linux x86_64 clients only. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can also set this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

►—`POSTSNAPshotcmd`— `"cmdstring"`—◄

Parameters

"cmdstring"

Specifies a command to process.

Use the `srvprepostsnapdisabled` option to prevent the Tivoli Storage Manager server administrator from executing operating system commands on the client system.

If the command string contains blanks, enclose the command string in quotation marks:

```
"resume database myDb"
```

If you placed quotation marks within the command string, then enclose the entire command string in single quotation marks:

```
'resume database "myDb"'
```

Examples

Options file:

```
postsnapshotcmd "any command"
```

The command string is a valid command for restarting your application.

Command line:

```
backup image -postsnapshotcmd="any command"
```

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

Related information

"Include options" on page 421

"Srvprepostscheddisabled" on page 531

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`.

Note:

1. Successful completion of the `preschedulecmd` command is considered to be a prerequisite to running the scheduled operation. If the `preschedulecmd` command does not complete with return code 0, the scheduled operation and any `postschedulecmd` and `postnschedulecmd` commands will not run. The client reports that the scheduled event failed, and the return code is 12. If you do not want the `preschedulecmd` command to be governed by this rule, you can create a script or batch file that invokes the command and exits with return code 0. Then configure `preschedulecmd` to invoke the script or batch file. The return code for the `preschedulecmd` command is not tracked, and does not influence the return code of the scheduled event.
2. The server can also define the `preschedulecmd` option (and the `prenschedulecmd` option).

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **Scheduler** tab, in the **Schedule Command** dialog box in the Preferences editor.

Syntax

► `PRESchedulecmd` `PRENSchedulecmd` `cmdstring` ►

Parameters

cmdstring

Specifies the command to process. Use only one `preschedulecmd` option. You can enter a command to be executed before a schedule using this option.

If the command string contains blanks, enclose the command string in quotation marks. If you placed quotation marks within the command string, then enclose the entire command string in single quotation marks.

Use a blank or null string for *cmdstring* if you want to prevent any commands from running that the Tivoli Storage Manager server administrator uses for

postschedulecmd and preschedulecmd. If you specify a blank or null string on either option, it prevents the administrator from using a command on both options.

If your administrator uses a blank or null string on the preschedulecmd option, you cannot run a pre-schedule command.

For Mac OS X, if the preschedulecmd schedule command is an AppleScript, you must use the **osascript** command to run the script. For example, if "Database Script" is an apple script, enter this command:

```
preschedulecmd osascript "/Volumes/La Pomme/Scripting/  
Database Script"
```

Examples

Options file:

```
preschedulecmd "<the quiesce command of your database product>  
database"
```

The command string is a valid command for quiescing your database.

Command line:

```
-preschedulecmd="'quiesce database'"
```

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

Related concepts:

Chapter 8, "Client return codes," on page 273

PreserveLastAccessDate

Use the `preserveLastAccessDate` option to specify whether a backup or archive operation changes the last access time.

A backup or archive operation can change the last access time of a file. After an operation, the backup-archive client can reset the last access time to the value before the operation. The last access time can be preserved, rather than modified, by the backup-archive client. Resetting the last access time requires extra processing for each file that is backed up or archived.

If you enable open file support, the last access date for files is always preserved regardless of the setting for `preserveLastAccessDate`. When open file support is enabled, do not use the `preserveLastAccessDate` option.

Use this option with the **incremental**, **selective**, or **archive** commands.

Note:

1. This option applies only to files; it does not apply to directories.
2. Resetting the last access date affects backup and archive performance.
3. Resetting the last access date can affect applications that rely on accurate last-access dates such as a Storage Resource Management (SRM) application.
4. On file systems that are not managed by the Tivoli Storage Manager for Space Management client or when non-root users back up or archive, the `ctime` attribute is reset. The last changed time and date (`ctime`) attribute is reset to the date and time of the backup or archive operation.

5. The `updatectime` option takes precedence over the `preserveastaccessdate` option. If both options are set to `yes`, the `preserveastaccessdate` option is ignored
6. On file systems that are not managed by the Tivoli Storage Manager for Space Management client, do not use `preserveastaccessdate yes` and the GPFS `mmbackup` command. The `mmbackup` command and `preserveastaccessdate yes` selects all files for each backup operation.
7. You cannot reset the last access date of read-only files. The `preserveastaccessdate` option ignores read-only files and does not change their date.

Supported Clients

This option is valid for all UNIX and Linux clients except Mac OS X.

The server can also define this option.

Options File

Place this option in the client user options file (`dsm.opt`). You can set this option on the Backup tab of the Preferences editor.

Syntax



Parameters

No A backup or archive operation can change the last access date. This value is the default.

Yes

A backup or archive operation does not change the last access date.

Examples

Options file:

```
preserveastaccessdate yes
```

Command line:

```
Incremental /proj/test/test_file -preserveastaccessdate=yes
```

Related information:

-  `mmbackup` command: Tivoli Storage Manager requirements
-  Guidance for integrating IBM Spectrum Scale AFM with IBM Spectrum Protect
-  Using IBM Spectrum Protect include and exclude options with IBM Spectrum Scale `mmbackup` command

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 the client prompt you before files are overwritten.


Use the `preservepath` option with the following commands:

- **restore**
- **restore backupset**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option. The server can also define 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 `SUBDIR=yes`, the client restores all files in the source directories to the single target directory.

Examples

Command line:

Assume that the server file space contains the following backup copies:

```

/fs/h1/m1/file.a
/fs/h1/m1/file.b
/fs/h1/m1/l1/file.x
/fs/h1/m1/l1/file.y
  
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -preser=complete
```

Restores these directories and files:

/u/ann/h1/m1/file.a
/u/ann/h1/m1/file.b

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -preser=nobase
```

Restores these directories and files:

/u/ann/file.a
/u/ann/file.b

This command:

```
dsmc res backupset /fs/h1/m1/ /u/ann/ -su=yes  
-preser=nobase -loc=file
```

Restores these directories and files:

/u/ann/file.a
/u/ann/file.b
/u/ann/file.x
/u/ann/file.y

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -preser=subtree
```

Restores these directories and files:

/u/ann/m1/file.a
/u/ann/m1/file.b

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -preser=none
```

Restores these directories and files:

/u/ann/file.a
/u/ann/file.b

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -su=yes -preser=complete
```

Restores these directories and files:

/u/ann/h1/m1/file.a
/u/ann/h1/m1/file.b
/u/ann/h1/m1/l1/file.x
/u/ann/h1/m1/l1/file.y

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -su=yes -preser=nobase
```

Restores these directories and files:

/u/ann/file.a
/u/ann/file.b
/u/ann/l1/file.x
/u/ann/l1/file.y

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -su=yes -preser=subtree
```

Restores these directories and files:


```
/u/ann/m1/file.a
/u/ann/m1/file.b
/u/ann/m1/l1/file.x
/u/ann/m1/l1/file.y
```

This command:

```
dsmc res /fs/h1/m1/ /u/ann/ -su=yes -preser=none
```

Restores these directories and files:

```
/u/ann/file.a
/u/ann/file.b
/u/ann/file.x
/u/ann/file.y
```

Presnapshotcmd

The `presnapshotcmd` option allows you to run operating system commands before the backup-archive client starts a snapshot.

This allows you to quiesce an application before the client starts the snapshot during a snapshot-based backup or archive.

AIX only: This option is only applicable to JFS2 snapshot-based file backup or archive and snapshot-based image backup. For a snapshot-based file backup or archive, use this option with the **backup** command, the `include.fs` option, or in the `dsm.sys` file.

Linux only: This option is only valid if the LVM is installed and configured on your system, allowing you to perform a snapshot-based image backup.

AIX and Linux only: For a snapshot-based image backup, use this option with the **backup image** command, the `include.image` option, or in the `dsm.sys` file.

If the `presnapshotcmd` fails it is assumed that the application is not in a consistent state and the client stops the operation and display the appropriate error message.

Supported Clients

This option is valid for AIX JFS2 and Linux x86_64 clients only. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set also this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

►►—PRESNAPshotcmd— —"*cmdstring*"——►►

Parameters

"cmdstring"

Specifies a command to process.

Use the `srvprepostsnapdisabled` option to prevent the Tivoli Storage Manager server administrator from running operating system commands on the client system.

If the command string contains blanks, enclose the command string in quotation marks:

```
"quiesce database myDb"
```

If you placed quotation marks within the command string, then enclose the entire command string in single quotation marks:

```
'resume database "myDb"'
```

Examples

Options file:

```
presnapshotcmd "any shell command or script"
```

Command line:

```
backup image -presnapshotcmd="any shell command or script"
```

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

Related information

"Include options" on page 421

"Srvprepostscheddisabled" on page 531

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 clients. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **Scheduler** tab, in the **Query Schedule Interval** field of the Preferences editor.

Syntax

►►—QUERYSCHEdperiod— —hours—————►►

Parameters

hours

Specifies the number of hours the client scheduler waits between attempts to contact the server for scheduled work. The range of values is 1 through 9999; the default is 12.

Examples

Options file:

```
querysch 6
```

Command line:

```
-queryschedperiod=8
```

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

Querysummary

The **querysummary** option provides statistics about files, directories and objects that are returned by the **query backup** or **query archive** commands.

The following statistics are provided by the **querysummary** option:

- The aggregate number of files and directories that are returned by the **query backup** or **query archive** command
- The aggregate amount of data of the objects that are returned by the **query backup** or **query archive** command
- The classic restore memory-utilization estimate to restore objects that are returned by the **query backup** or **query archive** command
- The total number of unique server volumes where the objects that are returned by the **query** command reside

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—QUERYSUMMARY—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc q ba '/usr/fs1/*' -sub=yes -querysummary
```

```
[root@kaveri:/home/cpark] $ dsmc q ba '/kalafs1/*' -sub=yes -querysummary
```

```
IBM Tivoli Storage Manager
```

```
Command Line Backup-Archive Client Interface
```

```
Client Version 6, Release 1, Level 0.0 0804FB
```

```
Client date/time: 08/19/08 08:09:48
```

```
(c) Copyright by IBM Corporation and other(s) 1990, 2008. All Rights Reserved.
```

```

Node Name: KAVERI
Session established with server TEMPLAR: AIX-RS/6000
  Server Version 5, Release 4, Level 3.0
  Server date/time: 08/19/08   08:09:02   Last access: 08/19/08   07:48:59

```

Size		Backup Date		Mgmt Class	A/I	File
----		-----		-----	---	----
4,096	B	08/07/08	12:07:30	BASVT2	A	/kalafs1/
256	B	08/07/08	12:07:30	BASVT2	A	/kalafs1/dir1
10,485,760	B	08/07/08	12:07:30	DEFAULT	A	/kalafs1/info1
5,242,880	B	08/07/08	12:07:30	DEFAULT	A	/kalafs1/info2
1,044	B	08/07/08	12:07:30	DEFAULT	A	/kalafs1/dir1/subfile1
1,044	B	08/07/08	12:07:30	DEFAULT	A	/kalafs1/dir1/subfile2

Summary Statistics

Total Files	Total Dirs	Avg. File Size	Total Data	Memory Est.
-----	-----	-----	-----	-----
4	2	3.75 MB	15.00 MB	1.07 KB

Estimated Number of Volumes: 2

[root@kaveri:/home/cpark] \$

Quiet

The quiet option limits the number of messages that are displayed on your screen during processing..

For example, when you run the **incremental**, **selective**, or **archive** commands, information might appear about each file that is backed up. Use the quiet option if you do not want to display this information

When you use the quiet option, error and processing information appears on your screen, and messages are written to log files. If you do not specify quiet, the default option, verbose is used.

Supported Clients

This option is valid for all clients. The server can also define the quiet option, overriding the client setting. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Command Line** tab, **Do not display process information on screen** checkbox of the Preferences editor.

Syntax

►►—QUIET—◄◄

Parameters

There are no parameters for this option.

Options file:
quiet

Command line:
-quiet

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

Quotesareliteral

The `quotesareliteral` option specifies whether single quotation marks (') or double quotation marks (") are interpreted literally, when they are included in a file list specification on a `filelist` option.

Ordinarily, the client requires you to use single or double quotation marks to delimit file specifications that contain space characters. Some file systems, such as the IBM General Parallel File System (GPFS), allow single and double quotation marks in file and directory names.

To prevent errors that would otherwise occur, when file specifications are included on a `filelist` option and they contain single quotation marks (') or double quotation marks ("), set `quotesareliteral` yes. When `quotesareliteral` is set to yes, quotation marks that are included in a file list specification on a `filelist` option are interpreted literally, as quotation marks, and not as delimiters.

This option applies to any command that accepts a `filelist` option as command parameter.

Supported Clients

This option is valid for all supported platforms. The option is applied to any command that takes a file list specification as a parameter.

Options File

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

Syntax



Parameters

no Specifies that single quotation marks (') and double quotation marks (") are interpreted as delimiters for file list specifications included on a `filelist` option. No is the default setting.

yes

Specifies that single quotation marks (') and double quotation marks (") are interpreted literally, and not as delimiters, for file list specifications that are included on a `filelist` option. Specify this value if you are backing up files from a file system that allows quotation marks in file or directory names.

Examples

Options file:

QUOTESARELITERAL YES

Command line:

Assuming that the file system allows quotation marks in paths, the following are examples of files in a file list specification that can be successfully processed if QUOTESARELITERAL is set to YES

Assume the command that is issued is `dsmc sel -filelist=/home/user1/important_files`, where `important_files.txt` contains the list of files to process.

```
/home/user1/myfiles/"file"1000  
/home/user1/myfiles/'file'  
/home/user1/myfiles/file'ABC  
/home/user1/myfiles/ABC"file"
```

Related information

For information about the `filelist` option, see “Filelist” on page 402.

For information about syntax for file specifications, see “Specifying input strings that contain blank spaces or quotation marks” on page 138.

“Wildcardsareliteral” on page 600

Removeoperandlimit

The `removeoperandlimit` option specifies that the client removes the 20-operand limit.

If you specify the `removeoperandlimit` option with the **incremental**, **selective**, or **archive** commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits.

The `removeoperandlimit` option can be useful if you generate scripts which can invoke the command-line client with a large number of operands. For example, you can prescan a directory tree looking for files to back up. As each *eligible* file is discovered, it is added to the operand list of a **selective** command. Later, this **selective** command is submitted by a controlling script. In this case, specifying the `removeoperandlimit` option removes the 20-operand limit.

Note:

1. The `removeoperandlimit` option *must* be placed immediately after the **incremental**, **selective**, or **archive** command before any file specifications.
2. This option does not accept a value. If this option is specified on a command, the 20-operand limit is removed.
3. Because it adversely affects performance to allow the shell to expand wild cards, use the `removeoperandlimit` option in backup or archive operations in which wild cards are not used.
4. The `removeoperandlimit` option is valid only on the **incremental**, **selective**, or **archive** commands in batch mode. It is not valid in the client options file (`dsm.opt`) or `dsm.sys` file.

Supported Clients

This option is valid for all UNIX and Linux clients.

Syntax

►►—REMOVEOperandlimit—►►

Parameters

There are no parameters for this option.

Examples

Command line:

`-removeoperandlimit`

Replace

The replace option specifies whether to overwrite existing files on your workstation, or to prompt you for your selection when you restore or retrieve files.

Important: The replace option does not affect recovery of directory objects. Directory objects are always recovered, even when specifying `replace=no`. To prevent overwriting existing directories, use the `filesonly` option.

You can use this option with the following commands:

- **restore**
- **restore backupset**
- **restore group**
- **retrieve**

Note: Replace prompting does not occur during a scheduled operation. If you set the replace option to prompt, the backup-archive client skips files without prompting you during a scheduled operation.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client user-options file (`dsm.opt`). You can set this option on the **Restore** tab, **Action for files that already exist** section of the Preferences editor.

Syntax

►►—REPlace—

Prompt
All
Yes
No

—►►

Parameters

Prompt

For nonscheduled operations, you specify whether to overwrite existing files. For scheduled operations, existing files are not overwritten and no prompts are displayed. This is the default.

All

All existing files are overwritten, including read-only files. If access to a file is denied, you are prompted to skip or overwrite the file. No action is taken on the file until there is a response to the prompt.

Yes

Existing files are overwritten, *except* read-only files. For nonscheduled operations, you specify whether to overwrite existing read-only files. For scheduled operations, existing read-only files are not overwritten and no prompts are displayed. If access to a file is denied, the file is skipped.

No Existing files are not overwritten. No prompts are displayed.

Examples

Options file:

```
replace all
```

Command line:

```
-replace=no
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the `dsm.opt` file unless overridden by the initial command line or by an option forced by the server.

Replserverguid

The `replserverguid` option specifies the globally unique identifier (GUID) that is used when the client connects to the secondary server during failover. The GUID is used to validate the secondary server to ensure that it is the expected server.

The replication GUID is different from the machine GUID of the server. It is generated one time for a server that is doing the replication and never changes.

This option must be specified within a **replservername** stanza in the client options file. The **replservername** stanza contains connection information about the secondary server.

This option is set by the Tivoli Storage Manager server administrator for the client node. During the normal (non-failover) logon process, the option is sent to the client and is saved in the client options file.

Do not edit this option during normal operations.

Edit this option only during situations such as the following ones:

- The primary server is offline and the information for the secondary server is not in the options file.
- The secondary server information is out-of-date or incorrect.

Any values that you edit are removed or updated the next time you log in to the primary server.

Supported Clients

This option is valid for all clients.

Options File

This option is placed in the `dsm.sys` file within the `replservername` stanza.

Syntax

►—`replserverguid—serverguid`—►

Parameters

serverguid

Specifies the GUID of the secondary server that is used during a failover.

Examples

Options file:

```
REPLSERVERGUID 91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.02
```

Command line:

Does not apply.

Options file:

The following example demonstrates how to specify options for three different servers in the `dsm.sys` file, and how to reference the secondary server. Connection information for multiple secondary servers is presented in stanzas. Each stanza is identified by the **replservername** option and the name of the secondary server. The **servername** stanza must contain the **myreplicationserver** option, which points to the secondary server that is specified by the **replservername** stanza. Only one secondary server can be specified per **servername** stanza.

```
REPLSERVERNAME TargetReplicationServer1
REPLTCPSERVERADDRESS TargetReplicationServer1
REPLTCPPOINT 1505
REPLSSLPORT 1506
REPLSERVERGUID 91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00
```

```
REPLSERVERNAME TargetReplicationServer2
REPLTCPSERVERADDRESS TargetReplicationServer2
REPLTCPPOINT 1505
REPLSSLPORT 1506
REPLSERVERGUID 91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.02
```

```
SErvername server_a
COMMMethod TCPip
TCPPOINT 1500
TCPServeraddress server_hostname1.example.com
PASSWORDAccess prompt
MYREPLICATIONServer TargetReplicationServer1
```

```
SErvername server_b
COMMMethod TCPip
TCPPOINT 1500
TCPServeraddress server_hostname2.example.com
PASSWORDAccess generate
```

INCLExcl	/adm/tsm/archive.excl
MYREPLICATIONServer	TargetReplicationServer2
Servname	server_c
COMMethod	TCPip
TCPPort	1500
TCPServeraddress	server_hostname3.example.com
PASSWORDAccess	generate
MYREPLICATIONServer	TargetReplicationServer1

Related concepts:

“Automated client failover configuration and use” on page 87

Related tasks:

“Configuring the client for automated failover” on page 90

Replservername

The replservername option specifies the name of the secondary server that the client connects to during a failover.

The replservername option begins a stanza in the client options file that contains connection information about the secondary server.

This option is set by the Tivoli Storage Manager server administrator for the client node. During the normal (non-failover) logon process, the option is sent to the client and is saved in the client options file.

Do not edit this option during normal operations.

Edit this option only during situations such as the following ones:

- The primary server is offline and the information for the secondary server is not in the options file.
- The secondary server information is out-of-date or incorrect.

Any values that you edit are removed or updated the next time you log in to the primary server.

Supported Clients

This option is valid for all clients.

Options File

This option is placed in the client-system options dsm.sys.

Syntax

►► replservername=repl_servername ◀◀

Parameters

repl_servername

Specifies the name of the secondary server to be used during a failover. This value is usually the name of the secondary server, not the host name of the server.

Examples

Options file:

```
REPLSERVERName    TargetReplicationServer1
```

Command line:

Does not apply.

Options file:

The following example demonstrates how to specify options for three different servers in the `dsm.sys` file, and how to reference the secondary server. Connection information for multiple secondary servers is presented in stanzas. Each stanza is identified by the **replservername** option and the name of the secondary server. The **servername** stanza must contain the **myreplicationserver** option, which points to the secondary server that is specified by the **replservername** stanza. Only one secondary server can be specified per **servername** stanza.

```
REPLSERVERNAME    TargetReplicationServer1
REPLTCPSEVERADDRESS TargetReplicationServer1
REPLTCPPEORT      1505
REPLSSLPEORT      1506
REPLSERVERGUID    91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00
```

```
REPLSERVERNAME    TargetReplicationServer2
REPLTCPSEVERADDRESS TargetReplicationServer2
REPLTCPPEORT      1505
REPLSSLPEORT      1506
REPLSERVERGUID    91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00.02
```

```
SErvername        server_a
COMMMethod         TCPip
TCPPEORT           1500
TCPSEVERADDRESS    server_hostname1.example.com
PASSWORDAccess     prompt
MYREPLICATIONServer TargetReplicationServer1
```

```
SErvername        server_b
COMMMethod         TCPip
TCPPEORT           1500
TCPSEVERADDRESS    server_hostname2.example.com
PASSWORDAccess     generate
INCLEXcl           /adm/tsm/archive.excl
MYREPLICATIONServer TargetReplicationServer2
```

```
SErvername        server_c
COMMMethod         TCPip
TCPPEORT           1500
TCPSEVERADDRESS    server_hostname3.example.com
PASSWORDAccess     generate
MYREPLICATIONServer TargetReplicationServer1
```

Related concepts:

“Automated client failover configuration and use” on page 87

Related tasks:

“Configuring the client for automated failover” on page 90

Replsslport

The `replsslport` option specifies the TCP/IP port on the secondary server that is SSL-enabled. The `replsslport` option is used when the client connects to the secondary server during a failover. This option is deprecated if you are connecting to a Tivoli Storage Manager server V7.1.8 and later V7 levels, and V8.1.2 and later levels.

The `replsslport` option is sent to the client by the primary server only if the secondary server is configured for SSL.

This option is applicable only when the client is configured to use SSL for secure communications between the Tivoli Storage Manager server and client. If the client is not configured to use SSL, the port that is specified by the `repltcpport` option is used. You can determine whether the client uses SSL by verifying the SSL client option.

This option must be specified within a **replservername** stanza in the client options file. The **replservername** stanza contains connection information about the secondary server.

During the normal (non-failover) logon process, this option is sent to the client and is saved in the client options file.

Do not edit this option during normal operations.

Edit this option only during situations such as the following ones:

- The primary server is offline and the information for the secondary server is not in the options file.
- The secondary server information is out-of-date or incorrect.

Any values that you edit are removed or updated the next time you log in to the primary server.

Supported Clients

This option is valid for all clients.

Options File

This option is placed in the `dsm.sys` file within the `replservername` stanza.

Syntax

►►—`replsslport—port_address`—————►►

Parameters

port_address

Specifies the TCP/IP port address that is enabled for SSL and that is used to communicate with the secondary server.

Examples

Options file:

`REPLSSLPORT 1506`

Command line:

Does not apply.

Options file:

The following example demonstrates how to specify options for three different servers in the `dsm.sys` file, and how to reference the secondary server. Connection information for multiple secondary servers is presented in stanzas. Each stanza is identified by the **replservername** option and the

name of the secondary server. The **servername** stanza must contain the **myreplicationserver** option, which points to the secondary server that is specified by the **replservername** stanza. Only one secondary server can be specified per **servername** stanza.

```

REPLSERVERNAME    TargetReplicationServer1
REPLTCPSERVERADDRESS TargetReplicationServer1
REPLTCPPOrt       1505
REPLSSLPORT       1506
REPLSERVERGUID    91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00

REPLSERVERNAME    TargetReplicationServer2
REPLTCPSERVERADDRESS TargetReplicationServer2
REPLTCPPOrt       1505
REPLSSLPORT       1506
REPLSERVERGUID    91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00.02

```

```

SErvername        server_a
COMMMethod        TCPip
TCPPOrt           1500
TCPSErveraddress  server_hostname1.example.com
PASSWORDAccess    prompt
MYREPLICATIONServer TargetReplicationServer1

```

```

SErvername        server_b
COMMMethod        TCPip
TCPPOrt           1500
TCPSErveraddress  server_hostname2.example.com
PASSWORDAccess    generate
INCLExcl          /adm/tsm/archive.excl
MYREPLICATIONServer TargetReplicationServer2

```

```

SErvername        server_c
COMMMethod        TCPip
TCPPOrt           1500
TCPSErveraddress  server_hostname3.example.com
PASSWORDAccess    generate
MYREPLICATIONServer TargetReplicationServer1

```

Related concepts:

“Automated client failover configuration and use” on page 87

Related tasks:

“Configuring the client for automated failover” on page 90

Repltcpport

The **repltcpport** option specifies the TCP/IP port on the secondary server to be used when the client connects to the secondary server during a failover.

This option must be specified within a **replservername** stanza in the client options file. The **replservername** stanza contains connection information about the secondary server.

This option is set by the Tivoli Storage Manager server administrator for the client node. During the normal (non-failover) logon process, the option is sent to the client and is saved in the client options file.

Do not edit this option during normal operations.

Edit this option only during situations such as the following ones:

- The primary server is offline and the information for the secondary server is not in the options file.

- The secondary server information is out-of-date or incorrect.

Any values that you edit are removed or updated the next time you log in to the primary server.

Supported Clients

This option is valid for all clients.

Options File

This option is placed in the `dsm.sys` file within the `replservername` stanza.

Syntax

►►—`repltcpport`—*port_address*—►►

Parameters

port_address

Specifies the TCP/IP port address that is used to communicate with the secondary server.

Examples

Options file:

```
REPLTCPPort 1500
```

Command line:

Does not apply.

Options file:

The following example demonstrates how to specify options for three different servers in the `dsm.sys` file, and how to reference the secondary server. Connection information for multiple secondary servers is presented in stanzas. Each stanza is identified by the **replservername** option and the name of the secondary server. The **servername** stanza must contain the **myreplicationserver** option, which points to the secondary server that is specified by the **replservername** stanza. Only one secondary server can be specified per **servername** stanza.

```
REPLSERVERNAME TargetReplicationServer1
REPLTCPSEVERADDRESS TargetReplicationServer1
REPLTCPPOINT      1505
REPLSSLPORT       1506
REPLSERVERGUID    91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00
```

```
REPLSERVERNAME TargetReplicationServer2
REPLTCPSEVERADDRESS TargetReplicationServer2
REPLTCPPOINT      1505
REPLSSLPORT       1506
REPLSERVERGUID    91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00.02
```

```
SErvername      server_a
COMMMethod      TCPip
TCPPOINT        1500
TCPSEVERADDRESS server_hostname1.example.com
PASSWORDAccess  prompt
MYREPLICATIONServer TargetReplicationServer1
```

```
SErvername      server_b
COMMMethod      TCPip
```

TCPPort	1500
TCPServeraddress	server_hostname2.example.com
PASSWORDAccess	generate
INCLExc1	/adm/tsm/archive.exc1
MYREPLICATIONServer	TargetReplicationServer2
SErvername	server_c
COMMMethod	TCPip
TCPPort	1500
TCPServeraddress	server_hostname3.example.com
PASSWORDAccess	generate
MYREPLICATIONServer	TargetReplicationServer1

Related concepts:

“Automated client failover configuration and use” on page 87

Related tasks:

“Configuring the client for automated failover” on page 90

Repltcpserveraddress

The `repltcpserveraddress` option specifies the TCP/IP address of the secondary server to be used when the client connects to the secondary server during a failover.

This option must be specified within a **replservername** stanza in the client options file. The **replservername** stanza contains connection information about the secondary server.

This option is set by the Tivoli Storage Manager server administrator for the client node. During the normal (non-failover) logon process, the option is sent to the client and is saved in the client options file.

Do not edit this option during normal operations.

Edit this option only during situations such as the following ones:

- The primary server is offline and the information for the secondary server is not in the options file.
- The secondary server information is out-of-date or incorrect.

Any values that you edit are removed or updated the next time you log in to the primary server.

Supported Clients

This option is valid for all clients.

Options File

This option is placed in the `dsm.sys` file within the **replservername** stanza.

Syntax

►►—REPLTCPserveraddress—*server_address*—◄◄

Parameters

server_address

Specifies a TCP/IP address for a server that is 1 - 64 characters in length.

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 use only IPv6 addresses if you specified the commmethod V6Tcpip option.

Examples

Options file:

```
REPLTCPServeraddress dsmchost.example.com
```

Command line:

Does not apply.

Options file:

The following example demonstrates how to specify options for three different servers in the dsm.sys file, and how to reference the secondary server. Connection information for multiple secondary servers is presented in stanzas. Each stanza is identified by the **replservername** option and the name of the secondary server. The **servername** stanza must contain the **myreplicationserver** option, which points to the secondary server that is specified by the **replservername** stanza. Only one secondary server can be specified per **servername** stanza.

```
REPLSERVERNAME TargetReplicationServer1
REPLTCPSEVERADDRESS TargetReplicationServer1
REPLTCPPOINT 1505
REPLSSLPORT 1506
REPLSERVERGUID 91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00.00
```

```
REPLSERVERNAME TargetReplicationServer2
REPLTCPSEVERADDRESS TargetReplicationServer2
REPLTCPPOINT 1505
REPLSSLPORT 1506
REPLSERVERGUID 91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00.02
```

```
SErvername server_a
COMMMethod TCPip
TCPPOINT 1500
TCPSEVERADDRESS server_hostname1.example.com
PASSWORDACCESS prompt
MYREPLICATIONServer TargetReplicationServer1
```

```
SErvername server_b
COMMMethod TCPip
TCPPOINT 1500
TCPSEVERADDRESS server_hostname2.example.com
PASSWORDACCESS generate
INCLEXCL /adm/tsm/archive.excl
MYREPLICATIONServer TargetReplicationServer2
```

```
SErvername server_c
COMMMethod TCPip
TCPPOINT 1500
TCPSEVERADDRESS server_hostname3.example.com
PASSWORDACCESS generate
MYREPLICATIONServer TargetReplicationServer1
```

Related concepts:

“Automated client failover configuration and use” on page 87

Related tasks:

“Configuring the client for automated failover” on page 90

Resourceutilization

Use the `resourceutilization` option in your option file to regulate the level of resources the Tivoli Storage Manager server and client can use during processing.

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the `dsm.sys` file within a server stanza. You can set this option on the **General** tab, in the **Resource Utilization** field of the Preferences editor.

Syntax

►►—`RESOURCEutilization—`*number*—◄◄

Parameters

number

Specifies the level of resources the Tivoli Storage Manager server and client can use during processing. The range of values that you can specify is 1 through 10.

Examples

Options file:

```
resourceutilization 7
```

Command line:

```
-resourceutilization=7
```

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

Regulating backup and archive sessions

When you request a backup or archive, the client can use more than one session to the server.

The default is to use a maximum of two sessions; one to query the server and one to send file data. The client can use only one server session if you set the **resourceutilization** option to 1.

A client can use more than the default number of sessions when it connects to a Tivoli Storage Manager server. For example, `resourceutilization 10` permits up to eight sessions with the server. Multiple sessions can be used for querying the server and sending file data.

Multiple query sessions are used when you specify multiple file specifications with a backup or archive command. For example, if you enter the following commands and you specify `resourceutilization 5`, the client might start a second session to query files on file space B.

```
inc /Volumes/filespaceA /Volumes/filespaceB
```

Whether the second session starts depends on how long it takes to query the server about files that are backed up on file space A. The client might also try to read data from the file system and send it to the server on multiple sessions.

Note: During a backup operation, if you enter multiple file specifications, the result might be that files from one file specification are stored on multiple tapes and interspersed with files from different file specifications. This can decrease restore performance. Setting the **collocatebyfilespec** option to yes eliminates interspersing of files from different file specifications, by limiting the client to one server session per file specification. Therefore, if you store the data to tape, files for each file specification are stored together on one tape (unless another tape is required for more capacity).

Related reference:

“Collocatebyfilespec” on page 334

Regulating restore sessions

When you request a restore, the default is to use a maximum of one session.

Additional restore sessions are based on:

- **resourceutilization** value
- how many tapes on which the requested data is stored
- how many tape drives are available
- the maximum number of mount points that are allowed for the node

Note:

1. If all of the files are on disk, only one session is used. There is no multi-session for a pure disk storage pool restore. However, if you are performing a restore in which the files are on 4 tapes and others are 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 by using the **MAXNUMMP** parameter. If the **resourceutilization** option value exceeds the value of the **MAXNUMMP** on the server for a node, the backup can fail with an Unknown System Error message.
3. You can get a multi-session restore from your single **restore** command, and from a single volume on the server, if that volume is device class FILE.

For example, if the data you want to restore is on 5 different tape volumes, the maximum number of mount points is 5 for your node, and **resourceutilization** is set to 3, then 3 sessions are used for the restore. If you increase the **resourceutilization** setting to 5, then 5 sessions are used for the restore. There is a 1 to 1 relationship between the number of restore sessions that are allowed and the **resourceutilization** setting. Multiple restore sessions are only allowed for no-query restore operations.

Multiple client session considerations

This topic lists some items to consider when working with multiple client sessions.

The following factors can affect the throughput of multiple sessions:

- The ability of the server to handle multiple client sessions. Is there sufficient memory, multiple storage volumes, and processor cycles to increase backup throughput?
- The ability of the client to drive multiple sessions (sufficient processor cycles, memory, etc.).

- The configuration of the client storage subsystem. File systems that are striped across multiple disks, using either software striping or RAID-5 can better handle an increase in random read requests than a single drive file system. Additionally, a single drive file system might not see performance improvement if it attempts to handle many random concurrent read requests.
- Sufficient bandwidth in the network to support the increased traffic.

Potentially undesirable aspects of running multiple sessions include:

- The client could produce multiple accounting records.
- The server might not start enough concurrent sessions. To avoid this, the server *maxsessions* parameter must be reviewed and possibly changed.
- A query node command might not summarize client activity.
- It is possible that files are restored instead of hard links.

Restoring files instead of hard links can occur when the following criteria are all true:

- You restore an entire file system.
- During the restore operation, the value of the *resourceutilization* option is greater than 1.
- The file system contained hard links when the file system was backed up.

The chance of restoring linked files instead of hard links increases as the number of sessions increases. When you restore a file system that contained hard links when the file system was backed up, set *resourceutilization*=1 to ensure that hard links are restored.

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 system options file after your client node successfully contacts the server.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client system-options file (*dsm.sys*) within a server stanza. You can set this option on the **Scheduler** tab, in the **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:

```
retry 10
```

Command line:

```
-retryperiod=10
```

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

Revokeremoteaccess

The `revokeremoteaccess` option restricts an administrator with client access privilege from accessing a client workstation that is running the web client.

This option does not restrict administrators with client owner, system, or policy privilege from accessing your workstation through the web client.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **Web Client** 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.

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.

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the `dsm.sys` file within a server stanza.

Syntax



Parameters

Yes

Specifies that the server disables the scheduling of commands using the `action=command` option on the **DEFINE SCHEDULE** server command.

No Specifies that the server does not disable the scheduling of commands using the `action=command` option on the **DEFINE SCHEDULE** server command. This is the default.

Examples

Options file:

```
schedcmddisabled no
```

Command line:

Does not apply.

Related information

“Query Schedule” on page 684

Schedcmdexception

The schedcmdexception option is used in conjunction with the schedcmddisabled option to disable the scheduling of commands by the server action=**command** option on the DEFINE SCHEDULE server command, except for specific command strings.

You must specify the exact string that matches the “objects” definition in the schedule for the scheduled server command to be accepted. If the string does not match exactly (for example, there is an extra space or the capitalization is different), the scheduled command action is blocked.

You can provide multiple schedcmdexception options in the options file. This option is not honored if schedcmddisabled is not enabled. The placement of this option in the options file is independent of the placement of the schedcmddisabled option.

Supported Clients

This option is valid for all clients. This option is not valid in the Tivoli Storage Manager server client options set.

Options File

Place this option in the dsm.sys file within a server stanza.

Syntax

►►—SCHEDCMDException—string—————►►

Parameters

string

For commands scheduled by the action=command option on the DEFINE SCHEDULE server command, this parameter indicates the objects pattern to enable if the schedcmddisabled=yes option is specified. This parameter is case sensitive, and must match the command string on the Tivoli Storage Manager server schedule definition.

Example

Options file:

```
schedcmddisabled yes  
schedcmdexception "start echo hello, world!"
```

Related information

“Schedcmddisabled” on page 503

Schedlogmax

The schedlogmax option specifies the maximum size of the schedule log (dsmsched.log) and web client log (dsmwebcl.log), in megabytes.

This option causes the log files that get created for scheduler events (dsmsched.log) and web client events (dsmwebcl.log) to wrap around when they reach their maximum size. As scheduler and web client events are logged, log records are

added to the end of the log files 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 beginning of the file. The end of the wrapped log is indicated by a record saying END OF DATA.

When you set the `schedlogmax` option, scheduler and web client log messages are not saved in a prune file. If you want to prune logs and save the pruned log entries to another file, see the `schedlogretention` option.

If you change from log wrapping (`schedlogmax` option) to log pruning (`schedlogretention` option), all existing log entries are retained and the log is pruned using the new `schedlogretention` criteria.

If you change from log pruning (`schedlogretention` option) to log wrapping (`schedlogmax` option), all records in the existing logs are copied to a file containing the pruned entries. For example, log records pruned from the `dsmsched.log` file are copied to `dsmsched.pru`. Log records pruned from `dsmwebcl.log` are copied to `dsmwebcl.pru`. The existing logs (`dsmsched.log` and `dsmwebcl.log`) are emptied, and logging begins using the new log wrapping criteria.

If you simply 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.

If neither `schedlogmax` nor `schedlogretention` is specified, the error log can grow without any limit on its size. You must manually manage the log contents to prevent the log from depleting disk resources. When the log has been created with neither option specified, if you later issue a command and specify the `schedlogretention` option, the log is pruned using the retention value specified. When the log has been created with neither option specified, if you later issue a command and specify the `schedlogmax` option, the existing log is treated as if it was a pruned log. That is, the content of the `dsmsched.log` file is copied to a file called `dsmsched.pru`, the content of `dsmwebcl.log` is copied to a file called `dsmwebcl.pru`, and new log entries are created in `dsmsched.log` and `dsmwebcl.log`, and both files wrap when they reach their maximum size.

Note: If you specify a non-zero value for `schedlogmax` (which enables log wrapping), you cannot use the `schedlogretention` option to create pruned logs. Logs can be pruned or wrapped, but not both.

Logs created with the `schedlogmax` option contain a log header record that contains information similar to this example record:

```
LOGHEADERREC 661 104857600 IBM Tivoli Storage Manager 7.1.0.1 11080B Tue Sep 9 06:46:53 2014
```

Note that the dates and time stamps in the LOGHEADERREC text are not translated or formatted using the settings specified on the `dateformat` or `timeformat` options.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza.

You can also set this option on the **Client Preferences > Scheduler** tab in the GUI, by selecting **Enable scheduler log file wrapping** and by specifying a non-zero **maximum size** for the log file. To prevent log file wrapping, set the **maximum size** to zero. When the maximum wrapping is set to zero, clearing or setting the **Enable scheduler log file wrapping** option has no effect; log wrapping does not occur if the **maximum size** is set to zero.

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.

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, the client attempts to create it.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **Scheduler** tab, in the **Schedule Log** text box, in the Preferences editor.

Note: Set the DSM_LOG environment variable to name a directory where the log is to be placed. The directory specified must have permissions which allow write access from the account under which the client is run. The root directory is not a valid value for DSM_LOG.

Syntax

►—SCHEDLOGName— *filespec*—►

Parameters

filespec

Specifies the path and file name where you want to store schedule log information when processing scheduled work. If any part of the path you specify does not exist, the client attempts to create it.

If you specify a file name only, the file is stored in your current directory. The default is the current working directory with a file name of dsmsched.log. The dsmsched.log file *cannot* be a symbolic link.

For Mac OS X, if you specify a file name only, the file is stored in your default folder. The default directories are:

```
~/Library/Logs/tivoli/tsm  
/Library/Logs/tivoli/tsm
```

Examples

Options file:

```
SCHEDLOGN /Users/user1/Library/Logs/schedlog.jan  
schedlogname /home/mydir/schedlog.jan
```

Command line:

```
-schedlogname=/Users/user1/Library/Logs/schedlog.jan
```

Command line:

```
-schedlogname=/home/mydir/schedlog.jan
```

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

Related information

See “Errorlogname” on page 387 for more information on placement of the dsmsched.log file.

Schedlogretention

The schedlogretention option specifies the number of days to keep entries in the schedule log (dsmsched.log) and the web client log (dsmwebcl.log), and whether to save the pruned entries in another file.

The schedule log (dsmsched.log) is pruned when the scheduler starts and after a scheduled event completes. Pruned entries are written to a file called dsmsched.pru.

The web client log (dsmwebcl.log) is pruned during the initial start of the client acceptor daemon. Pruned entries are written to a file called dsmwebcl.pru.

If you change from log pruning (schedlogretention option) to log wrapping (schedlogmax option), all records in the existing log are copied to the pruned log (dmsched.pru and dsmwebcl.pru), and the existing logs (dmsched.log and dsmwebcl.log) are emptied, and logging begins using the new log wrapping criteria.

If you change from log wrapping (schedlogmax option) to log pruning (schedlogretention option), all existing log entries are retained and the log is pruned using the new schedlogretention criteria. Pruned entries are saved in their corresponding *.pru files.

If neither schedlogmax nor schedlogretention is specified, the logs can grow without any limit on their size. You must manually manage the log contents to prevent the log from depleting disk resources. When the log has been created with neither option specified, if you later issue a command and specify the schedlogretention option, the log is pruned using the retention value specified. When the log has been created with neither option specified, if you later issue a command and specify the schedlogmax option, the existing log is treated as if it was a pruned log. That is, the content of the dmsched.log file is copied to a file called dmsched.pru, the content of dsmwebcl.log is copied to dsmwebcl.pru, and new log entries are created in both dmsched.log and dsmwebcl.log, and both files wrap when they reach their maximum size.

Note: If you specify schedlogretention option to create pruned logs, you cannot specify the schedlogmax option. Logs can be pruned or wrapped, but not both.

Supported Clients


This option is valid for all clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza.

You can also set this option on the **Client preferences > Scheduler** tab in the GUI, by selecting **Prune old entries** and by specifying a value for **Prune entries older than**. Selecting the **Save pruned entries** option saves the pruned scheduler log entries in the dmsched.pru log file. Selecting **Save pruned entries** also saves web client log entries in the dsmwebcl.pru log file.

Syntax



Parameters

N or days

Specifies how long to wait before pruning the log.

N Do not prune the log. This permits the log to grow indefinitely. This is the default.

days

Specifies the number of days to keep log file entries before pruning. The range of values is zero through 9999.

D* or *S

Specifies whether to save the pruned entries. Use a space or comma to separate this parameter from the previous one.

D Discards the log entries when pruning the log. This is the default.

S Saves the log entries when pruning the log.

Pruned entries are copied to the file of pruned entries (`dsmsched.pru` or `dsmsched.pru`), which is stored in the same directory as the log.

Examples

Options file:

```
    schedlogretention 30 S
```

Command line:

```
-schedlogretention=30,S
```

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

Schedmode

The `schedmode` option specifies whether you want to use the polling mode (your client node periodically queries the server for scheduled work), or the prompted mode (the server contacts your client node when it is time to start a scheduled operation).

All communication methods can use the client polling mode, but only TCP/IP can use the server prompted mode.

This option applies only if you are using the TCP/IP communication method, and the **`schedule`** command is running.

Your administrator can specify that the server support both modes or just one mode. If your administrator specifies that both modes are supported, you can select either schedule mode. If your administrator specifies only one mode, you must specify that mode in your `dsm.sys` file or scheduled work is not processed.

If you specify prompted mode, you should consider supplying values for the `tcpclientaddress` and `tcpclientport` options in your `dsm.sys` file or on the `schedule` command; the client can then be contacted at either an address or a port of your choice (useful for client systems with multiple network interface cards).

Note:

1. When changing the setting of this option in the `dsm.sys` file you must stop and restart the scheduler service for the setting to take effect.
2. The server can also define this option.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza. You can set this option on the **Scheduler** tab, in the **Schedule Mode** section in 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:

1. If you use the **dsmc schedule** command and both `schedmode prompted` and `commethod V6Tcpip` are specified, the client and Tivoli Storage Manager server must be configured for IPv6. Additionally, the client host name must be set up for the IPv6 address.

Examples

Options file:

```
schedmode prompted
```

Command line:

```
-schedmod=po
```

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

Related reference:

“`Cadlistenonport`” on page 332

“`Tcpclientaddress`” on page 545

“`Tcpclientport`” on page 545

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 `dsm.sys` file within a server stanza for the scheduler. You can set this option on the **Scheduler** tab in the **Schedule Command** section in the Preferences editor.

Syntax



Parameters

No Specifies that the client does not disable the execution of restore and retrieve schedule operations. This parameter is the default.

Yes Specifies that the client disables the execution of restore and retrieve schedule operations.

Examples

Options file:

`schedrestretrdisabled yes`

Command line:

Does not apply.

Scrolllines

The `scrolllines` option specifies the number of lines of information that are displayed on your screen at one time.

Use this option when you set the `scrollprompt` option to *Yes*.

You can use the `scrolllines` option with the following commands only:

- **delete filespace**
- **query archive**
- **query backup**
- **query backupset**
- **query filespace**
- **query group**
- **query image**
- **query nas**
- **query node**
- **query options**

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client user-options file (`dsm.opt`). You can set this option in **Command Line > Number of lines to display** in the Preferences editor.

Syntax

►—SCROLLLines— *number* —►

Parameters

number

Specifies the number of lines of information that are displayed on your screen at one time. The range of values is 1 through 80; the default is 20.

Examples

Options file:

scrolllines 25

Command line:

-scroll=25

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

Scrollprompt

The scrollprompt option specifies whether you want the backup-archive client 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 group**
- **query image**
- **query nas**
- **query node**
- **query options**

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Command Line** tab, **Pause after displaying the following number of lines** field of the Preferences editor.

Syntax



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 on the screen:

Press 'Q' to quit, 'C' to continuous scroll, or 'Enter' to continue.

Examples

Options file:

```
scrollprompt yes
```

Command line:

```
-scrollp=yes
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the `dsm.opt` file unless overridden by the initial command line or by an option forced by the server.

Servename

In your `dsm.sys` file, the `servername` option specifies the name you want to use to identify a server and to begin a stanza containing options for that server. You can name and specify options for more than one server.

The following example demonstrates how to specify options for two different servers:

```
Servename    server_a
COMMMethod   TCPip
TCPPort      1500
TCPServeraddress server_hostname2.domain.company.com
PASSWORDAccess prompt
GRoups       tsm
USERS        sullivan mushock tallan
INCLExcl     /adm/tsm/backup.excl

Servename    server_b
COMMMethod   SHAREdmem
shmpor       1520
PASSWORDAccess generate
GRoups       system tsm
INCLExcl     /adm/tsm/archive.excl
```

In your client user-options file (`dsm.opt`), the `servername` option specifies which server, of those named in your `dsm.sys` file, to contact for backup-archive services. When specified in a client user-options file (`dsm.opt`) or on the command line, the `servername` option overrides the default server specified in your client system options file.

Note:

1. You cannot use the `servername` option to override the server that is specified for migration in your client system options file.
2. If the Tivoli Storage Manager server name changes or backup-archive clients are directed to a different Tivoli Storage Manager server, all clients must have a new password initialized for the new server name.

Supported Clients

This option is for all UNIX and Linux clients.

Options File

Place this option in the client user options file (`dsm.opt`) and in the client system options file (`dsm.sys`). In the `dsm.sys` file, the `servername` option is the beginning of a server stanza.

Do not modify this option in `dsm.opt` when you are running the Backup-Archive client in a command-line session or when you are running the Backup-Archive client GUI.

Syntax

►►—SErvername— —*servername*—►►

Parameters

servername

In your `dsm.sys` file, specify the name you want to assign to a particular server. In your client user-options file (`dsm.opt`) or on the command line, specify the name of the server you want to contact for backup-archive services. The value of *servername* in `dsm.opt` must match a *servername* value in `dsm.sys`, or the client cannot contact the server.

A server name is not case sensitive; it can have up to 64 characters.

Examples

Options file:

```
servername server_a
```

Command line:

```
-se=server_b
```

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

Sessioninitiation

Use the `sessioninitiation` option to control whether the server or client initiates sessions through a firewall. The default is that the client initiates sessions. You can use this option with the **schedule** command.

For the client scheduler, you do not need 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, the backup-archive client GUI, and web client GUI still attempts to initiate sessions.

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

Note: If you set the sessioninitiation option to serveronly, the client setup wizard and scheduler service are unable to authenticate to the Tivoli Storage Manager server. In this case, you can execute the scheduler from the command line (dsmc schedule) and enter the password for your node when prompted.

A similar problem can occur if an encryption key is required for backup operations. In this case, you can execute the scheduler from the command line (dsmc schedule) and enter the encryption key when prompted. After the password and encryption key are updated, you must restart the scheduler.

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

Note:

1. The Tivoli Storage Manager server can specify SESSIONINITiation=clientorserver or SESSIONINITiation=serveronly on the **register node** and **update node** commands. If the server specifies SESSIONINITiation=clientorserver, the client can decide which method to use. If the server specifies SESSIONINITiation=serveronly, all sessions are initiated by the server.
2. If sessioninitiation is set to serveronly, the value for the tcpclientaddress client option must be the same as the value for the HLAddress option of the **update node** or **register node** server command. The value for the tcpclientport client option must be the same as the value for the LLAddress option of the **update node** or **register node** server command.


Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the dsm.sys file within a server stanza. You can set this option on the **Scheduler** tab, **Session Initiation** field of the Preferences editor.

Syntax



Parameters

Client

Specifies that the client initiates sessions with the server by communicating on the TCP/IP port defined with the server option `TCPPORT`. This is the default. Server prompted scheduling can be used to prompt the client to connect to the server.

SERVEROnly

Specifies that the server will not accept client requests for sessions. All sessions must be initiated by server prompted scheduling on the port defined on the client with the `tcpclientport` option. Except for client acceptor daemon-managed schedulers, the command-line client, the backup-archive client GUI, and web client GUI still attempt to initiate sessions.

If the server `AUTHENTICATION` option is set to `LDAP`, do not set the client `sessioninitiation` option to `serveronly`; if you do, schedules cannot run.

Examples

Options file:

```
sessioninitiation serveronly
```

Command line:

```
schedule -sessioninitiation=serveronly
```

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

Related information

“Configuring the scheduler” on page 61

“`Tcpclientport`” on page 545

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 `dsm.sys` file must match the value specified for `shmport` in the server options file.

Supported Clients

This option is valid for AIX, HP-UX, Linux, and Solaris clients only.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza.

Syntax

```
▶▶—SHMPort— —port_number————▶▶
```

Parameters

port_number

Specifies the port number. You can specify a value from 1000 to 32767. The default value is 1510.

Examples

Options file:

shimport 1580

Command line:

Does not apply.

Showmembers

Use the showmembers option to display all members of a group.

You can use the showmembers option with the **query group**, and **restore group** commands.

The showmembers option is not valid with the inactive option. If you want to display members of a group that are not currently active, use the pitdate and pittime options.

Supported Clients

This option is valid for all UNIX and Linux clients except Mac OS X.

Syntax

►►—SHOWMembers—►►

Parameters

There are no parameters for this option.

Examples

Command line:

restore group /virtfs/* -pick -showmembers

Skipacl

The skipacl option allows you to include or exclude access control list (ACL) data during a backup or archive operation; by default, ACL data is included.

When this option is set to yes, the backup-archive client does not include ACL data when it backs up or archives files and directories. The default is no, which enables the ACL data to be included when objects are copied to the server. You should only set the skipacl to yes when ACLs are not defined on the file system, or when you are certain that you do not need the ACL data when the files are retrieved or restored.

Supported Clients

This option is valid for all UNIX and Linux clients. On Linux and AIX systems, setting skipacl to yes also omits the extended attributes.

Options File

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

Syntax



Parameters

No If you specify *No*, the ACL data is backed up. This is the default.

Yes

If you specify *Yes*, the ACL data is not backed up, and consequently, it cannot be restored. `skipacl=yes` overrides `skipaclupdatecheck` settings.

Examples

Options file:

`skipacl yes`

Skipaclupdatecheck

The `skipaclupdatecheck` option disables checksum and size comparisons of ACL data.

When set to *yes* (default is *no*), the backup-archive client will not perform checksum and size comparisons before or after backup and during incremental processing (ACL checksum from previous backup and current ACL) to detect ACL updates. However, current ACL data is backed up if the file is selected for backup due to other reasons. If only ACLs are updated on a file, the next incremental backup will not recognize this ACL update, and the file is not backed up.

Supported Clients

This option is valid for all UNIX and Linux clients.

Options File

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

Syntax



Parameters

No If you specify *No*, the client performs checksum and size comparisons of the ACL data, before and after backup and during incremental processing. This is the default.

Yes

If you specify Yes, the client does not perform checksum and size comparisons of the ACL data.

Examples

Options file:

```
skipaclup yes
```

Snapdiff

Using the `snapdiff` option with the **incremental** command streamlines the incremental backup process. The command runs an incremental backup of the files that were reported as changed by NetApp instead of scanning all of the volume for changed files.

You must configure a user ID and password on the backup-archive client to enable snapshot difference processing. For more information about setting up the `snapdiff` option, see “Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 106.

Use this option with an incremental backup of a NAS file server volume, instead of a simple incremental backup or an incremental backup with the `snapshotroot` option, whenever the NAS file server is running ONTAP 7.3.0, or later. Do not use the `snapdiff` and `snapshotroot` options together.

Restriction: Incremental backups with snapshot difference processing are only available with the 64-bit AIX backup-archive client and the Linux x86_64 backup-archive client.

The first time that you run an incremental backup with the snapshot difference option, a snapshot is created (the base snapshot) and a traditional incremental backup is run by using this snapshot as the source. The name of the snapshot that is created is recorded in the Tivoli Storage Manager server database. The initial incremental backup must complete without failure in order for the next backup operation to use snapshot difference processing.

The second time an incremental backup is run with this option, a newer snapshot is either created, or an existing one is used (depending on the value set for the `diffsnapshot` option) to find the differences between these two snapshots. The second snapshot is called the *diffsnapshot*, or differences snapshot. The client then incrementally backs up the files that are reported as changed, by NetApp, to the Tivoli Storage Manager server. The file system that you select for snapshot difference processing must be mounted to the root of the volume. You cannot use the `snapdiff` option for any file system that is not mounted to the root of the volume. After you backed up the data with the `snapdiff` option, the snapshot that was used as the base snapshot is deleted from the snapshot directory.

On AIX and Linux systems, the snapshot directory is in `.snapshot`.

The client does not delete any snapshots that it did not create.

When a snapshot-differential-incremental backup operation completes, the client ensures that only the most recently-registered base snapshot persists on the filer volume. All snapshots that are created by a Tivoli Storage Manager snapshot-differential-incremental backup begin with the characters "TSM_". If you use a snapshot tool other than the backup-archive client to produce snapshots,

ensure that you do not use the string "TSM_" at the beginning of the snapshot name. If the snapshot names begin with "TSM_", the files are deleted when Tivoli Storage Manager initiates the next snapshot-differential-incremental backup operation.

To run a snapshot-differential-incremental backup of read-only NetApp filer volumes, the `useexistingbase` option must be specified to prevent an attempt to create a snapshot on the read-only volume. Also, specify the name of the base snapshot to use (`basesnapshotname` option) and the name of the differential snapshot to use (`diffsnapshotname` option).

For NAS and N-Series file servers that are running ONTAP 7.3.0, or later, you can use the `createnewbase` option to back up any files that were skipped because of one of the following reasons:

- A file is excluded because the include-exclude file has an exclude rule in effect. A file is excluded when you did not change the include-exclude file, but you removed the rule that excluded the file. The NetApp API detects file changes only between two snapshots, not changes to the include-exclude file.
- If you added an include statement to the option file, that include option does not take effect unless NetApp detects that the file changes occurred. The client does not inspect each file on the volume during backup.
- You used the **`dsmdc delete backup`** command to explicitly delete a file from the Tivoli Storage Manager server inventory. NetApp does not detect that a file was manually deleted from the server. Therefore, the file remains unprotected in Tivoli Storage Manager storage until it is changed on the volume and the change is detected by NetApp, signaling the client to back it up again.
- Policy changes such as changing the policy from `mode=modified` to `mode=absolute` are not detected.
- The entire file space is deleted from the Tivoli Storage Manager inventory. This action causes the snapshot difference option to create a snapshot to use as the source, and runs a full incremental backup.
- A file is excluded from backup because the file name contains a character that is not in the 7 bit-ASCII character set. The `createnewbase` option creates a base snapshot and uses it as a source to run a full incremental backup. NetApp controls what constitutes a changed object.

Tip: You can use the `snapdiffhttps` option to run snapshot-differential-incremental backups of NetApp filers with a secure HTTPS connection. To successfully run snapshot-differential-incremental backups, previous releases of the backup-archive client required HTTP administrative access to be enabled on the NetApp filer. With the `snapdiffhttps` option, you can establish a secure administrative session with the NetApp filer regardless of whether HTTP administrative access is enabled on the filer.

In the list of options that are used by the traditional **`incremental`** command, the last column shows the interaction of each option with the `snapdiff` option. The following information describes the definitions of *valid*, *not valid*, and *no effect*:

Valid Processing runs normally when the option is used.

Not valid

If the option is used with the `snapdiff` option, an error message is generated.

No effect

The option can be used, but it is ignored.

Table 74. Incremental command: Related options

Option	Where specified	With snapdiff
asnodename "Asnodename" on page 319	Client system options file (dsm.sys) or command line.	Valid
automount "Automount" on page 328	Client options file (dsm.opt).	No effect
basesnapshotname "Basesnapshotname" on page 330	Client options file (dsm.opt) or command line.	Valid
changingretries "Changingretries" on page 333	Client system options file (dsm.sys) or command line.	No effect
compressalways "Compressalways" on page 338	Client options file (dsm.opt) or command line.	Valid
compression "Compression" on page 339	Client system options file (dsm.sys) within a server stanza, or command line.	Valid
createnewbase "Createnewbase" on page 342	Command line only.	Valid
diffsnapshot "Diffsnapshot" on page 354	Command line only.	Valid
diffsnapshotname "Diffsnapshotname" on page 355	Client options file (dsm.opt) or command line.	Valid
dirsonly "Dirsonly" on page 357	Command line only.	Valid
domain "Domain" on page 361	Client system options file (dsm.sys), client user-options file (dsm.opt), or command line.	Valid
efsdecrypt "Efsdecrypt" on page 376	Client system options file (dsm.sys), client user-options file (dsm.opt), or command line.	No effect
enablelanfree "Enablelanfree" on page 382	Client system options file (dsm.sys) or command line.	Valid
encryptiontype "Encryptiontype" on page 383	system-options file (dsm.sys) within a server stanza.	Valid
encryptkey "Encryptkey" on page 384	System-options file (dsm.sys) within a server stanza.	Valid
exclude.fs.nas "Exclude options" on page 389	Client system options file (dsm.sys).	No effect
filelist "Filelist" on page 402	Command line only.	Not valid
filesonly "Filesonly" on page 406	Command line only.	Valid
followsymboliclink "Followsymbolic" on page 407	Client options file (dsm.opt).	No effect
include.fs.nas "Include options" on page 421	Client system options file (dsm.sys) or command line.	No effect
incl excl "Incl excl" on page 420	Client system options file (dsm.sys).	Valid, but only when a file change is detected by NetApp.
incrbydate "Incrbydate" on page 436	Command line only.	Not valid

Table 74. Incremental command: Related options (continued)

Option	Where specified	With snapdiff
memoryefficientbackup "Memoryefficientbackup" on page 452	This option is allowed in both dsm.sys and dsm.opt, but the value in dsm.opt is ignored if it is also in dsm.sys. You can also place this option within a server stanza, or on the initial command line.	No effect
monitor "Monitor" on page 457	Command line only.	Not valid
nojournal "Nojournal" on page 463	Command line only.	Not valid
postsnapshotcmd "Postsnapshotcmd" on page 476	Client system options file (dsm.sys) or with the include.fs option.	Valid
preservelastaccessdate "Preservelastaccessdate" on page 479	Client user-options file (dsm.opt) or command line.	Valid
presnapshotcmd "Presnapshotcmd" on page 483	Client system options file (dsm.sys) or with the include.fs option.	Valid
removeoperandlimit "Removeoperandlimit" on page 488	Command line only.	Valid
skipaclupdatecheck "Skipaclupdatecheck" on page 518	Client options file (dsm.opt).	Valid
snapdiffhttps "Snapdiffhttps" on page 524	Command line only.	Valid
snapshotcachesize "Snapshotcachesize" on page 525	Client system options file (dsm.sys) or with the include.fs option.	No effect
snapshotproviderfs "Snapshotproviderfs" on page 526	System-options file (dsm.sys) within a server stanza or with the include.fs option.	Not valid
snapshotproviderimage "Snapshotproviderimage" on page 527	Client system options file (dsm.sys) or with the include.image option.	Not valid
snapshotroot "Snapshotroot" on page 528	Command line only.	Not valid
subdir "Subdir" on page 539	Client options file (dsm.opt) or command line.	Not valid
tapeprompt "Tapeprompt" on page 541	Client options file (dsm.opt) or command line.	Valid
toc "Toc" on page 552	Command line only.	Not valid
useexistingbase "Useexistingbase" on page 558	Command line only.	Valid
virtualfsname "Virtualfsname" on page 565	Command line only.	Not valid
virtualmountpoint "Virtualmountpoint" on page 566	Client system options file (dsm.sys).	Not valid

Supported Clients

This option is valid for the following clients:

- AIX 64-bit clients

- Linux x86_64 clients

Syntax

►—SNAPdiff—◄

Parameters

There are no parameters for this option.

Examples

Command line:

Perform a snapshot-differential-incremental backup of an NFS mounted file system /vol/vol1 hosted on the file server homestore.example.com, where /net/home1 is the mount point of /vol/vol1.

```
incremental -snapdiff -diffsnapshot=latest /net/home1
```

Important: For Linux clients, when you export the path of a volume on NetApp 7-mode filers, the export path must equal the volume name (for example, /vol_name), or the export path must begin with "/vol" and include one subdirectory level (for example, /vol/vol_name). Do not add a slash ("/") to the end of the export path.

Command line:

Run a one-time full incremental backup after detecting that the NetApp server has migrated to a unicode-enabled file server from a server that did not support unicode file names.

```
dsmc incremental -snapdiff -createnewbase=migrate /net/home1
```

Run a snapshot-differential-incremental backup after detecting that the NetApp server has migrated to a unicode-enabled file server from a server that did not support unicode file names. This command suppresses the warning message.

```
dsmc incremental -snapdiff -createnewbase=ign /net/home1
```

Perform a full incremental backup because you made some include or exclude changes:

```
dsmc incremental -snapdiff -createnewbase=yes /net/home1
```

Related concepts:

“Snapshot differential backup with an HTTPS connection” on page 167

“SnapMirror support for NetApp snapshot-assisted progressive incremental backup (snapdiff)” on page 110

Related tasks:

“Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 106

Related reference:

“Snapdiffhttps” on page 524

“Basesnapshotname” on page 330

“Diffsnapshotname” on page 355

“Useexistingbase” on page 558

“Diffsnapshot” on page 354

“Set Password” on page 734

Snapdiffhttps

Specify the `snapdiffhttps` option to use a secure HTTPS connection for communicating with a NetApp filer during a snapshot differential backup.

When you specify this option, the backup-archive client can establish a secure administrative session with the NetApp filer regardless of whether HTTP administrative access is enabled on the NetApp filer.

Important: The default communication protocol that the backup-archive client uses to establish the administrative session with the NetApp filer is HTTP. To use a secure HTTPS connection, you must specify the `snapdiffhttps` option whenever you run a snapshot differential backup.

Restrictions:

The following restrictions apply to snapshot differential backups with HTTPS:

- The HTTPS connection is used only to securely transmit data over the administrative session between the backup-archive client and the NetApp filer. The administrative session data includes information such as filer credentials, snapshot information, and file names and attributes that are generated by the snapshot differencing process. The HTTPS connection is not used to transmit normal file data that is accessed on the filer by the client through file sharing. The HTTPS connection also does not apply to normal file data transmitted by the client to the Tivoli Storage Manager server through the normal Tivoli Storage Manager client/server protocol.
- The `snapdiffhttps` option does not apply to vFilers because the HTTPS protocol is not supported on the NetApp vFiler.
- The `snapdiffhttps` option is available only by using the command-line interface. It is not available for use with the backup-archive client GUI.

Supported Clients

This option is valid for the following clients:

- AIX 64-bit clients
- Linux x86_64 clients

Options File

This option is valid only on the command-line interface. You cannot enter it in a client options file.

Syntax

▶▶—SNAPDIFFHTTPS—————▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

Issue the following command on an AIX or Linux system, with an NFS mounted file system `/vol/vol1` hosted on the file server `homestore.example.com`, where `/net/home1` is the mount point of `/vol/vol1`.

```
dsmc incr /net/home1 -snapdiff -snapdiffhttps
```

Related concepts:

“Snapshot differential backup with an HTTPS connection” on page 167

Related reference:

“Snapdiff” on page 519

Snapshotcachesize

Use the `snapshotcachesize` option to specify an appropriate size to create the snapshot.

The size estimation is needed for storing the original data blocks for modified and deleted data for the point in time when the snapshot was taken.

For snapshot-based file backup or archive, use the `snapshotcachesize` option with the `include.fs` option, or in the server stanza in the `dsm.sys` file.

For snapshot-based image backups, use the `snapshotcachesize` option with the **backup image** command, the `include.image` option, or in your `dsm.sys` file.

Supported Clients

This option is valid for AIX and Linux clients *only*. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Place this option in the server stanza in the `dsm.sys` file. You can set this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

►►—SNAPSHOTCACHESize— *size* —————►►

Parameters

size

Specifies an appropriate size to create the snapshot for storing the original data blocks for modified and deleted data for the point in time when the snapshot was taken. The value is the percent of the file system size that is changed due to file system activity. The range of values is 1 to 100 percent. For AIX JFS2 and Linux the default value is 100 percent of the file system size. If a sufficient amount of free space is not available to create the snapshot, the command fails with an error message. You can then either increase the size of the volume group or retry the operation. If based on your experience with your AIX JFS2 file system activity, you find that a snapshot size of 100 percent is not necessary, you can fine-tune the value.

Examples

Options file:

```
snapshotcachesize 95
AIX only: include.fs /kalafs1
          snapshotproviderfs=JFS2 snapshotcachesize=95
AIX only: include.image /kalafs2
          snapshotcachesize=95
Linux only: include.image /linuxfs1
           snapshotcachesize=100
```

Command line:

```
-snapshotcachesize=95
```

Related information

See “Include options” on page 421 for more information about `include.fs`.

Snapshotproviderfs

Use the `snapshotproviderfs` option to enable snapshot-based file backup and archive operations, and to specify a snapshot provider.

You must be a root user to perform a snapshot-based file backup or archive operation. If you are not a root user, the operation fails with an error message.

Supported Clients

This option is valid for AIX clients only. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Specify this option in the server stanza of the system-options file, `dsm.sys`, to enable snapshots for all JFS2 file systems on the client. You can override the client-wide option for a specific operation by specifying this option on the command line for the backup and archive commands. You can also override the client-wide option for a specific file system by using the `include.fs` statement in the `dsm.sys` file. You can also set this option using the Preferences editor.

Syntax

►►—SNAPSHOTPROVIDERFS— —*value*—————►►

Parameters

value

Specifies one of the following values:

JFS2

Specifies that you want to perform a snapshot-based file backup or archive while the file system is available to other system applications. Valid for JFS2 file systems on AIX clients *only*.

NONE

Specifies that no snapshots should be used. A file backup or archive operation is performed using the specified file system. This is the default.

Examples

Options file:

```
snapshotproviderfs JFS2
include.fs /kalafs1 snapshotproviderfs=JFS
```

Command line:

```
-SNAPSHOTPROVIDERFs=JFS2
```

Snapshotproviderimage

Use the `snapshotproviderimage` option to enable snapshot-based image backup, and to specify a snapshot provider.

You must be a root user to perform a snapshot-based image backup operation. If you are not a root user, the operation fails with an error message.

Supported Clients

This option is valid for AIX and Linux clients only. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Specify this option in the server stanza of the system-options file, `dsm.sys`, to enable snapshots for all the file systems on the client. You can override the client-wide option for a specific operation by specifying this option on the command line for the **backup image** command. You can also override the client-wide option for a specific file system using the `include.image` statement in the `dsm.sys` file. You can also set this option using the Preferences editor.

Syntax

►►—SNAPSHOTPROVIDERImage— *value* —————►►

Parameters

value

Specifies one of the following values:

JFS2

Specifies that you want to perform a snapshot-based image backup while the file system is available to other system applications. This is the default for JFS2 file systems. Valid for AIX clients *only*.

LINUX_LVM

Specifies that you want to perform a snapshot-based image backup while the file system is available to other system applications. This is the default for file systems residing on logical volumes created by the Linux Logical Volume Manager. Valid for Linux clients *only*.

NONE

Specifies that you do not want to perform a snapshot-based image backup operation. This performs a static image backup operation using the specified file system. This is the default for file systems other than AIX JFS2 and Linux LVM.

Examples

Options file:

```
snapshotprovideri JFS2  
include.image /kalafs1 snapshotprovideri=JFS2
```

Command line:

```
-SNAPSHOTPROVIDERImage=NONE
```

Snapshotroot

Use the `snapshotroot` option with the **incremental**, **selective**, or **archive** commands with an independent software vendor application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server.

The `snapshotroot` option can be used to back up NFS mounted file systems. Both the backup specification (source) and the `snapshotroot` value can be an NFS mounted file specification. For example, the `snapshotroot` option can be used to backup an NFS file system that is hosted on a network-attached storage (NAS) that supports snapshot.

This option should be used with an incremental backup of a NAS file server volume instead of a simple incremental or incremental with `snapshotroot` option whenever the NAS file server is running ONTAP V7.3 for performance reasons. The `snappdiff` and `snapshotroot` options should not be used together.

In the following example, filesystem `test495` is NFS-mounted from a NAS file server `philo` and `/philo/test945/.snapshot/backupsnap` represents the snapshot that is created at the NAS file server.

```
dsmc i /philo/test945 -snapshotroot=/philo/test945/.snapshot/backupsnap
```

You can also specify a directory with the `snapshotroot` option when you backup each file set as a separate file space.

The `snapshotroot` option does not provide any facilities to take a volume snapshot, only to manage data that is created by a volume snapshot.

For example, consider an application that takes a snapshot of the `/usr` file system and mounts it as `/snapshot/day1`. If you back up this data by using the following command, a unique file space that is called `/snapshot/day1` is created on the server.

```
dsmc incremental /snapshot/day1
```

However, you might want to associate the snapshot data with the data already processed for the `/usr` file system. Using the `snapshotroot` option, you can associate the data with the file space corresponding to the `/usr` file system on the Tivoli Storage Manager server:

```
dsmc incremental /usr -snapshotroot=/snapshot/day1
```

On a subsequent day, you can back up a snapshot that was written to an alternative location, but managed under the same file space on the server:

```
dsmc incremental /usr -snapshotroot=/snapshot/day2
```

You can perform incremental backups, selective backups, or archives of a single directory, directory structure, or single file by using the `snapshotroot` option. In all

instances, the `snapshotroot` option must identify the root of the logical volume that was created by the snapshot. For example:

```
dsmc incremental /usr/dir1/* -subdir=yes
  -snapshotroot=/snapshot/day1
dsmc selective /usr/dir1/sub1/file.txt
  -snapshotroot=/snapshot/day1
dsmc archive /usr/dir1/sub1/*.txt
  -snapshotroot=/snapshot/day1
```

If you want to include or exclude specific file specifications, the include and exclude statements should contain the name of the file system that was the source of the snapshot (the `/usr` file system), and not the name of the target of the snapshot (`/snapshot/day1`). Doing this allows you to preserve a set of include and exclude statements regardless of the name of the logical volume to which the snapshot is written. The following are examples of include and exclude statements.

```
include /usr/dir1/*.txt lyrmgmtclass
exclude /usr/mydocs/*.txt
```

The following include-exclude statements are not valid because they contain the name of the snapshot:

```
include /snapshot/day1/dir1/*.txt lyrmgmtclass
exclude /snapshot/day1/mydocs/*.txt
```

You must use the `snapshotroot` option with a single file specification for an incremental, selective, or archive operation. You cannot specify multiple file specifications or no file specifications. For example, these commands are valid:

```
dsmc incremental /usr -snapshotroot=/snapshot/day1
dsmc incremental /usr/dir1/* -snapshotroot=/snapshot/day1
```

The following command is invalid because it contains two file specifications:

```
dsmc incremental /usr/dir1/* /home/dir2/*
  -snapshotroot=/snapshot/day1
```

The following command is invalid because it contains no file specification:

```
dsmc incremental -snapshotroot=/snapshot/day1
```

Notes:

1. Ensure that the `snapshotroot` option references a snapshot of the correct volume. Ensure that `snapshotroot` location refers to the root of the snapshot. If these rules are not followed, unintended results, such as files that expire incorrectly, can result.
2. If you specify the `filelist` option and the `snapshotroot` option, all files that are specified in the `filelist` option are assumed to be in the same file system. If there are entries in the `filelist` in a different file system, they are skipped and an error is logged. If the `filelist` contains files that were created in the file system after the snapshot was taken, these entries are also skipped, and an error is logged.
3. You cannot use the `snapshotroot` option with the `snapdiff` option.
4. You can use the `snapshotroot` option with the `preschedulecmd` and `postschedulecmd` options, or in an automated script that you run with the client scheduler.

Supported Clients

This option is valid for the following clients:

- UNIX and Linux clients except Mac OS X.

Syntax

►►—SNAPSHOTRoot =— —*snapshot_volume_name*—————►◄

Parameters

snapshot_volume_name

Specifies the root of the logical volume that is created by the independent software vendor snapshot application.

Examples

Command line:

```
dsmc incremental /usr -SNAPSHOTRoot=/snapshot/day1
```

Srvoptsetencryptiondisabled

The `srvoptsetencryptiondisabled` option allows the client to ignore encryption options in a client options set from the Tivoli Storage Manager server.

If the option is set to yes in the client options file, the client ignores the following options in a client options set from the server:

- `encryptkey generate`
- `exclude.encrypt`
- `include.encrypt`

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client options file (`dsm.sys`) within a server stanza.

Syntax

►►—SRVOPTSETENCryptiondisabled—

<i>no</i>
<i>yes</i>

—————►◄

Parameters

yes

The backup-archive client ignores the values of the listed encryption options in a client options set from the Tivoli Storage Manager server.

no

The backup-archive client processes the setting of the listed encryption options in a client options set from the Tivoli Storage Manager server. This is the default.

Examples

Options file:

```
srvoptsetencryptiondisabled no
```

Command line:

Does not apply.

Srvprepostscheddisabled

The `srvprepostscheddisabled` option specifies whether to prevent the pre-schedule and post-schedule commands specified by the Tivoli Storage Manager administrator from executing on the client system, when performing scheduled operations.

The `srvprepostscheddisabled` option can be used in conjunction with the `schedcmddisabled` and `srvprepostscheddisabled` options to disable the execution of any unwanted operating system command by a Tivoli Storage Manager administrator on a client node.


Supported Clients

This option is valid for all backup-archive clients that use the Tivoli Storage Manager client scheduler. The server cannot define this option.

Options File

Place this option in the `dsm.sys` file within a server stanza for the scheduler. You can set this option on the **Scheduler** tab of the Preferences editor, in the **Schedule Command** section.

Syntax



Parameters

No Specifies that the client 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 overrides the corresponding command defined in the client option file. This is the default.

Yes

Specifies that the client 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 `srvprepostscheddisabled` options.

Examples

Options file:

srvprepostscheddisabled yes

Command line:

Does not apply.

Srvprepostsnapdisabled

The `srvprepostsnapdisabled` option specifies whether to prevent the pre-snapshot and post-snapshot commands specified by the Tivoli Storage Manager administrator from executing on the client system, when performing scheduled image snapshot backup operations.

The `srvprepostsnapdisabled` option can be used in conjunction with the `schedcmddisabled` and `srvprepostsnapdisabled` options to disable the execution of any unwanted operating system command by a Tivoli Storage Manager administrator on a client node.


Supported Clients

This option is valid for Linux clients that support the image snapshot backup command. The server cannot define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the `dsm.sys` file within a server stanza for the scheduler. You can set this option on the **Snapshot** tab of the Preferences editor, in the **Snapshot Options** section.

Syntax



Parameters

No Specifies that client 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 overrides the corresponding command defined in the client option file. This is the default.

Yes

Specifies that the client does not allow 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 `srvprepostsnapdisabled` options.

Examples

Options file:

srvprepostsnapdisabled yes

Command line:

Does not apply.

Ssl

Use the `ssl` option to enable Secure Sockets Layer (SSL) to provide secure client and server communications. When the backup-archive client communicates with a Tivoli Storage Manager server V7.1.7 and earlier levels, and V8.1.1 and earlier V8 levels, it determines whether SSL is enabled. When the backup-archive client communicates with a Tivoli Storage Manager server V7.1.8 and later V7 levels, and V8.1.2 and later levels, SSL is always used and this option controls whether object data is encrypted or not. For performance reasons, it might be desirable to not encrypt the object data.


Supported Clients

This option is valid for all supported clients.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can also set this option on the **Communication** tab of the Preferences editor.

Syntax



Parameters for communicating with a Tivoli Storage Manager server V7.1.7 and earlier levels, and V8.1.1 and earlier V8 levels.

No Specifies that the backup-archive client does not use SSL to encrypt information. *No* is the default.

Yes

Specifies that the backup-archive client uses SSL to encrypt information.

To enable SSL, specify `SSL Yes` and change the value of the `TCPPORT` option. Changing the value of the `TCPPORT` option is generally necessary because the Tivoli Storage Manager server is typically set up to listen for SSL connections on a separate port.

Parameters for communicating with a Tivoli Storage Manager server V7.1.8 and later V7 levels, and V8.1.2 and later levels.

No Specifies that the backup-archive client does not use SSL to encrypt object data when communicating with the server. All other information is encrypted. *No* is the default.

Yes

Specifies that the backup-archive client uses SSL to encrypt all information, including object data, when communicating with the server.

To use SSL for all data, specify `SSL Yes`.

Examples

Options file:

ssl yes

Command line:

Does not apply.

Related information

“Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer” on page 68.

“Sslrequired” on page 536

“Tcpport” on page 547

Sslacceptcertfromserv

Use the `sslacceptcertfromserv` option to control whether the backup-archive client or the API application accept and trust the Tivoli Storage Manager server's Secure Sockets Layer (SSL) public certificate the first time they connect. This option applies only the first time that the backup-archive client or the API application connects to the Tivoli Storage Manager server. When the SSL public certificate is accepted, future changes to the certificate are not automatically accepted, and must be manually imported to the backup-archive client. You can use this option to connect only to a Tivoli Storage Manager server V7.1.8 and later.

Supported Clients

This option is valid for all supported clients.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza.

Syntax



Parameters

Yes

Specifies that the backup-archive client does automatically accept the Tivoli Storage Manager server's public certificate. Yes is the default.

No Specifies that the backup-archive client does not automatically accept the Tivoli Storage Manager server's public certificate.

To disable `SSLACCEPTCERTFROMSERV`, specify `sslacceptcertfromserv no`.

Examples

Options file:

sslacceptcertfromserv no

Command line:
Does not apply.

Related information

“Ssl” on page 533

“Sslrequired” on page 536

Ssldisablelegacytls

Use the `ssldisablelegacytls` option to disallow the use of SSL protocols that are lower than TLS 1.2.

Supported Clients

This option is valid for all supported clients.

Options File

Place this option in the `dsm.sys` file. You can also set this option in the GUI by selecting the **Require TLS 1.2 or above** check box on the **Communication** tab of the Preferences editor. You cannot set this option on the command line.

Syntax



Parameters

No Specifies that the backup-archive client does not require TLS 1.2 for SSL sessions. It allows connection at TLS 1.1 and lower SSL protocols. When the backup-archive client communicates with a Tivoli Storage Manager server V7.1.7 and earlier levels, and V8.1.1 and earlier V8 levels, No is the default.

Yes Specifies that the backup-archive client requires that all SSL sessions use TLS 1.2 (or higher) protocol. When the backup-archive client communicates with a Tivoli Storage Manager server V7.1.8 and later V7 levels, and V8.1.2 and later levels, Yes is the default.

Examples

Options file:
`ssldisablelegacytls yes`

Command line:
Does not apply.

Related information

“Ssl” on page 533.

“Sslrequired” on page 536

“Tcpport” on page 547

Sslfipsmode

The `sslfipsmode` option specifies whether the client uses SSL Federal Information Processing Standards (FIPS) mode for Secure Sockets Layer (SSL) communications with the server. The default is `no`.


Supported clients

This option is supported on all clients.

Options File

Set this option in the client options file. You cannot specify it as a command-line parameter and you cannot set this option in a client options set.

Syntax



Parameters

No Specifies that the client does not use SSL FIPS mode for secure communications with the server. SSL in FIPS mode is supported only by Tivoli Storage Manager Version 6.3 and newer versions. Set this client option to `no` if the client uses SSL to connect to a server that is not at V6.3, or newer.

Yes

Specifies that the client uses SSL FIPS mode for secure communications with the server. Setting this option to `yes` restricts SSL session negotiation to use only FIPS-approved cipher suites. SSL FIPS mode is only supported by Tivoli Storage Manager V6.3 (or newer).

Example

To enable SSL FIPS mode on the client:

```
SSLFIPSMODE yes
```

Sslrequired

The `sslrequired` option specifies the conditions when SSL is or is not required when the client logs on to the Tivoli Storage Manager server or storage agents. To actually enable SSL so client-to-server and client-to-storage-agent communications are secure, you must set the client `ssl` option to `yes`. When communicating with the Tivoli Storage Manager server V7.1.8 and later V7 levels, and V8.1.2 and later levels, this option no longer applies since SSL is always used.

Supported Clients

This option is supported on all clients.

Options File

Place this option in the client options file or in the GUI, on the Communications tab. You cannot set this option on the command line.

Syntax



Parameters

Default

This setting indicates that SSL is required to secure communications between the client and server, and client and storage agents, if AUTHENTICATION=LDAP is set on the server. To secure communications by using SSL, you must also set `ssl=yes` on the client.

If AUTHENTICATION=LOCAL is set on the server, this setting indicates that SSL is not required. Even though SSL is not required when AUTHENTICATION=LOCAL and `sslrequired=default`, you can still use SSL by setting the client `ssl` option to `yes`.

Yes

Indicates that SSL is always required to secure communications between the client and server, and between the client and storage agents. `sslrequired=yes` has no dependency on the server AUTHENTICATION option. If you set `sslrequired=yes` on the client, you must also set `ssl=yes` on the client.

No Indicates that you do not require SSL to be used to secure communications between the client and server or between the client and storage agents. Choose this option only if you use a virtual private network or other method to secure your session communications. You can still enable SSL by setting `ssl=yes` on the client; but `sslrequired=no` specifies that SSL is not a prerequisite.

SERVERonly

Indicates that SSL is required for client-to-server communications and not for server-to-storage agent communications. To use SSL for client to server communications, set `sslrequired=serveronly` and `ssl=yes`. The server setting for the AUTHENTICATION option can be either LOCAL or LDAP.

For client to storage agent communications, use the client `lanfreessl` option to enable SSL.

The following table describes the situations under which authentication succeeds or fails, depending on the settings of the SSLREQUIRED option on the server, and client, and the setting of the `ssl` option on the client. The table results assume that valid credentials are supplied.

Table 75. Effects of server and client SSL settings on success or failure of login attempts

SSLREQUIRED option (server setting)	sslrequired option (client setting)	ssl option (client setting)	Authentication success or failure
Yes	Yes	Yes	Authentication succeeds
Yes	Yes	No	Authentication fails; the client rejects the session
Yes	No	Yes	Authentication succeeds

Table 75. Effects of server and client SSL settings on success or failure of login attempts (continued)

SSLREQUIRED option (server setting)	sslrequired option (client setting)	ssl option (client setting)	Authentication success or failure
Yes	No	No	Authentication fails; the server rejects the session
No	Yes	Yes	Authentication succeeds
No	Yes	No	Authentication fails; the client rejects the session
No	No	Yes	Authentication succeeds
No	No	No	Authentication succeeds

The following text describes how setting SSLREQUIRED=DEFAULT and SSLREQUIRED=SERVERONLY on the server affects the ssl option on the client.

If the server sets SSLREQUIRED=DEFAULT and AUTHENTICATION=LDAP, the client must set ssl=yes or authentication fails.

If the server sets SSLREQUIRED=DEFAULT and AUTHENTICATION=LOCAL, the client can set ssl=yes or ssl=no.

If the server sets SSLREQUIRED=SERVERONLY, you must set ssl=yes on the client. The client lanfreessl option can be set to yes, to secure communications with a storage agent, or to no if secure communications with storage agents is not needed.

Examples

Options file:

```
sslrequired yes
sslrequired no
sslrequired default
sslrequired serveronly
```

Command line:

Not applicable; you cannot set this option on the command line.

Stagingdirectory

The stagingdirectory option defines the location where the client can keep any data that it generates to perform its operations. The data is deleted when processing is complete.

This option is used during snapshot difference backup operations.

Supported Clients

This option is valid for AIX and Linux clients. The server can also define this option.

Options File

Place this option in the client options file (dsm.opt). When stagingdirectory is specified on the command line, it overrides the values that are specified in the options file.

Syntax

►—STAGINGDIRectory—path—►

Parameters

path

Specifies the directory path where the client writes staging data. If you do not specify a staging directory, the client stores temporary data in the temporary file system (typically /tmp).

Examples

Options file:

```
stagingdirectory /usr/tsmdata
```

```
stagingdirectory /private/tmp
```

Command line:

```
-stagingdir="/tmp/tsmtempdata"
```

Related reference:

“Diffsnapshot” on page 354

“Snapdiff” on page 519

Subdir

The `subdir` option specifies whether you want to include subdirectories of named directories for processing.

You can use the `subdir` option with the following commands:

- **archive**
- **delete archive**
- **delete backup**
- **incremental**
- **query archive**
- **query backup**
- **restore**
- **restore backupset**
- **restore group**
- **retrieve**
- **selective**

If you set the `subdir` option to `yes` when backing up a specific path and file, the backup-archive client recursively searches all of the subdirectories under that path, and looks for any instances of the specified file that exist under any of those subdirectories. For example, assume that a file called `myfile.txt` exists on a client in the following directories:

```
//myfile.txt  
/dir1/myfile.txt  
/dir1/dir_a/myfile.txt  
/dir1/dir_b/myfile.txt
```

Performing a selective backup of that file, as follows, backs up all four instances of `myfile.txt`:

```
dsmc sel /myfile.txt -subdir=yes
```

Similarly, the following command displays all instances of `myfile.txt` if you specify `subdir=yes` in the client options file or in a client options set.

```
dsmc rest "/Users/van/dir1/*" /Users/mike/ -su=yes
dsmc rest "/Users/van/dir1/file*" /Users/mike/ -su=yes
dsmc rest "/Users/van/dir1/file1*" /Users/mike/ -su=yes
```

To restore the structure:

```
/path2/dir1
/path2/dir1/file1
/path2/dir1/dir2
/path2/dir1/dir2/file1
```

enter any of the following commands:

```
dsmc rest "/path/dir1/*" /path2/ -su=yes
dsmc rest "/path/dir1/file*" /path2/ -su=yes
dsmc rest "/path/dir1/file1*" /path2/ -su=yes
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the `dsm.opt` file unless overridden by the initial command line or by an option forced by the server.

Related information

“Preservepath” on page 480

Tapeprompt

The `tapeprompt` option specifies whether you want the backup-archive client 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 backup-archive client GUI, the Media Mount dialog can display the Information Not Available value in the Device and Volume Label fields if you perform a standard (also known as classic) restore or retrieve operation. This value means that this information is only available for no-query restore or retrieve operations; not a standard restore or retrieve operation. The **Device** field displays the name of the device on which to mount the media needed to process an object. The **Volume Label** field displays the name of the volume needed to process an object.

Tape prompting does not occur during a scheduled operation regardless of the setting for the `tapeprompt` option.

The `tapeprompt` option can be used with the following commands:

- **archive**
- **delete archive**
- **delete backup**
- **incremental**
- **restore**
- **retrieve**
- **selective**

Note: The server can also define this option.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **General** tab, **Prompt before mounting tapes** check box 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.

Note: For API applications, this permits backup directly to tape.

Yes

You are prompted when a tape is required to back up, archive, restore, or retrieve data. At the prompt, you can wait for the appropriate tape to be mounted, always wait for a tape to be mounted, skip a particular object, skip all objects on a single tape, skip all objects on all tapes, or cancel the entire operation.

Examples

Options file:

```
tapeprompt yes
```

Command line:

```
-tapep=yes
```

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

Tcpadminport

Use the `tcpadminport` option to specify a separate TCP/IP port number on which the server waits for requests for administrative client sessions, allowing secure administrative sessions within a private network.

The client `tcpadminport` setting depends on how the Tivoli Storage Manager server `tcpadminport` and `adminonclientport` options are configured. The Tivoli Storage Manager server has a `tcpadminport` setting that indicates on which port the server listens for administrative sessions, and the `adminonclientport` setting, which can be either `yes` or `no`.

If `tcpadminport` is not set on the server, then administrative sessions are allowed on the same port as client sessions.

If `tcpadminport` is set on the server, then administrative sessions are allowed on the port specified by that setting. In this case, if `adminonclientport yes` is in effect, then administrative sessions can connect on either the regular client port or the port specified by `tcpadminport`. If `adminonclientport no` is in effect, then administrative sessions can connect only on the port specified by `tcpadminport`.

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the `dsm.sys` file within a server stanza. You can set this option on the **Communication** tab, in the **Admin Port** field in the Preferences editor.

Syntax

►►—TCPADMINPort—*admin_port_address*—►►

Parameters

admin_port_address

Specifies the port number of the server. The default value is the value of the `tcpport` option.

Examples

Options file:

`tcpadminport 1502`

Tcpbuffsize

The `tcpbuffsize` option specifies the size of the internal TCP/IP communication buffer used to transfer data between the client node and server. Although it uses more memory, a larger buffer can improve communication performance.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **Communication** tab, in the **Buffer Size** field in the Preferences editor.

Syntax

►►—TCPBuffsize— *size* —►►

Parameters

size

Specifies the size, in kilobytes, that you want to use for the internal TCP/IP communication buffer. The range of values is 1 through 512; the default is 32.

Depending on the operating system communication settings, your system might not accept all values in the range of 1 through 512.

Examples

Options file:

tcpb 32

Command line:

-tcpbuffsize=32

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

Tcpcadaddress

The `tcpcadaddress` option specifies a TCP/IP address for `dsmcad`. Normally, this option is not needed. Use this option only if your client node has more than one TCP/IP address, or if TCP/IP is not the default communication method.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the `dsm.sys` file within a server stanza.

Syntax

►—TCPCADAddress— *—cad_address—*◄◄

Parameters

cad_address

Specifies a TCP/IP Internet domain name or a numeric IP address. If you specify an IPv6 addresses, you must specify the `commethod V6Tcpip` option.

Examples

Options file:

tcpcada dsmclnt.example.com

Command line:

-tcpcadaddress=192.0.2.0

-tcpcadaddress=mycompany.example.com

-tcpcadaddress=2001:0DB8:0:0:0:0:0:0

This option is valid only on the initial command line of the `dsmcad` program. It is not valid with other `dsm` modules.

Related information

See “`Commmethod`” on page 335 to determine if your client node has more than one TCP/IP address, or if TCP/IP is not the default communication method.

Tcpclientaddress

The `tcpclientaddress` option specifies a TCP/IP address if your client node has more than one address, and you want the server to contact an address other than the one that was used to make the first server contact.

The server uses this address when it begins the server prompted scheduled operation.

Use this option only if you use the prompted parameter with the `schedmode` option.

If `sessioninitiation` is set to `serveronly`, the value for the `tcpclientaddress` client option should be the same as the value for the `HLAddress` server setting.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the `dsm.sys` file *within* a server stanza. You can set this option on the **Scheduler** tab, **Your TCP/IP address** field of the Preferences editor.

Syntax

►►—TCPCLIENTAddress— —*client_address*—►►

Parameters

client_address

Specifies the TCP/IP address you want the server to use to contact your client node. Specify a TCP/IP Internet domain name or a numeric IP address. The numeric IP address can be either a TCP/IPv4 or TCP/IPv6 address. You can only use IPv6 addresses if you specified the `commethod V6Tcpip` option.

Examples

Options file:

For Mac clients:

```
tcpclienta dsmclnt.mycompany.mydomain.com
or
tcpclienta 192.0.2.21
```

Options file:

For AIX, HP, Linux, and Solaris clients:

```
tcpclienta dsmclnt.example.com
```

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

Tcpclientport

The `tcpclientport` option specifies a TCP/IP port number for the server to contact the client when the server begins the server prompted scheduled operation.

Use this option only if you specify the prompted parameter with the `schedmode` option.

If `sessioninitiation` is set to `serveronly`, the value for the `tcpclientport` client option should be the same as the value for the `LLAddress` server option.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the `dsm.sys` file within a server stanza. You can set this option on the **Scheduler** tab, in the **Your TCP/IP port** field in the Preferences editor.

Syntax

►—TCPCLIENTPort— *—client_port_address—*—————►

Parameters

client_port_address

Specifies the TCP/IP port address you want the server to use to contact your client node. The range of values is 1 through 32767; the default is 1501.

Examples

Options file:

`tcpclientp 1502`

Command line:

`-tcpclientport=1492`

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

Tcpnodelay

The `tcpnodelay` option specifies whether the client disables the delay of sending successive small packets on the network, per transaction.

Change the value from the default of `yes` only under one of the following conditions:

- You are directed to change the option by IBM technical support.
- You fully understand the effects of the TCP Nagle algorithm on network transmissions. Setting the option to `no` enables the Nagle algorithm, which delays sending small successive packets.

Supported Clients

This option is valid for all UNIX and Linux clientst.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza. You can set this option on the **Communication** tab in the Preferences editor. Select **Send transaction to the server immediately**.

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 the Tivoli Storage Manager server. You can obtain this address from your administrator.


Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza. You can set this option on the **Communication** tab, in the **Server Port** field in the Preferences editor.

Syntax



Parameters

port_address

Specifies the TCP/IP port address that is used to communicate with a server. The range of values is 1 through 32767; the default is 1500.

Examples

Options file:

tcpp 1501

Command line:

Does not apply.

Tcpserveraddress

The tcpserveraddress option specifies the TCP/IP address for a Tivoli Storage Manager server. You can obtain this server address from your administrator.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza. You can set this option on the **Communication** tab, in the **Server Address** field in the Preferences editor.

If this option is not specified, the client attempts to contact a server running on the same computer as the backup-archive client.

Syntax

►—TCPServeraddress— *server_address*—►

Parameters

server_address

Specifies a 1 to 64 character TCP/IP address for a server. Specify a TCP/IP domain name or a numeric IP address. The numeric IP address can be either a TCP/IP v4 or TCP/IP v6 address. You can only use IPv6 addresses if you specified the commethod V6Tcpip option.

Examples

Options file:

tcps dsmchost.example.com

Command line:

Does not apply.

Tcpwindowsize

Use the tcpwindowsize option to specify, in kilobytes, the size you want to use for the TCP/IP sliding window for your client node.

The sending host cannot send more data until it receives an acknowledgment and a TCP receive window update. Each TCP packet contains the advertised TCP receive window on the connection. A larger window allows the sender to continue sending data and can improve communication performance.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza. You can set this option on the **Communication** tab, **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 the client to use the operating system default TCP window size. Values from 1 to 2048 indicate that the window size is in the range of 1KB to 2MB. For Mac OS X, the range of values is 1 through 128. If you specify a value less than 1, the TCP window size defaults to 1. If you specify a value greater than 2048, the TCP window size defaults to 2048.

For backup-archive clients, the default value for this parameter is 63 KB.

For Tivoli Storage Manager for Virtual Environments: Data Protection for VMware, the default value for this parameter is 512 KB.

Notes:

- 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.
- 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.
- Depending on the operating system communication settings, your system might not accept all values in the range of values.
- The tcpwindowsize option overrides the operating system's default TCP/IP session send and receive window sizes.

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 and enter system time.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time the client is called. Consult the documentation on your local system for details about setting up your locale definition.

Note: The `timeformat` option does not affect the web client. The web client uses the time format for the locale that the browser is running in. If the browser is not running in a locale that the client supports, the web client uses the time format for US English.

You can use the `timeformat` option with the following commands:

- **delete archive**
- **delete backup**
- **expire**
- **query archive**
- **query backup**
- **query filespace**
- **query image**
- **query nas**
- **restore**
- **restore image**
- **restore nas**
- **retrieve**
- **set event**

When you include the `timeformat` option with a command, it must precede the `fromtime`, `pittime`, and `totime` options.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client user-options file (`dsm.opt`). You can set this option on the **Regional Settings** tab, **Time Format** field of the Preferences editor.

Syntax

►►—`TIMEformat`— *—format_number—*—————►►

Parameters

format_number

Displays time in one of the formats listed here. Select the format number that corresponds to the format you want to use. When you include the `timeformat` option in a command, it must precede the `fromtime`, `pittime`, and `totime` options.

- 0** Use the locale-defined time format (does not apply to Mac OS X). This value is the default if the locale-specified format consists of digits, separator characters, and, if applicable, the AM or PM string.

- 1** 23:00:00

This is the default if the locale-specified format does not consist of digits, separator characters, and, if applicable, the AM or PM string.

- 2 23,00,00
- 3 23.00.00
- 4 12:00:00 A/P
- 5 A/P 12:00:00

Examples

Options file:

timeformat 4

Command line:

-time=3

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This is the value from the dsm.opt file unless overridden by the initial command line or by an option forced by the server.

Additional considerations for specifying time and date formats

The date or time format you specify with this option must be used when using options that take date and time as input. Examples are: totime, fromtime, todate, fromdate, and pittime.

For example, if you specify the timeformat option as TIMEFORMAT 4, the value that you provide on the fromtime or totime option must be specified as a time such as 12:24:00pm. Specifying 13:24:00 would not be valid because TIMEFORMAT 4 requires an hour integer that is 12 or less. If you want to specify up to 24 hour values on an option, and if you want to use commas as separators, you must specify TIMEFORMAT 2.

Configuring date and time formats in the system locale configuration file

You can specify date and time formats by configuring them in your system's locale file. If you specify time and date formats in the locale file, they must be defined using a subset of number-producing format specifiers that are supported by the C language strftime() function. You can use the following specifiers to set date and time formats in configuration settings for your locale.

Date specifiers

- %Y - the year, in four digits. For example, 2011.
- %y - the year, last two digits only. For example, 11 not 2011.
- %m - the month, as a decimal number (1-12).
- %d - the day of the month (1-31).

In the date specifiers, you can specify only one year specifier. Do not specify both %Y and %y. The E modifier (a capital E) can precede the date specifiers to produce the locale's alternative form for the year, month, or day. If no alternative form exists, the E modifier is ignored. Separate the specifiers with a single 7-bit ASCII character. Commonly used separators include colons (:), commas (,), periods (.), hyphens (-), or forward slash (/) characters. Do not use multibyte characters as separators.

Time specifiers

- %H - the hour, in 24-hour form (00-23).
- %I - the hour, in 12-hour form (00-12).
- %M - minutes after the hour (00-59).
- %S - seconds after the minute (00-59)
- %p - adds the AM (before noon) or PM (after noon) indicator.

In the time specifiers, you can specify only one hour specifier. Do not specify both %I and %H.

The O modifier (a capital O) can precede the time specifiers to produce the locale's alternative form for the hour, minutes, or seconds. The O modifier cannot precede the %p specifier. Separate the specifiers with a single 7-bit ASCII character. Commonly used separators include colons (:), commas (,), or periods (.). Do not use multibyte characters as separators. Do not specify a separator between the %p specifier and the separator that precedes or follows it.

Time format examples, configured in the locale settings

To set a particular time format, edit the configuration file for your locale and modify the `t_fmt` line to support your needs. Whatever time format you select applies both to output and to input. After the locale configuration file has been edited, the **localedef** command must be run to create the final locale file.

Table 76. Sample time format settings in the locale configuration (`t_fmt` line)

Example	Result
"%H:%M:%S"	Displays time in the form <i>hh:mm:ss</i> with <i>hh</i> ranging from 0 through 23.
"%H,%M,%S"	Displays time in the form <i>hh,mm,ss</i> with <i>hh</i> ranging from 0 through 23.
"%I,%M,13p"	Displays time in the form <i>hh,mm,ssA/P</i> with <i>hh</i> ranging from 1 through 12 and <i>A/P</i> is the local abbreviation for ante-meridian (AM in English) or post-meridian (PM in English).

Date format examples, configured in the locale settings

To set a particular date format, edit the configuration file and modify the `d_fmt` line as needed to support your needs. Whatever date format you select applies both to output and to input.

Table 77. Sample date format settings in the locale configuration (`d_fmt` line)

Example	Result
"%m/%d/%y"	Displays the date in the form MM/DD/YY.
"%d.%m.%Y"	Displays the date in the form DD.MM.YYYY.

Toc

Use the `toc` option with the **backup nas** command or the `include.fs.nas` option to specify whether the backup-archive client 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 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 only valid for AIX and Solaris clients. The Tivoli Storage Manager API does not support this option.

Options File

Place the `include.fs.nas` statement containing the `toc` value in the `dsm.sys` file within a server stanza.

Syntax



Parameters

Yes

Specifies that the client saves TOC information during a NAS file system image backup. However, the backup fails if an error occurs during creation of the TOC.

No Specifies that the client does not save TOC information during a NAS file system image backup.

Preferred

Specifies that the client 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, the client performs a full image backup and creates a TOC.

Examples

Options file:

```
include.fs.nas netappsj/vol/vol0 homemgmtclass toc=yes
```

Command line:

```
backup nas -nasnodename=netappsj /vol/vol0 -toc=yes
```

Todate

Use the `todate` option with the `totime` option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation.

Use the `todate` and `totime` options with the `fromtime` and `fromdate` options to request a list of backed up or archived files within a period of time. For example, you might request a list of files that were backed up between 6:00 AM on July 1, 2002 and 11:59 PM on July 30, 2002.

Use the `todate` option with the following commands:

- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—TODate =— —*date*—————►►

Parameters

date

Specifies an ending date. Enter the date in the format you selected with the `dateformat` option.

When you include `dateformat` with a command, it must precede the `fromdate`, `pitdate`, and `todate` options.

Examples

Command line:

```
dsmc restore "/Users/agordon/Documents/*" -todate=12/11/2003
```

Command line:

```
dsmc restore "/home/user1/*" -todate=12/11/2003
```

Totime

Use the `totime` option with the `todate` option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation. The backup-archive client ignores this option if you do not specify the `todate` option.

Use the `totime` and `todate` options with the `fromtime` and `fromdate` options to request a list of files that were backed up within a period of time. For example, you might request a list of files that were backed up between 6:00 AM on July 1, 2003 and 11:59 PM on July 30, 2003.

Use the `totime` option with the following commands:

- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►► `-TOTime =` `—time—` ◄◄

Parameters

time

Specifies an ending time. If you do not specify a time, the time defaults to 23:59:59. Specify the time in the format you selected with the `timeformat` option.

When you include the `timeformat` option in a command, it must precede the `fromtime`, `pittime`, and `totime` options.

Examples

Command line:

```
dsmc restore "/Users/van/Documents/myfiles/*" -todate=09/17/2003  
-totime=23:00:00
```

Command line:

```
dsmc restore "/home/user1/*" -todate=09/17/2003 -totime=23:00:00
```

Txnbytelimit

The `txnbytelimit` option specifies the number of kilobytes the client program buffers before it sends a transaction to the server.

A *transaction* is the unit of work exchanged between the client and server. A transaction can contain more than one file or directory, called a *transaction group*.

You can control the amount of data sent between the client and server, before the server commits the data and changes to the server database, using the `txnbytelimit` option. Controlling the amount of data sent changes the speed of the client to perform the transactions. The amount of data sent applies when files are batched together during backup or when receiving files from the server during a restore procedure.

After the `txngroupmax` number is reached, the client sends the files to the server, even if the transaction byte limit is not reached.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza. You can set this option on the **General** tab, in the **Transaction Buffer Size** field in the Preferences editor.

Syntax

►►—`TXNByteLimit`— *number* —————►◄

Parameters

number

Specifies the number of kilobytes the client program sends to the server before committing the transaction. The range of values is 300 through 34359738368 (32 GB). The default is 25600 KB. The number can be specified as an integer or as an integer with one of the following unit qualifiers:

- K or k (kilobytes)
- M or m (megabytes)
- G or g (gigabytes)

If no unit qualifier is specified, the integer is in kilobytes.

Restriction: The `txnbytelimit` option does not support decimal numbers, and only one-unit letters are allowed. For example: K, M, or G.

Examples

Options file:

```
txnb 25600
txnb 2097152
txnb 2097152k
txnb 2048m
txnb 2g
txnb 32G
```

Command line:

```
-txnb=25600
-txnb=16G
```

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

Type

Use the `type` option with the **query node** command to specify the type of node to query. Use this option with the **set event** command to activate, hold, or release.

Supported Clients

This option is also valid for the **set password** command with TSM on AIX clients.

This option is only valid for AIX and Solaris clients. The Tivoli Storage Manager API does not support this option.

Syntax



Parameters

nas

Specifies all NAS nodes registered at the server.

server

Specifies client nodes that are other Tivoli Storage Manager servers.

client

Specifies client nodes that are backup-archive clients.

Examples

Command line:

query node -type=nas

Updatectime

| Use the updatectime option to check the change time (ctime) attribute during an
| incremental backup operation.

Use this option with the **incremental**, **selective**, or **archive** commands.

Supported Clients

This option is valid for AIX and Linux clients on GPFS file systems only. The server can also define this option.

Options File

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

Syntax



Parameters

no The backup-archive client does not check the change time (ctime attribute) during a backup operation. This value is the default.

yes

The backup-archive client checks the change time (ctime attribute) during a backup operation. If the ctime attribute changed since the last backup operation, the ctime attribute is updated on the Tivoli Storage Manager server. The object is not backed up unless it has either ACLs or extended attributes. The client checks files and directories.

Examples

Options file:

updatect yes

Command line:

dsmc incr /proj/gpfs/test/ -updatectime=yes

Useexistingbase

The useexistingbase option is used when you back up snapshots that are on NetApp filer volumes. The useexistingbase option indicates that the latest snapshot that exists on the volume being backed up, is to be used as the base snapshot, during a snapshot differential backup operation.

If this option is not specified, a new snapshot is created on the volume that is being backed up. Because target filer volumes are read only volumes, useexistingbase must be specified when performing snapshot differential backups of target filer volumes. If useexistingbase is not specified, snapshot differential backups of a target filer volume fail because the new snapshot cannot be created on the read only volume.

When backing up target filer volumes, use both the useexistingbase option and the diffsnapshot=latest option to ensure that the most recent base and most recent differential snapshots are used during the volume backup

Supported Clients

This option can be used with supported x86_64 Linux and AIX clients.

Options File

This option is only valid on the command line.

Syntax

►►—USEEXISTINGBase—————◄◄

Parameters

This option has no parameters

Examples

Options file:

Does not apply.

Command line:

```
dsmc incr \\DRFiler\UserDataVol_Mirror_Share -snapdiff
-useexistingbase -basenameshotname="nightly.?"
```

Related information

Basesnapshotname

Usereplicationfailover

The usereplicationfailover option specifies whether automated client failover occurs on a client node.


Use this option to enable a client node for failover or to prevent it from failing over to the secondary server. This option overrides the configuration that is provided by the Tivoli Storage Manager server administrator settings on the primary server.

Supported Clients

This option is valid for all clients.

Options File

Place this option within a server stanza in the dsm.sys file.

Syntax**Parameters***Yes*

Specifies that you want the client to automatically fail over to the secondary server if the primary server is unavailable. The client uses the configuration that is provided by the primary server to connect to the secondary server. This value is the default.

No Specifies that the client does not automatically fail over to the secondary server.

Examples**Options file:**

```
USEREPLICATIONFailover no
```

Command line:

Does not apply.

Related concepts:

“Automated client failover configuration and use” on page 87

Related tasks:

“Configuring the client for automated failover” on page 90

Users (deprecated)

This option is deprecated.

V2archive

Use the v2archive option with the **archive** command to archive only files to the server.

The backup-archive client will not process directories that exist in the path of the source file specification.

This option differs from the filesonly option in that the filesonly option archives the directories that exist in the path of the source file specification.

The v2archive and dirsonly options are mutually exclusive and an error message is displayed if you use both options in the same **archive** command.

If you use this option, you might want to consider the following:

- You might experience performance problems when retrieving large amounts of data archived with this option.
- You might want to use this option only if you are concerned about expiration performance on a server that already contains extremely large amounts of archived data.
- If there are multiple files with the same name for the v2archive option, the files are archived multiple times, with their directory structure. The v2archive option archives only the files.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—V2archive—◄◄

Parameters

There are no parameters for this option.

Examples

This command:

```
dsmc archive "/Users/user2/Documents/*" -v2archive -su=y.
```

Archives these files:

```
/Users/user2/Documents/file1  
/Users/user2/Documents/file2  
/Users/user2/Documents/file3  
/Users/user2/Documents/dir2/file4  
/Users/user2/Documents/dir2/file5
```

Note: The client does not archive /Users/user2/Documents and /Users/user2/Documents/dir2.

This command:

```
dsmc archive "/home/relx/dir1/*" -v2archive -su=y.
```

Archives these files:

```
/home/relx/dir1/file1  
/home/relx/dir1/file2  
/home/relx/dir1/file3  
/home/relx/dir1/dir2/file4  
/home/relx/dir1/dir2/file5
```

Note: The client does not archive `/home/relx/dir1` and `/home/relx/dir1/dir2`.

Vappmc

Use the `vappmc` option to store VMware vCloud vApp backups by using a management class other than the default management class. The default management class is defined on the server, for the domain that the node belongs to.

If you set this option in the options file for a data mover node, all vApp back ups that are performed by the data mover are stored as indicated by the specified management class.

Supported Clients

This option can be used with supported Linux clients that are configured to back up VMware vCloud vApps.

Options File

Place this option in the client options file, either `dsm.opt` or `dsm.sys`, or on the command line. The server can also define this option.

Syntax

►►—VAPPMc—*mgmtclassname*—————►►

Parameters

mgmtclassname

Specifies a management class that applies to VMware vApp backups. If you do not set this option, the default management class of the node is used.

Tip: The management class that you specify on this option affects the storage of vApp data; the `vmmc` option does not apply to vApp data storage processing. However, the `vmctlmc` option does specify the management class that is used to store control files for both virtual machine backups and vApp backups.

Examples

Options file:

Specify a management class to use when the server stores all vApp backups:

```
vappmc MCPRDVPAPPs
```

Command line:

```
dsmc backup vapp "org1,vdc1,vapp1" -vappmc=myManagementClass
```

Vcdhost

The `vcdhost` option specifies the host name of the VMware vCloud Director server that manages vApps that you want to protect.

Supported Clients

This option is valid for Linux clients that are configured to perform vCloud Director vApp backups.

Options file

Set this option in the client options file (`dsm.opt`), or on the command line. This option is required, either in the client options file or on the command line, to use the **backup vapp** or **restore vapp** commands.

Syntax

►►—VCDHost— —vCloud_server—►►

Parameters

vCloud_server

Specifies the VMware vCloud Director server. You can specify either a host name or an IP address.

Examples

The following examples show how you can use this option.

Options file:

`VCDHost vcloud1.example.com`

Command line:

`-vcdhost=vcloud2.example.com`

Vcdpw

The `vcdpw` option specifies the password for the VMware vCloud system administrator account that you specify on the `vcduser` option.

The `vcdpw` option is required when you use the **Backup VAPP** or **Restore VAPP** commands.

Supported Clients

This option is valid for Linux clients that are configured to perform vCloud Director vApp backups.

Options file

You can set this option in the client options file (`dsm.opt`). However, storing credentials in clear text is not a best practice. Alternatively, set the `vcdhost` option in the client options file. Then, include that host name and the vCloud administrator name, and the vCloud administrator password, on a **set password** command and specify the type parameter as `TYPE=VCD`. All three of these credentials are required to use the **backup vapp** or **restore vapp** commands.

Syntax

►►—VCDPw— —password—◄◄

Parameters

vCloud_pw

Specifies the password for the VDCUser (vCloud system administrator).

Examples

The following examples show how you can use this option.

Options file:

Do not set this option in the client options file.

Command line:

-VCDPw=password

Vcduser

The vcduser option specifies a VMware vCloud system administrator account name.

Supported Clients

This option is valid for Linux clients that are configured to perform vCloud Director vApp backups.

Options file

You can set this option in the client options file (dsm.opt). However, storing credentials in clear text is not a best practice. Alternatively, set the vcdhost option in the client options file. Then, include that host name and the vCloud administrator name, and the vCloud administrator password, on a **set password** command and specify the type parameter as TYPE=VCD. All three of these credentials are required to use the **backup vapp** or **restore vapp** commands.

Syntax

►►—VCDUser— —vCloud_administrator—◄◄

Parameters

vCloud_administrator

Specifies the name of a vCloud system administrator account.

Examples

The following examples show how you can use this option.

Options file:

VCDUser administrator

Command line:

-vcduser=administrator

Verbose

The verbose option specifies that you want to display detailed processing information on your screen. This is the default.

When you run the **incremental**, **selective**, or **archive** commands, information is displayed about each file that is backed up. Use the quiet option if you do not want to display this information.

The following behavior applies when using the verbose and quiet options:

- If the server specifies either the quiet or verbose option in the server client option set, the server settings override the client values, even if **force** is set to *no* on the server.
- If you specify quiet in your dsm.opt file, and you specify -verbose on the command line, -verbose prevails.
- If you specify both -quiet and -verbose on the same command, the last option encountered during options processing prevails. If you specify -quiet -verbose, -verbose prevails. If you specify -verbose -quiet, -quiet prevails.

The information is displayed on your screen in the Scheduler Status window. This option only applies when you are running the scheduler and the client is performing scheduled work.

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Command Line** tab, **Do not display process information on screen** checkbox of the Preferences editor.

Syntax

►►—Verbose—————►►

Parameters

There are no parameters for this option.

Examples

Options file:

verbose

Command line:

-verbose

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

Verifyimage

Use the `verifyimage` option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume.

If bad sectors are detected on the target volume, the backup-archive client issues a warning message on the console and in the error log.

Supported Clients

This option is valid for AIX, HP-UX, all Linux clients, and Solaris *only*. The Tivoli Storage Manager API does not support this option.

Syntax

►►—VERIFYImage—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:
`dsmc restore image /usr -verifyimage`

Virtualfsname

Use the `virtualfsname` option with the **backup group** command to specify the name of the virtual file space for the group on which you want to perform the operation. The `virtualfsname` cannot be the same as an existing file space name.

Supported Clients

This option is valid for all UNIX and Linux clients except for Mac OS X.

Syntax

►►—VIRTUALFname =— —*fsname*—◄◄

Parameters

fsname

Specifies the name of the container for the group on which you want to perform the operation.

Examples

Command line:
`backup group -filelist=/Users/van/Documents/filelist1 -groupname=group1
-virtualfsname=virtfs -mode=full
backup group -filelist=/home/dir1/filelist1 -groupname=group1
-virtualfsname=virtfs -mode=full`

Virtualmountpoint

The `virtualmountpoint` option defines a virtual mount point for a file system if you want to consider files for backup that begin with a specific directory within that file system.

Using the `virtualmountpoint` option to identify a directory within a file system provides a direct path to the files you want to back up, saving processing time. It is more efficient to define a virtual mount point within a file system than it is to define that file system using the `domain` option, and then to use the `exclude` option in your include-exclude options list to exclude the files that you do not want to back up.

Use the `virtualmountpoint` option to define virtual mount points for multiple file systems, for local and remote file systems, and to define more than one virtual mount point within the same file system. Virtual mount points cannot be used in a file system handled by automounter.

You can use the `virtualmountpoint` option to back up unsupported file systems, with certain limitations. For information about using `virtualmountpoint` with unsupported file systems, see “File system and ACL support” on page 153.

Note: If the directory that you want to specify as a virtual mount point is a symbolic link, set the `followsymbolic` option to *Yes*. If that option is set to *no* (the default), you are not permitted to use a symbolic link as a virtual mount point. Also, if you back up a file system, then add a virtual mount point, and then do another incremental on the file system, the files and directories in the virtual mount point directory are expired, because they are logically contained within the virtual mount point directory and not the file system.

After you define a virtual mount point, you can specify the path and directory name with the `domain` option in either the default client options file or on the **incremental** command to include it for incremental backup services. When you perform a backup or archive using the `virtualmountpoint` option, the **query filespace** command lists the virtual mount point in its response along with other file systems. Generally, directories that you define as virtual mount points are treated as actual file systems and require that the `virtualmountpoint` option is specified in the `dsm.sys` file to restore or retrieve the data.

Note: When you specify a `virtualmountpoint` option, the path that it specifies is added to the default backup domain (`domain all-local`). The `virtualmountpoint` path is always considered a local “mount point” regardless of the real file system type it points to.

Supported Clients

This option is valid for all UNIX clients except Mac OS X. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza.

Syntax

Options File

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

Syntax

►—VIRTUALnodename— —*nodename*—►

Parameters

nodename

Specifies a 1- to 64-character name that identifies the node for which you want to request Tivoli Storage Manager services. There is no default.

Examples

Options file:

virtualnodename cougar

Command line:

-virtualn=banshee

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

Vmbackdir

The `vmbackdir` option specifies the temporary disk location where the client saves control files that are created during full VM backup and restore operations of VMware and Microsoft Hyper-V virtual machines.

Supported Clients

When a client on a data mover node starts a full VM backup of a virtual machine, the client creates metadata in files that are associated with the backed up virtual machine and its data. The files that contain the metadata are referred to as *control files*.

During full VM backup operations, the metadata is saved on a disk in the data mover node until the backup completes and both the virtual machine data and the control files are saved to server storage. During a full VM restore operation, the control files are copied from the server and are temporarily stored on the data mover disk, where they are used to restore the virtual machine and its data. After a backup or a restore operation completes, the control files are no longer needed and the client deletes them from their temporary disk location.

The directory that is specified by this option must be on a drive that contains sufficient free space to contain the control information from a full VM backup.

This option is valid for Linux clients that are installed on a vStorage backup server.

Options File

Set this option in the client options file, or specify it on the command line as an option for the **backup vm** or **restore vm** commands.

Syntax

►►—VMBACKDir—directory—►►

Parameters

directory

Specifies the path where the control files are stored on the backup server.

The default is /tmp/tsmvmbackup/fullvm/

Examples

Options file:

```
VMBACKD /tmp/tsmvmbackup/
```

Command line:

```
dsmc backup vm -VMBACKUPT=fullvm -VMBACKD=/home/vmware/control_files
```

```
dsmc restore vm -VMBACKUPT=fullvm -VMBACKD=/home/mine/bkup_ctrl
```

Vmbackupmailboxhistory

The vmbackupmailboxhistory option specifies whether mailbox history is automatically uploaded with the virtual machine (VM) backup if IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server is detected on a VM.

Supported Clients

This option is valid on clients that act as a data mover for VMware guest backups.

Options File

Place this option in the client system-options file (dsm.sys) within a server stanza.

Syntax

►►—VMBACKUPMAILBoxhistory—

Yes
No

—►►

Parameters

Yes

The mailbox history is automatically uploaded with the VM backup if Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server is detected on a VM.

No The mailbox history is not automatically uploaded with the VM backup.

Examples

Options file:

```
vmbackupmailboxhistory yes
```

Vmbackuptype

Use the `vmbackuptype` option with the **backup VM** or **restore VM** command to specify the type of virtual machine backup or restore to complete. You can also use this option on **query VM** commands to filter the query results to include only virtual machines that were backed up by a specific backup type. For examples, see the **query VM** command description.

You can specify a VMware full VM backup.

Supported Clients

This option is valid on Linux clients that are installed on a vStorage backup server. The server can also define this option.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza.

Syntax

►►—VMBACKUPTYPE—FULLvm—►►

Parameters

FULLvm

Specify this value to run a traditional full VM backup of a VMware virtual machine. This is the default backup type for Linux clients.

Examples

Options file:

```
VMBACKUPT full
```

Command line:

```
dsmc backup vm vm1 -VMBACKUPT=full -vmchost=virtctr  
-vmcuser=virtctr_admin -vmcpw=xxxxx
```

Performs a full virtual-machine backup of `vm1.example.com` using the VMware VirtualCenter machine `virtctr.example.com`, to the Tivoli Storage Manager server, using machine name `vm1`.

Vmchost

Use the `vmchost` option with the **backup VM**, **restore VM**, or **query VM** commands to specify the host name of the VMware VirtualCenter or ESX server that you want to backup, restore, or query.

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 command is valid for clients that are configured to perform an off-host backup of a VMware virtual machine. The server can also define this option.

Options File

Place this option in the client options file (dsm.opt), the client system options file (dsm.sys), or on the command line.

Syntax

▶▶—VMCHost— —*hostname*————▶▶

Parameters

hostname

Specifies the host name of the VMware VirtualCenter or ESX server that you want to backup, restore, or query.

Examples

Options file:

```
VMCH vcenter.storage.usca.example.com
```

Command line:

```
-VMCH=esx1.storage.usca.example.com
```

Vmcpw

Use the vmcpw option with the **backup VM**, **restore VM**, or **query VM** commands to specify the password for the VMware VirtualCenter or the ESX user ID that is specified with the vmcuser option.

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 only on supported Linux clients that are installed on a vStorage backup server that is used to backup a VMware virtual machine.

Options File

Place this option in the client system options file (dsm.sys), or on the command line.

1. Click **Edit > Client Preferences > VM Backup**. In the **Password** field, type the password that you want to have saved.
2. Click **OK**.

As an alternative to the preferences editor, you can store the password locally by using the **set password** command. For example:

```
dsmc SET PASSWORD -type=vm  
vcenter.us.ibm.com Administrator secret
```

Syntax

►—VMCPw— *pwname*—►

Parameters

pwname

Specifies the password for the VMware VirtualCenter or ESX server that you want to backup, restore, or query.

Examples

Options file:

VMCPw SECRET

Command line:

-VMCPw=SECRET

Related reference:

“Set Password” on page 734

Vmctlmc

This option specifies the management class to use when backing up virtual machine control files.

By default, virtual machine control files are bound to the default management class. The `vmmc` option can be used to specify a different management class to which virtual machine data and virtual machine control files are bound. The `vmctlmc` option overrides the default management class and the `vmmc` option for the virtual machine control files.

Under certain conditions, it might be desirable or necessary to bind the control files to a different management class than the data files.

The `vmctlmc` option is required if virtual machine data files are backed up to tape. Virtual machine control files must be backed up to a disk-based storage pool that does not migrate to tape. The storage pool can be composed of random access volumes and sequential file volumes; the storage pool can also be a deduplicated pool. Use the `vmctlmc` option to specify a management class that stores data in such a storage pool.

Restriction: The management class that is specified by the `vmctlmc` option determines only the destination storage pool for virtual machine control files. Retention of the control files is determined by the `vmmc` option, if specified, or by the default management class. The retention for the virtual machine control files always matches the retention of the virtual machine data files.

Supported Clients

This option is valid for clients that act as data mover nodes that protect VMware virtual machines.

The option can only be used for virtual machine backups that use an incremental-forever backup mode.

This option is available only if you have a license to use either Tivoli Storage Manager for Virtual Environments: Data Protection for VMware or Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V.

Options File

Place this option in the system options file `dsm.sys`.

Syntax

►►—VMCTLmc—*class_name*—————►►

Parameters

class_name

Specifies a management class that applies to backing up virtual machine control files. If you do not set this option, the management class that is specified on the `vmmc` option is used. If you do not set this option and the `vmmc` option is not set, the default management class of the node is used.

Examples

Options file:

`vmctlmc diskonlymc`

Command line:

Does not apply.

Vmcuser

Use the `vmcuser` option with the **backup VM**, **restore VM**, or **query VM** commands to specify the user name of the VMware VirtualCenter or ESX server that you want to backup, restore, or query.

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 clients that are configured as to perform an off-host backup of VMware virtual machines. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`), the client system options file (`dsm.sys`), or on the command line.

Syntax

►►—VMCUser— *username*—————►►

Parameters

username

Specifies the user name of the VMware VirtualCenter or ESX server that you want to backup, restore, or query.

When working with a virtual center, a user id with access to the Windows system hosting the virtual center is required. This user id must either have administrator privileges, or the minimum privileges that are identified in technote 1659544.

Examples

Options file:

```
VMCUser administrator
```

Command line:

```
backup vm -VMCUser=domainname\administrator
```

Command line:

Example of connecting to an ESX server:

```
backup vm -VMCUser=root
```

Vmdatastorethreshold

Use the `vmdatastorethreshold` option to set the threshold percentage of space usage for each VMware datastore of a virtual machine.

When you specify this option, space usage is checked before a virtual machine snapshot is created. If the threshold is exceeded, the virtual machine is not backed up. By setting this option, you can prevent out-of-space errors when you back up virtual machines.

To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

Supported clients

You can use this option with supported x86_64 Linux clients.

Options file

You can specify this option in the client system-options file (`dsm.sys`) or on the command line by using the **backup vm** command. You can also include this option on the Tivoli Storage Manager Version 7.1.5 or later server in a client option set. You cannot set this option in the Preferences Editor.

Syntax

►—VMDATASTOREThreshold—*percent*————►

Parameters

percent

Specifies the threshold percentage of each VMware datastore of the virtual machine to be backed up. You can specify an integer from 0 - 100. The default value is 100. If you do not set this option, the client begins a virtual machine backup without first verifying the existing space usage.

Requirements:

- Ensure that the threshold is low enough so that the snapshot does not use up all the available space in the VMware datastores. Otherwise, you will run out of space on the VMware datastores and the snapshot will not be created.
- If you use multiple clients that act as data mover nodes, you must add this option to the options file for each data mover.
- The client checks the data usage of the VMware datastore that contains the virtual machine disk snapshots. By default, the snapshots are created in the same directory as that of the parent virtual disk (.vmdk) file.

If you change the snapshot location to a new directory on the same datastore or on another datastore with the `workingDir` option in the VM configuration file, ensure that the path of the working directory is correct. If the path is incorrect, the client might validate the data usage of the wrong datastore.

If you use the `EXCLUDE.VMDISK` option to exclude one or more disks from a backup, the threshold check is still run on these disks. Even though these disks are not backed up, VMware still takes a snapshot of these disks.

Independent disks are not checked during space verification processing because a snapshot of these disks does not use any VMware datastore space.

Example 1

Virtual machine `vm1` spans `datastore1` and `datastore2`. Set the `vmdatastorethreshold` option to 90 to ensure that both VMware datastores are at most 90% full before the virtual machine is backed up.

Options file:

```
vmdatastorethreshold 90
```

Command line:

```
dsmc backup vm vm1 -vmdatastorethreshold=90
```

Example 2

The datastore threshold of `datastore2` is set to 85. The datastore threshold is exceeded during the backup of virtual machine `vm5`. The following error message is displayed:

```
ANS14200E The virtual machine 'vm5' could not be backed up because the
data usage of datastore 'datastore2' exceeded the datastore threshold
of 85%.
```

Increase the value of the `vmdatastorethreshold` option to 95 and restart the backup.

Options file:

```
vmdatastorethreshold 95
```

Command line:

```
dsmc backup vm vm5 -vmdatastorethreshold=95
```

Related reference:

“Backup VM” on page 634

Vmdefaultdvportgroup

Use this option to specify the port group for the NICs to use during **restore vm** operations for a virtual machine that was connected to a distributed virtual port

group when it was backed up, but the target host for the restore operation does not contain a similar distributed virtual port group.

This option does not apply to backup or restore operations for Microsoft Hyper-V virtual machines.

Supported clients

This option is valid for Linux clients that are installed on a vStorage backup server.

Options file

Place this option in the client options file (`dsm.opt`), in the client system options file (`dsm.sys`), or specify it as a command-line parameter on the **restore vm** command.

Syntax

►—VMDEFAULTDVPORTGROUP—*portgroup_name*—◄

Parameters

portgroup_name

Specifies the name of the port group to use. The port group name is case sensitive.

Examples

Option file:

```
VMDEFAULTDVPORTGROUP dvPortGroup
```

Command line:

```
dsmc restore vm vm123 -VMDEFAULTDVPORTGROUP=dvPortGroup
```

Related reference:

“Vmdefaultnetwork” on page 577

“Vmdefaultdvswitch”

Vmdefaultdvswitch

Use this option to specify the distributed virtual switch (dvSwitch) that contains the port group that you set on the `vmdefaultdvportgroup` option. The option has no effect unless you also specify the `vmdefaultdvportgroup` option.

This option does not apply to backup or restore operations for Microsoft Hyper-V virtual machines.

Supported clients

This option is valid for Linux clients that are installed on a vStorage backup server.

Options file

Place this option in the client options file (`dsm.opt`), in the client system options file (`dsm.sys`), or specify it as a command-line parameter on the **restore vm** command.

Syntax

►►—VMDEFAULTDVSWITCH—*dvSwitch*—◄◄

Parameters

dvSwitch

Specifies the name of the virtual switch to use. The virtual switch name is case sensitive.

Examples

Option file:

```
VMDEFAULTDVSWITCH dvSwitch
```

Command line:

```
dsmc restore vm vm123 -VMDEFAULTDVSWITCH=dvSwitch -VMDEFAULTVPORTGROUP=dvPortGroup
```

Related reference:

“Vmdefaultdvportgroup” on page 575

Vmdefaultnetwork

Use this option to specify the network for NICs to use during a **restore vm** operation, for a virtual machine that had been connected to a distributed virtual port group when it was backed up, but the target host for the restore operation does not have any distributed switch port groups configured.

This option does not apply to restore operations for Microsoft Hyper-V virtual machines.

Supported clients

This option is valid for Linux clients that are installed on a vStorage backup server.

Options file

Place this option in the client options file (`dsm.opt`), in the client system options file (`dsm.sys`), or specify it as a command-line parameter on the **restore vm** command.

Syntax

►►—VMDEFAULTNETWORK—*vm_network_name*—◄◄

Parameters

vm_network_name

Specifies the name of the virtual machine network to use. The network name is case sensitive. If the name contains space characters, enclose it in quotation marks.

Examples

Option file:

```
VMDEFAULTNETWORK "VM Network"
```

Command line:

```
dsmc restore vm vm123 -VMDEFAULTNETWORK="VM Network"
```

Related reference:

“Vmdefaultdvportgroup” on page 575

“Vmdefaultdvswitch” on page 576

Vmenabletemplatebackups

The `vmenabletemplatebackups` option specifies whether the client backs up VMware template virtual machines when it protects virtual machines in a vCenter server. VMware templates virtual machines cannot be backed up when they are in an ESXi host because ESXi does not support templates.

When this option is enabled, you can include VMware template machines in full VM backup operations. You use the existing **Backup VM** command and the `DOMAIN.VMFULL` option to specify the virtual machines to include in the backup operation.

Incremental backups are not supported and snapshots are not taken, so you must use `MODE=FULL` or `MODE=IFFULL`. Use `MODE=IFFULL` to force a new backup of VMware template virtual machines, even if they were not changed since the last backup.

When `vmenabletemplatebackups` is enabled, any backup process that is initiated by using `MODE=INCREMENTAL` is performed by using `MODE=FULL`. Likewise, any backup process that is initiated by using `MODE=IFINCREMENTAL` is processed by using `MODE=IFFULL`. VMware template VMs are included in a backup only if they were changed since the last backup occurred.

With this option enabled, make sure that the `vmvstortransport` options include `NBDSSL` or `NBD`. Using only the `SAN` or `HOTADD` transport modes with this option enabled causes backups of the template machines to fail.

Supported clients


This option can be used with supported x86_64 Linux clients.

Options file

You can set this option on the command line, in the client system options file (`dsm.sys`), client options file (`dsm.opt`), or on the server in a client options set.

You can also set it in the preferences editor on the VM Backup tab (select the **Backup virtual machine templates** option).

Syntax



Parameters

No Specifies that template virtual machines are not included in full VM backup operations; this is the default setting.

Yes

Specifies that template VMs are included in full VM backup operations.

Examples**Options file**

```
vmenabletemplatebackups yes
```

Command line

Back up a VMware template VM

```
dsmc backup vm vmname -VMENABLETEMPLATEBACKUPS=YES
```

where *vmname* is the template machine name.

Command line

Restore a VMware template VM to the same location and name

```
dsmc restore vm vmname -VMENABLETEMPLATEBACKUPS=YES
```

where *vmname* is the template machine name.

Command line

Restore a template virtual machine to a new location

```
dsmc restore vm vmname -vmname=win7x64  
-datastore=datastore22 -host=supersht.labx.com  
-datacenter="Lab Center" -VMENABLETEMPLATEBACKUPS=YES
```

where *vmname* is the template machine name. “win7x64” is the new template VM name. The new data center, host, and datastore are also included.

Related reference:

“Backup VM” on page 634

“Restore VM” on page 715

“Domain.vmfull” on page 369

Vmlimitperdatastore

The `vmlimitperdatastore` specifies the maximum number of virtual machines in a datastore that can be included in a parallel backup operation.

This option works together with the `vmmaxparallel` and `vmlimitperhost` options to optimize the performance of parallel backups, and to reduce the processor load that parallel backs can create on a host in the vSphere infrastructure.


Supported clients

This option can be used with supported x86_64 Linux clients.

Options file

This option is valid in the client system options file (`dsm.sys`), in the client options file (`dsm.opt`) or on the command line for **Backup VM**. It can also be included on the server in a client options set. It cannot be set in the Preferences Editor.

Syntax



Parameters

integer

Specifies the maximum number of virtual machines in any one data store, that are included during a parallel backup operation. The maximum that you can specify is 50 VMs. The default is 0 (zero).

Specifying 0 means that you are not concerned about how many VMs can be backed up in parallel from a datastore. Instead, you want to limit the maximum number of VMs to include in a parallel backup by using the value that you specify on the `vmmaxparallel` option. The `vm-limit-per-datastore` option is enforced even when virtual machine data exists in two or more datastores. Adjust the values of `vmmaxparallel`, `vm-limit-per-host`, and `vm-limit-per-datastore` to find the values that provide optimal performance for parallel backups, in your environment.

Examples

Options file

```
VMLIMITPERD 5
```

Related reference:

“**Backup VM**” on page 634

“`Domain.vmfull`” on page 369

“`Vmmaxparallel`” on page 581

“`Vmlimitperhost`”

Vmlimitperhost

The `vm-limit-per-host` specifies the maximum number of virtual machines in an ESX server that can be included in a parallel backup operation.

This option works together with the `vmmaxparallel` and `vm-limit-per-datastore` options to optimize the performance of parallel backups, and to reduce the processor load that parallel backs can create on a host in the vSphere infrastructure.

Supported clients

This option can be used with supported x86_64 Linux clients.

Options file

This option is valid in the client system options file (`dsm.sys`), in the client options file (`dsm.opt`) or on the command line for **Backup VM**. It can also be included on the server in a client options set. It cannot be set in the Preferences Editor.

Syntax



Parameters

integer

Specifies the maximum number of virtual machines in any one ESX server, that can be included in a parallel backup operation. The maximum that you can specify is 50 VMs. The default is 0 (zero).

Specifying 0 means that you are not concerned about how many VMs can be backed up in parallel from an ESX server. Instead, you want to limit the maximum number of VMs to include in a parallel backup by using the limit you specify on the `vmmaxparallel` option.

Adjust the values of `vmmaxparallel`, `vm-limitperhost`, and `vm-limitperdatastore` to find the values that provide optimal performance for parallel backups, in your environment.

Examples

Options file

```
VMLIMITPERH 5
```

Related reference:

“**Backup VM**” on page 634

“`Domain.vmfull`” on page 369

“`Vmmaxparallel`”

“`Vm-limitperhost`” on page 580

Vmmaxparallel

This option is used to configure parallel backups of several virtual machines by using a single instance of the backup-archive client. The `vmmaxparallel` option specifies the maximum number of virtual machines that can be backed up to the server, at any one time.

This option works together with the `vm-limitperhost` and `vm-limitperdatastore` options to optimize the performance of parallel backups, and to reduce the processor load that parallel backs can create on a host in the vSphere infrastructure.


Supported clients

This option can be used with supported x86_64 Linux clients.

Options file

This option is valid in the client system options file (`dsm.sys`) or on the command line for the **Backup VM** command. It can also be included on the server in a client options set. It cannot be set in the Preferences Editor.

Syntax



Parameters

integer

Specifies the maximum number of virtual machines that can be backed up, at any one time, during a parallel backup operation. The default is 1. The maximum is 50.

Adjust the values of `vmmaxparallel`, `vmllimitperhost`, and `vmllimitperdatastore` to find the values that provide optimal performance for parallel backups, in your environment.

Note: When you use client-side data deduplication, a data deduplication session is started for each VM. This data deduplication session is not counted as one of the `vmmaxparallel` sessions.

The `MAXNUMMP` server parameter specifies the maximum number of mount points a node is allowed to use on the server when the copy destination of the storage pool is FILE or TAPE. The `MAXNUMMP` parameter must be equal to or greater than the `VMMAXPARALLEL` setting. When multiple instances of the client are backing up files, or when a single client performs parallel backups, more mount points might be needed.

If the number of mount points that are requested exceeds the `MAXNUMMP` value, the server issues an error (ANS0266I). In response to the error, the client reduces `VMMAXPARALLEL` to match the number that is specified by `MAXNUMMP` and continues the backup with the reduced number of sessions. If additional ANS0266I errors are detected, the client reduces `VMMAXPARALLEL` by 1 and attempts to continue the backup. If `VMMAXPARALLEL` is decremented to 1 and the client receives more ANS0266I errors, the client ends the backup and issues the following error:

ANS5228E A backup VM operation failed because `VMMAXPARALLEL` was reduced to 1 and the client still cannot obtain a server mount point. Contact your server administrator if you need the value that is currently set for `MAXNUMMP` increased, so your node can support additional parallel backup sessions.

Examples

Options file

```
VMMAXP 10
```

Related reference:

“Backup VM” on page 634

“Domain.vmfull” on page 369

“Vmllimitperhost” on page 580

“Vmllimitperdatastore” on page 579

Vmmaxvirtualdisks

The `vmmaxvirtualdisks` option specifies the maximum size of the VMware virtual machine disks (VMDKs) to include in a backup operation.

Use the `vmmaxvirtualdisks` option with the `mskipmaxvirtualdisks` option to specify how the client processes large VMDKs during a backup operation:

- Set the `vmmaxvirtualdisks` option to specify the maximum size of the VMDKs to include.
- Set the `mskipmaxvirtualdisks` option to back up the VMDKs that do not exceed the maximum size (and exclude any VMDKs that exceed the size), or fail the operation.


Supported clients

This option is valid for 64-bit Linux clients that are configured as data movers that back up VMware virtual machines.

Options file

Set the `vmmaxvirtualdisks` option in the client system options file (`dsm.sys`). You can also specify this option as a command-line parameter on the **backup vm** command.

Syntax



Parameters

size

Specifies the maximum size, in terabytes, of the VMDKs to include in a backup operation. The range is an integer 2 - 8; the default is 2. The maximum is 8.

To ensure that the VMware VMDK size that is included in backup operations is always the maximum size, specify 999. Use this value as the most effective method to ensure that the maximum value is always set. This value prevents the need to continuously modify the option files.

When you also specify the `mskipmaxvirtualdisks yes` option, VMDKs that are the specified maximum size or smaller are backed up and VMDKs that are larger than the specified maximum size are excluded.

When you also specify the `mskipmaxvirtualdisks no` option, backup operations fail if a VMDK is larger than the specified maximum size.

Examples

Options file:

```
vmmaxvirtualdisks 3
```

Command line:

Back up VMDKs that are 5 TB or smaller and exclude VMDKs that are larger than 5 TB:

```
backup vm VM1 -vmmaxvirtualdisks=5 -mskipmaxvirtualdisks=yes
```

Back up VMDKs that are 3 TB or smaller and fail the backup operation if a VMDK is larger than 3 TB:

```
backup vm VM1 -vmmaxvirtualdisks=3 -mskipmaxvirtualdisks=no
```

Back up VMDKs that are 8 TB or smaller and exclude VMDKs that are larger than 8 TB:

```
backup vm VM1 -vmmaxvirtualdisks=8 -vmskipmaxvirtualdisks=yes
```

or

```
backup vm VM1 -vmmaxvirtualdisks=999 -vmskipmaxvirtualdisks=yes
```

Vmmc

Use the `vmmc` option to store virtual machine backups by using a management class other than the default management class. The `vmmc` option is only valid if the `vmbackuptype=fullvm` or `vmbackuptype=hypervfull` option is set.

Supported Clients

This option is valid for clients that are configured to back up VMware virtual machines. The server can also define this option.

Options File

Place this option in the client options file `dsm.opt`, in the client system options file `dsm.sys`, or on the command line.

Syntax

►►—VMMC—*management_class_name*—————►►

Parameters

management_class_name

Specifies a management class that applies to the backed up virtual machine data. If you do not set this option, the default management class of the node is used.

Examples

Task: Run a backup of the virtual machine that is named `myVirtualMachine` and save the backup according to the management class that is named `myManagmentClass`.

```
dsmc backup vm "myVirtualMachine" -vmmc=myManagmentClass
```

Vmnoprddisks

This option enables the client to restore configuration information for the pRDM volumes that are associated with a VMware virtual machine, even if the LUNs that were associated with the volumes cannot be found. Because pRDM volumes are not included in virtual machine snapshot, only the configuration information can be restored, and not the data that was on the volumes.

This option does not apply to backups of Microsoft Hyper-V virtual machines.

Supported Clients

This option is valid for Windows and Linux clients that are installed on a vStorage backup server.

Options File

Place this option in the client options file (dsm.opt), in the client system options file (dsm.sys), or specify it as a command-line parameter on the **restore vm** command.

Syntax



Parameters

YES

Specify this value if you must restore a virtual machine that you backed up with `-vmprocesswithprdm=yes`, and the original LUNs that were mapped by the raw device mappings file cannot be located. This setting causes the client to skip attempts to locate the missing LUNs used by the pRDM volumes, and restore the configuration information (disk labels) that were associated with them. The pRDM volumes are restored as thin-provisioned VMFS VMDKs. You can then use the vSphere client to create the necessary pRDM mappings.

NO Setting `-vmnoprdmdisk=no` causes restore operations for virtual machines that were backed up with `-processvmwithprdm=yes` to fail if the original LUNs that were mapped to by the raw device mappings file cannot be located. This value is the default value.

Examples

Option file:

```
VMNOPRDMDISKS YES
```

Command line:

```
dsmc restore vm vm123 -vmnoprdmdisks=yes
```

Related information

[“Vmprocessvmwithindependent” on page 587](#)

[“Vmprocessvmwithprdm” on page 588](#)

Vmnovrdmdisks

This option enables the client to restore configuration information and data for vRDM volumes that are associated with a VMware virtual machine, even if the LUNs that were associated with the volumes cannot be found.

This option does not apply to backups of Microsoft Hyper-V virtual machines.

Supported Clients

This option is valid for Windows and Linux clients that are installed on a vStorage backup server.

Options File

Place this option in the client options file (dsm.opt), in the client system options file (dsm.sys), or specify it as a command-line parameter on the **restore vm** command.

Syntax



Parameters

YES

Specify this value if you must restore a virtual machine that you backed up, and the original LUNs that were mapped by the raw device mappings file cannot be located. This setting causes the client to skip attempts to locate the missing LUNs used by the vRDM volumes, and restore the configuration information (disk labels) and the data that was backed up. The vRDM volumes are restored as thin-provisioned VMFS VMDKs.

NO Setting `-vmnovrdmdisk=no` causes restore operations for virtual machines that had vRDM volume to fail, if the original LUNs that were mapped to by the raw device mappings file cannot be located. This value is the default value.

Examples

Option file:

```
VMNOVRDMDISKS YES
```

Command line:

```
dsmc restore vm vm123 -vmnovrdmdisks=yes
```

Related information

“Vmprocessvmwithprdm” on page 588

Vmpreferdagpassive

The `vmpreferdagpassive` option specifies whether to back up an active copy or passive copy of a database that is part of a Microsoft Exchange Server Database Availability Group (DAG).

This option applies to Microsoft Exchange Server workloads that run inside VMware virtual machine guests that are protected by Tivoli Storage Manager for Virtual Environments.

Use the `vmpreferdagpassive` option with the **backup vm** command.

Supported Clients

This option is valid on clients that act as a data mover for VMware guest backups.

Options File

Place this option in the client system-options file (`dsm.sys`) within a server stanza.

Syntax



Parameters

No Back up the Microsoft Exchange Server database in a DAG regardless of whether it is an active copy or passive copy. This value is the default.

Yes

Skip the backup for an active database copy in a DAG if a valid passive copy is available on another server. If no valid passive copy is available, the active database copy is backed up.

Examples

Options file:

`vmpreferdagpassive yes`

Vmprocessvmwithindependent

Use this option to control whether full VMware virtual machine backups are processed if the machine is provisioned with one or more independent disk volumes.

Independent disk volumes do not support snapshots. Any independent disk volumes found on a virtual machine are not be processed as part of the backup operation. When the virtual machine is restored, the backup-archive client recovers the virtual machine, and only the volumes that participated in snapshot operations are restored. Configuration information and content of the independent disk volumes is not preserved in the information that is stored on the Tivoli Storage Manager server. Users must recreate the independent disk volumes on the restored machine.

If the virtual machine also contains one or more raw device mapping (RDM) volumes configured in physical-compatibility mode (pRDM), use the `vmprocessvmwithprdm` option to control whether the client backs up any files on the virtual machine if an independent disk is present.

This option is only valid for VMware backups and does not pertain to Microsoft Hyper-V backups.


Supported Clients

This option is valid for Windows and Linux clients that are configured as a VMware backup server. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`), in the client system options file (`dsm.sys`), or on the command-line.

Syntax



Parameters

No The backup of the virtual machine fails if one or more independent disk volumes are detected. No is the default.

Yes

Virtual machines that contain one or more independent disk volumes are backed-up. However, the independent disk volumes are not be processed as part of the virtual machine backup operation.

If the virtual machine also contains one or more raw device mapping (RDM) disks that are provisioned in physical-compatibility mode, the `VMPROCESSVMWITHPRDM` option must also be specified.

Examples

Option file:

```
VMPROCESSVMWITHINDEPENDENT Yes
```

Command line:

```
dsmc backup vm vmlocal -vmbackuptype=fullvm -vmprocessvmwithindependent=yes
```

Related information

`"Vmprocessvmwithprdm"`

Vmprocessvmwithprdm

Use this option to control whether full VMware virtual machine backups are processed if the virtual machine has one or more raw device mapping (RDM) volumes provisioned in physical-compatibility mode (pRDM).

pRDM volumes do not support snapshots. Any pRDM volumes found on a virtual machine are not processed as part of the backup operation. When the virtual machine is restored, the backup-archive client recovers the virtual machine, and only the volumes that participated in snapshot operations are restored. Configuration information and content of the pRDM volumes is not preserved in the information stored on the Tivoli Storage Manager server. Users must re-create the pRDM volumes on the restored machine.

This option does not apply to virtual machines that have one or more RDM volumes that are provisioned in virtual-compatibility mode (vRDM). Because vRDM volumes do support snapshot operations, they are included in a full VMware virtual machine backup.

If the virtual machine also contains one or more independent disks, use the `vmprocessvmwithindependent` option to control whether the client backs up any files on the virtual machine if an independent disk is present.

This option is only valid for VMware backups and does not pertain to Microsoft Hyper-V backups.

Supported Clients

This option is valid for Windows and Linux clients that are configured as a VMware backup server. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`), in the client system options file (`dsm.sys`), or on the command line.

Syntax



Parameters

No The backup of the virtual machine fails if one or more pRDM volumes are detected. No is the default.

Yes

Virtual machines that contain one or more raw device mapping (RDM) volumes that are provisioned in physical-compatibility mode (pRDM) are backed up. However, the pRDM volumes are not processed as part of the virtual machine backup operation.

If the virtual machine also contains one or more independent disks, the `vmprocessvmwithindependentdisk` option must also be specified.

Examples

Option file:

```
VMPROCESSVMWITHPRDM Yes
```

Command line:

```
dsmc backup vm vmlocal -vmbackuptype=fullvm -vmprocessvmwithprdm=yes
```

Related information

“Vmprocessvmwithindependent” on page 587

Vmskipctlcompression

Use the `vmskipctlcompression` option for VM backups to specify whether control files (*.ctl) are compressed during VM backup. The option does not affect the compression of data files (*.dat)

You can compress virtual machine control files and data files only when the files are stored in a storage pool that is enabled for client-side deduplication. Use the following options configuration to compress data files and not compress control files:

```
compression yes  
vmskipctlcompression yes
```

You must direct the data files to a storage pool that is enabled for client-side deduplication. You can direct the control files to a storage pool that is not enabled for client-side deduplication

You must be licensed to use IBM Tivoli Storage Manager for Virtual Environments to use this option.

Supported Clients

Options file

Place this option in the client options file (`dsm.opt`), or on the command line.

Syntax



Parameters

Yes

Do not compress control files (*.ctl) during VM backup. The option does not affect compression of data files (*.dat).

No Control files (*.ctl) can be compressed during VM backup. Whether control files are compressed depends on the value of the compression option.

Vmskipmaxvirtualdisks

The `vmskipmaxvirtualdisks` option specifies how the backup operation processes VMware virtual machine disks (VMDKs) that exceed the maximum disk size.

Use the `vmskipmaxvirtualdisks` option with the `vmmaxvirtualdisks` option to specify how the client processes large VMDKs during a backup operation:

- Set the `vmskipmaxvirtualdisks` option to back up the VMDKs that do not exceed the maximum size (and exclude any VMDKs that exceed the size), or fail the operation.
- Set the `vmmaxvirtualdisks` option to specify the maximum size of the VMDKs to include.

In V7.1.3 and earlier, the `vmskipmaxvirtualdisks` option was named `vmskipmaxvmdks`. In V7.1.4 and later, `vmskipmaxvirtualdisks` is the preferred option name. However, the client still processes backup operations with the `vmskipmaxvmdks` name.


Supported clients

This option is valid for 64-bit Linux clients that are configured as data movers that back up VMware virtual machines.

Options file

Set the `vmskipmaxvirtualdisks` option in the client system options file (`dsm.sys`). You can also specify this option as a command-line parameter on the **backup vm** command.

Syntax



Parameters

No Specifies that backup operations fail if a VMware virtual machine has one or more VMDKs that are larger than the maximum size. This setting is the default value.

Yes

Specifies that backup operations include VMware VMDKs that are the maximum size (or smaller) and exclude any VMDKs that are larger than the maximum size.

Examples

Options file:

```
vmskipmaxvirtualdisks yes
```

Command line:

Fail a backup operation if a VMDK is larger than 2 TB:

```
backup vm VM1 -vmskipmaxvirtualdisks=no
```

Fail a backup operation if a VMDK is larger than 5 TB:

```
backup vm VM1 -vmskipmaxvirtualdisks=no -vmmaxvirtualdisks=5
```

Back up VMDKs that are 8 TB or smaller and exclude VMDKs that are larger than 8 TB:

```
backup vm VM1 -vmskipvirtualdisks=yes -vmmaxvirtualdisks=8
```

Vmskipmaxvmdks

The `vmskipmaxvmdks` option specifies how the backup operation processes VMware virtual machine disks (VMDKs) that exceed the maximum disk size.

In V7.1.4 and later, `vmskipmaxvmdks` is renamed `vmskipmaxvirtualdisks`. Although `vmskipmaxvirtualdisks` is the preferred name, the client still processes backup operations with the `vmskipmaxvmdks` name.

See “`Vmskipmaxvirtualdisks`” on page 590 for information about setting this option.

Vmtagdatamover

Use the `vmtagdatamover` option to enable tagging support in the backup-archive client (data mover). When this option is enabled, the client manages backups of VMware virtual machines according to the backup management tags that are set by the IBM Data Protection extension of the vSphere Web Client or with tools such as VMware vSphere PowerCLI version 5.5 R2 or later.

If you use the IBM Data Protection extension to manage backups, you can exclude virtual machines from backup operations or set the retention policy of virtual machine backups.

When tagging support is enabled, the following categories and tags are created in the virtual machine inventory.

Category	Tag	Tag description
Backup Management (IBM Spectrum Protect)	Excluded	The object is excluded from backups by IBM Spectrum Protect
Management Class (IBM Spectrum Protect)	<i>Management_class_name</i>	The policy used for retention settings in IBM Spectrum Protect

Where:

Excluded

Excludes a virtual machine from scheduled backup services.

Management_class_name

The name of the retention policy that defines how long backup versions are kept on the Tivoli Storage Manager server or how many backup versions can exist on the server before they are expired.

If you do not use the *Management_class_name* tag value, the default retention policy for the VMware datacenter node is used.

You can use scripts to apply these categories and tags to virtual machines to exclude the virtual machines from scheduled backup services, or to specify the retention policy for the backups.

When the `vmtagdatamover` option is set to *yes*, all tags that are assigned to a virtual machine are backed up during **backup vm** operations. The tags are restored when the **restore vm** command is run.

Requirements:

- VMware vCenter Server must be at Version 6.0 Update 1 or later.
- Extra permissions are required for the account that is used for backup or restore operations. These new vCenter permissions are required to perform category and tagging operations. Ensure that the following user permissions are set on the root vCenter Server:

```
Inventory Service > vSphere Tagging > Assign or Unassign vSphere Tag
Inventory Service > vSphere Tagging > Create vSphere Tag
Inventory Service > vSphere Tagging > Create vSphere Tag Category
Inventory Service > vSphere Tagging > Delete vSphere Tag
Inventory Service > vSphere Tagging > Delete vSphere Tag Category
Inventory Service > vSphere Tagging > Modify UsedBy Field For Tag
Inventory Service > vSphere Tagging > Modify UsedBy Field For Category
```

For more information about setting vCenter permissions for backup and restore operations, see technote 7047438.

- In order for the Data Protection for VMware vSphere GUI to function correctly with tagging support, ensure that the following requirements are met during the installation of the GUI:
 - At least one data mover and the Data Protection for VMware vSphere GUI must be installed on the same server. This data mover node must be configured so that the vCenter server credentials are saved. You can save the credentials by running the configuration wizard to save the data mover node password, or by using the **dsmc set password** command in the data mover command line.

If you use other data movers, running on virtual machines or physical machines as additional data movers, you can install them on other servers. For tagging support, all these data movers must also be configured with the `vmtagdatamover=yes` option. These additional data movers do not require the Data Protection for VMware vSphere GUI to be installed on the same server in order for them to work correctly as tag-based data movers.
 - For Linux data movers, ensure that you specify the data mover installation directory and the Java™ shared library `libjvm.so` in the `LD_LIBRARY_PATH` environment variable. The path to `libjvm.so` is used for tagging support when you enable the `vmtagdatamover` option on the data mover. For instructions, see "Setting up the data mover nodes in a vSphere environment".
 - On Linux operating systems, the Data Protection for VMware vSphere GUI must be installed by using the default user name (`tdpvmware`).

- On Linux data mover nodes, the default password file (/etc/adsm/TSM.sth) must be used.


Supported clients

This option can be used with supported x86_64 Linux clients.

Options file

You can specify this option in the client system-options file (dsm.sys) or on the command line for the **backup vm** command. You can also include this option on the Tivoli Storage Manager server in a client option set. You cannot set this option in the Preferences Editor.

Syntax



Parameters

No The client ignores any backup management settings or tags that are attributed to the VMware asset. This value is the default.

Yes

The client manages backups based on the backup management settings in the IBM Data Protection extension or based on the tag values that are attributed to the VMware asset.

When tagging support is enabled, some client options might be affected by the backup management settings:

- When you use backup management settings or tags to control which VMware virtual machines are backed up, the tag values might overlap the domain.vmfull client option setting. While the domain.vmfull option defines what virtual machines the client protects, the "Excluded" tag overrides what is defined by the domain.vmfull option.

For example, the following options file statement specifies what is backed up during full virtual machine backup operations:

```
DOMAIN.VMFULL VMHOSTCLUSTER=cluster01,cluster02;VM=Dept20*
```

If you use backup management settings or tags to exclude virtual machine Dept204, the Dept204 virtual machine is not backed up.

- The retention policy setting in the IBM Data Protection extension or the tag setting for the Management Class (IBM Spectrum Protect) category overrides the include.vm and vmc client options, but does not override the vmctlmc option or the Tivoli Storage Manager for vCloud Director vappmc option.

Example

Options file:

```
vmtagdat yes
```

Command line:

```
-vmtagdat=yes
```

Related tasks:

➡ Enabling tagging support

Related reference:

“Domain.vmfull” on page 369

“Include.vm” on page 430

“Vmmc” on page 584

“Vappmc” on page 561

“Vmctlmc” on page 572

“Set Vmtags” on page 740

Vmverifyifaction

Use this option to specify the action to perform if the data mover detects integrity problems with the latest CTL and bitmap files for a virtual machine.

This option affects backup processing for a VM guest only when all of the following conditions are true:

- The previous backup operation for the VM guest was an incremental-forever-incremental backup (mode=ifincremental)
- The current backup operation for the VM guest is an incremental-forever-incremental backup
- The data mover detected an integrity problem with the CTL and bitmap data from the previous incremental-forever-incremental backup operation
- The `vmverifyiflatest` option is set to yes

If all of these conditions are not true for a virtual machine, the backup occurs as it normally would; the action that is specified by this option is not initiated.

Supported clients


This option is valid for Linux clients that act as a data mover for VMware guest backups.

Options file

Set this option in the client options file (`dsm.opt`) or the client system options file (`dsm.sys`).

This option can also be included in a client options set, as a parameter on a **backup vm** command, or on the **options** parameter in a schedule definition.

Syntax



Parameters

FAILbackup

This action fails the backup operation. The following messages are written to the data mover error log file (`dsmerror.log`):

ANS9921E Virtual machine disk, *vm_name (disk_label)*,
verification check failed (*xxx/yy*).

The xxx/yyy in the message indicate the size of the bitmap (xxx) and CTL files (yyy).

ANS9919E Failed to find the expected control files for *vm_name*

Perform a full VM backup (set `-mode=IFFull` for the affected virtual machines at a time of your choosing. An alternative is to use the `-vmverifyifaction=forcefull` on the next scheduled incremental-forever-incremental operation to force a full backup of those VMs, if you determine that your scheduled backup window can contain the full VM backups for these VMs. This value is the default action value.

FORCEfull

This action changes the backup mode from `-mode=ifincremental` to `-mode=iffull`; the current backup becomes a full VM backup. The full VM backup is initiated for you. The following messages are written to the data mover error log file (`dsmerror.log`):

ANS9921E Virtual machine disk, *vm_name* (*disk_label*),
verification check failed (*xxx/yyy*)

The xxx/yyy in the message indicate the size of the bitmap (xxx) and CTL files (yyy).

ANS9919E Failed to find the expected control files for *vm_name*

ANS9922I VMVERIFYIFlatest is enabled for *vm_name* (action: FORCEFULL).

ANS9920W Forcing a full vm backup for *vm_name*

Use this option if your current backup window can contain a full VM backup of the affected virtual machines.

PREview

This action does not perform any backups. Instead, the CTL and bitmap data for each VM guest that is processed by the **backup vm** command is restored to a temporary location, where it is checked for integrity. If the integrity check fails, the following messages are written to the data mover error log file (`dsmerror.log`):

ANS9921E Virtual machine disk, *vm_name* (*disk_label*),
verification check failed (*xxx/yyy*)

The xxx/yyy in the message indicate the size of the bitmap (xxx) and CTL files (yyy).

ANS9919E Failed to find the expected control files for *vm_name*

ANS9922I VMVERIFYIFlatest is enabled for *vm_name* (action: PREVIEW)

Use this option to validate the integrity of the incremental-forever-incremental backups (`-mode=ifincremental`) that you previously created for one or more a virtual machines.

If the messages indicate that some VMs failed the integrity checks, start a full VM backup (`-mode=iffull`) at a time of your choosing. Alternatively, set `-vmverifyifaction=forcefull` on the next scheduled incremental-forever-incremental operation to force a full backup of those VMs. The backup window must be large enough to accommodate one or more full VM backups.

Vmverifyiflatest

This option applies only to VMware virtual machine (VM) backup operations that use the incremental-forever-incremental backup mode (that is, a **backup vm** command with `-mode=IFIncremental` specified). If this `vmverifyiflatest` option is

enabled, the data mover runs an integrity check on the CTL and bitmap files that were created on the server during the last backup, if the last backup was an incremental-forever-incremental backup.

If the files pass the integrity tests, the virtual machine is restorable. The current backup proceeds and adds another snapshot to the chain of snapshots for the virtual machine.

If the files fail the integrity tests, the virtual machine is not restorable. The data mover then performs another action, which you specified on the `vmverifyifaction` option. You can set `vmverifyifaction` to create a full VM backup immediately, or you can fail the backup completely, and run a full VM backup at another time. A third parameter can be set to just verify the CTL and bitmap files for a virtual machine, without creating a new backup snapshot.

Verification can be performed only if the previous backup operation for the VM used `mode=IFIncr`, and if the current backup operation also uses `mode=IFIncr`. This option has no effect on the other virtual machine backup modes.

Important:

If this option is set to `no`, VM backup processing continues without any verification tests. The processing resources that are involved in performing the integrity checks is negligible. To ensure the continued integrity of your incremental-forever-incremental backup chain, set or use the default value (`vmverifyiflatest yes`). Do not set this option to `no`, unless you are directed to do so, by IBM support.

Supported clients


This option is valid for Linux clients that act as a data mover for VMware guest backups.

Options file

Set this option in the client options file (`dsm.opt`) or the client system options file (`dsm.sys`).

This option can also be included in a client options set, as a parameter on a **backup vm** command, or on the **options** parameter in a schedule definition.

Syntax



Parameters

YES

This setting specifies that validation of the CTL and the bitmap data is performed for each VM that is processed by the current incremental-forever-incremental (`mode=IFIncr`) backup operation, if the previous backup operation for that VM was also an incremental-forever-incremental backup. This value is the default value.

NO This setting specifies that validation of CTL and bitmap data does not occur

during incremental-forever-incremental backup processing. Do not set this value unless directed to do so by IBM support.

Examples

Options file:

```
vmverifyiflatest yes
```

Command line:

```
dsmc backup vm vm1 -mode=ifincremental -vmverifyiflatest=yes
```

Vmstortransport

The `vmstortransport` option specifies the preferred transports order (hierarchy) to use when backing up or restoring VMware virtual machines. If you do not include a given transport using this option, that transport is excluded and is not used to transfer data.

The transport order that you specify determines how the VMware API for Data Protection (VADP) accesses virtual disk data, but it does not influence the data path that is used between the backup-archive client and the Tivoli Storage Manager server. Valid transports include any order or combination of the following options:

- nbd** Network based data transfer. Access virtual disk data using the LAN. This transport path is generally available in all configurations.
- nbdssl** Same as `nbd`, but the data is encrypted before being sent over the LAN. Encryption can decrease performance.
- san** Storage Area Network transfer: Access virtual disk data using the SAN.
- hotadd** If you use the backup-archive client in a virtual machine, the `hotadd` transport allows the transport of backed up data to dynamically added storage.

Separate each transport option from the others with a colon, for example, `san:nbd:nbdssl:hotadd`.

If you do not specify a transport hierarchy, the default transport selection order is `san:hotadd:nbdssl:nbd`.


The first transport that is available is used to transfer the data. If you want to prevent data transport over a particular path, do not include it in the transport list. For example, if it is important to not disrupt LAN traffic, omit the `nbd` transports from the hierarchy.

Set this option in `dsm.sys`.

Supported clients

This option is valid for Linux clients that are configured to back up or restore virtual machine files using VADP.

Syntax



Examples

If the SAN is available, do not transport backups or restores over the LAN
 VMVSTORTRANSPORT san

The backup-archive client is running in a virtual machine, but do not use the hotadd transport
 VMVSTORTRANSPORT nbdssl:nbd

Use the LAN transport, even if nbdssl is available, to obtain better performance
 VMVSTORTRANSPORT nbd

The SAN transport is preferred, but use nbd when the SAN is not available, and do not use nbdssl or hotadd
 VMVSTORTRANSPORT san:nbd

Vmtimeout

VMTIMEOut specifies the maximum time, in seconds, to wait before abandoning a **backup vm** operation, when the INCLUDE.VMTSMVSS option is used to provide application protection. To use this option, the Tivoli Storage Manager for Virtual Environments license must be installed.

Each **backup vm** operation that is performed on a virtual machine that is protected by a INCLUDE.VMTSMVSS option is subject to a timer. The timer value determines how many seconds the client should wait for the application to quiesce activity and truncate its logs so the backup can be performed. The default time out value is sufficient for most environments. However, if your application data cannot be backed up because the application needs additional time to prepare for the snapshot, you can increase the time out value. This timer applies only to **backup vm** operations when the INCLUDE.VMTSMVSS option is set for a virtual machine.

Supported clients

This option can be used with supported x86_64 Linux clients.

Options file

Place this option in the client options file. It cannot be set on the command line or in the Preferences editor.

Syntax



Parameters

time_out

Specifies the time to allow, in seconds, for backup operations to complete when a virtual machine is protected by the application protection option, INCLUDE.VMTSMVSS. The value specified must be an integer between 180 and 500. The default is 180 seconds.

Examples

Options file

VMTIMEout 500

Command line

Not applicable; this option cannot be set on the command line.

Related reference:

"INCLUDE.VMTSMVSS" on page 435

Webports

The webports option enables the use of the web client outside a firewall.

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 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 and the web client agent service.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the `dsm.sys` file within a server stanza. To set this option in the Client Preferences editor, click **Edit > Client Preferences > Web Client**, and specify the ports in the **Web Agent Port** and **Web Client Acceptor Port** fields.

Syntax

►►—WEBPorts— —cadport— —agentport—►►

Parameters

cadport

Specifies the required client acceptor port number. The range of values is 1000 through 32767. If a value is not specified, the default, zero (0), causes TCP/IP to randomly assign a free port number.

agentport

Specifies the required web client agent service port number. The range of

values is 1000 through 32767. If a value is not specified, the default, zero (0), causes TCP/IP to randomly assign a free port number.

Examples

Options file:

```
webports 2123 2124
```

Command line:

Does not apply.

Wildcardsareliteral

The `wildcardsareliteral` option specifies whether question marks (?) and asterisks (*) are interpreted literally, when they are included in a file list specification on a `filelist` option.

Ordinarily, the client does not accept wildcard characters (?) and (*) in a file list specification that is included on a `filelist` option. Some file systems, such as the IBM General Parallel File System (GPFS), allow these characters in file and directory names. To prevent errors that would otherwise occur, when file specifications are included on a `filelist` option and they contain wildcard characters, set `wildcardsareliteral` yes. When `wildcardsareliteral` is set to yes, question marks (?) and asterisks (*) that are included in a file list specification on the `filelist` option are interpreted literally, and not as wildcard characters.

This option applies to any command that accepts a `filelist` option as command parameter.

Supported Clients

This option is valid for all supported platforms. The option is applied to any command that takes a file list specification as a parameter.

Options File

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

Syntax



Parameters

no Specifies that question marks and asterisks are interpreted as wildcards when used in a file list specification that is included on a `filelist` option. No is the default. If a file list specification on a `filelist` option includes a question mark or asterisk, an error occurs and the file specification cannot be processed.

yes

Specifies that asterisks and question marks in a file list specification that is included on a `filelist` option are interpreted literally, and not as wildcard characters. Specify this value if you are backing up files from a file system that allows wildcard characters in file or directory names.

Examples

Options file:

```
WILDCARDSARELITERAL YES
```

Command line:

Assuming that the file system allows wildcard characters in paths, the following are examples of files in a file list specification that can be successfully processed if WILDCARDSARELITERAL is set to YES.

Assume that the command issued is `dsmc sel -filelist=/home/user1/important_files`, where `important_files.txt` contains the list of files to process.

`important_files.txt` contains the following list of files:

```
/home/user1/myfiles/file?9000  
/home/user1/myfiles/?file  
/home/user1/myfiles/**README**version2  
/home/user1/myfiles/ABC?file*
```

If both WILDCARDSARELITERAL and QUOTESARELITERAL are both set to YES, the following backups can be successfully processed:

```
/home/user1/myfiles/"file?  
/home/user1/myfiles/?file'  
/home/user1/myfiles/**"README Tomorrow"**  
/home/user1/myfiles/file*
```

Related information

For information about the `filelist` option, see “Filelist” on page 402.

For information about syntax for file specifications, see “Specifying input strings that contain blank spaces or quotation marks” on page 138.

“Quotesareliteral” on page 487

Chapter 11. Using commands

The backup-archive client provides a command-line interface (CLI) that you can use as an alternative to the graphical user interface (GUI). This topic describes how to start or end a client command session and how to enter commands.

The following is a list of tasks related to entering commands.

- “Start and end a client command session” on page 606
- “Enter client command names, options, and parameters” on page 607
- “Wildcard characters” on page 610

The following table provides an alphabetical list of the commands and a brief description.

Table 78. Commands

Command	Description
archive “Archive” on page 611	Archives files from a workstation to Tivoli Storage Manager storage.
backup fastback “Backup FastBack” on page 616	Backs up volumes specified by the fbpolycname, fbclientname and fbvolumename options for long term retention.
backup group “Backup Group” on page 618	Creates and backs up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.
backup image “Backup Image” on page 620	Creates an image backup of one or more file systems or logical volumes that you specify.
backup nas “Backup NAS” on page 626	Creates an image backup of one or more file systems belonging to a Network Attached Storage (NAS) file server.
backup vapp Backup Vapp	Backs up a single VMware vApp, a list of vApps, or all vApps within an organization virtual data center.
backup vm “Backup VM” on page 634	Backs up virtual machines specified in the vmlist option.
cancel process “Cancel Process” on page 641	Displays a list of current NAS (if NDMP support is enabled) image backup and restore processes for which the administrative user has authority.
cancel restore “Cancel Restore” on page 642	Displays a list of restartable restore sessions from which you can select one to cancel.
delete access “Delete Access” on page 642	Deletes authorization rules for files that are stored on the server. On those clients that support image backup, this command deletes authorization rules for images that are stored on the server.
delete archive “Delete Archive” on page 643	Deletes archived files from Tivoli Storage Manager server storage.
delete backup “Delete Backup” on page 645	Deletes active and inactive backup files from Tivoli Storage Manager server storage.

Table 78. Commands (continued)

Command	Description
delete file space “Delete Filespace” on page 648	Deletes file spaces in Tivoli Storage Manager server storage.
delete group “Delete Group” on page 649	Deletes a group backup on the Tivoli Storage Manager server.
expire “Expire” on page 651	Inactivates backup objects that you specify in the file specification or with the <code>filelist</code> option.
help “Help” on page 652	Displays a Table of Contents of help topics for the command-line client.
incremental “Incremental” on page 654	Backs up all new or changed files or directories in the default client domain or from file systems, directories, or files you specify, unless you exclude them from backup services.
loop “Loop” on page 660	Starts an interactive command session.
macro “Macro” on page 661	Executes commands within a macro file that you specify.
monitor process “Monitor Process” on page 662	Displays a list of current NAS image backup and restore processes from which you can select one to cancel.
preview archive “Preview Archive” on page 663	Simulates an archive command without sending data to the server.
preview backup “Preview Backup” on page 664	Simulates a backup command without sending data to the server.
query access “Query Access” on page 665	Displays a list of current authorization rules.
query archive “Query Archive” on page 665	Displays a list of archived files.
query backup “Query Backup” on page 668	Displays a list of backup versions.
query backupset “Query Backupset” on page 670	Queries a backup set from a local file or the Tivoli Storage Manager server. On those clients that support tape devices, this command can query a backup set from a tape device.
query file space “Query Filespace” on page 674	Displays a list of file spaces in Tivoli Storage Manager storage. You can also specify a single file space name to query.
query group “Query Group” on page 676	Displays information about group backups and their members.
query image “Query Image” on page 678	Displays information about image backups.
query inclexcl “Query Inclexcl” on page 679	Displays a list of include-exclude statements in the order in which they are processed during backup and archive operations.
query mgmtclass “Query Mgmtclass” on page 680	Displays information about available management classes.
query node “Query Node” on page 681	Displays all the nodes for which an administrative user ID has authority to perform operations.
query options “Query Options” on page 682	Displays all or part of your options and their current settings.

Table 78. Commands (continued)

Command	Description
query restore “Query Restore” on page 683	Displays a list of your restartable restore sessions in the server database.
query schedule “Query Schedule” on page 684	Displays information about scheduled events for your node.
query session “Query Session” on page 684	Displays information about your session, including the current node name, when the session was established, server information, and server connection information.
query systeminfo “Query Systeminfo” on page 686	Gathers Tivoli Storage Manager system information and outputs this information to a file or the console.
query vapp Query Vapp	Determines which VMware vApps are backed up.
query vm “Query VM” on page 691	Verifies the successful backups of the virtual machines from the vStorage backup server.
restart restore “Restart Restore” on page 692	Displays a list of restartable restore sessions from which you can one to restart.
restore “Restore” on page 693	Restores copies of backup versions of your files from a Tivoli Storage Manager server.
restore backupset “Restore Backupset” on page 698	Restores a backup set from the Tivoli Storage Manager server or a local file. On those clients that support tape devices, this command can restore a backup set from a tape device.
restore group “Restore Group” on page 705	Restores specific members or all members of a group backup.
restore image “Restore Image” on page 707	Restores a file system or raw volume image backup.
restore nas “Restore NAS” on page 710	Restores the image of a file system belonging to a Network Attached Storage (NAS) file server.
restore vapp Restore Vapp	Restores a single VMware vApp.
restore vm “Restore VM” on page 715	Restores a full VM backup, and returns the full VM backup files to the <code>vmbackdir</code> directory on the vStorage backup server.
retrieve “Retrieve” on page 720	Retrieves copies of archived files from the Tivoli Storage Manager server.
schedule “Schedule” on page 723	Starts the client scheduler on the workstation.
selective “Selective” on page 725	Backs up selected files.
set access “Set Access” on page 728	Authorizes another user to access your backup versions or archived copies. On those clients that support image backup, this command can set authorization rules for images that are stored on the server.
set event “Set Event” on page 731	Allows you to specify the circumstances for when archived data is deleted.
set netappsvm Set Netappsvm	Associates the login credentials for a cluster management server with a NetApp storage virtual machine and the data SVM name (data Vserver). This command must be entered before you can create a snapshot difference incremental backup of a clustered NetApp volume.

Table 78. Commands (continued)

Command	Description
set password “Set Password” on page 734	Changes the Tivoli Storage Manager password for your workstation.

For proper operation, the was node must be restored to the same location and under the same name.

Important: To avoid problems, restore your data at the Network Deployment Manager node or Application Server node level only.

Related reference:

“Reading syntax diagrams” on page xiv

Start and end a client command session

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

Use batch mode when you want to enter a *single* client command. The backup-archive client processes the command and returns to the command prompt.

Use interactive mode when you want to enter a *series* of commands. Since the client establishes connection to the server only once for interactive mode, a series of commands can be processed more quickly. The client processes the commands and returns to the `tsm>` prompt.

Process commands in batch mode

Some options are valid *only* on the initial command line and not in interactive mode. These options generally affect the operation of the entire session.

For example, the command **dsmc query session -errorlogname=myerror.log** is accepted and it does name the error log. However, it is accepted simply because it appears in the initial command, even though the option is not valid for the query command.

There are also some options that are always valid on the initial command line as well as on individual commands in interactive mode. Therefore, certain options are accepted on the initial command line even though they have no effect on the command being entered. For example, **dsmc query session -subdir=yes** is a valid command, but in this case the *-subdir* option has no effect on the command that was entered.

When you enter a *single* command in batch mode, precede it with the executable program name, **dsmc**. For example, to process the **incremental** command in batch mode, you would enter:

```
dsmc incremental
```

The backup-archive client 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. The client only prompts you for your password if you are registering your workstation with a server or manually changing your password.

Related concepts:

Chapter 10, “Processing options,” on page 287

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, the backup-archive client prompts you before you enter the first command.

Type your user ID and password and press Enter.

You can also enter your password using the password option with the **loop** command, but your password appears on the screen. For example, if your password is *secret*, enter:

```
dsmc loop -password=secret
```

To end an interactive session, enter *quit* at the prompt.

Note for UNIX and Linux clients:

In loop mode, following a restore operation directly from tape, the mount point is not released in case additional restore requests are made to that volume. If you request a backup operation in the same session and that mount point is the only one available, the backup operation will stop with the following message:

```
Waiting for mount of offline media
```

In this case, the mount point is not released until one of the following conditions is met:

- The device class MOUNTRETENTION limit is satisfied.
- The client idletimeout period is satisfied.
- The dsmc loop session is closed after the restore operation completes, allowing you to start a subsequent loop mode session to perform the backup operation.

Enter client command names, options, and parameters

A client command can include one or more of these components: *Command name*, *options*, and *parameters*. The topics that follow describe each of these components.

Command name

The first part of a command is the command name. The command name consists of a single word, such as **help** or **schedule**, or an action word and an object for that action, such as **query archive**.

Enter the full command name, or its minimum abbreviation.

For example, you can enter any of the following versions of the **query schedule** command:

```
query schedule
q sc
q sched
query sc
```

Options

When you enter options with a command, always precede the option with a dash (-). Do not put a space between the dash and the option name.

Enter more than one option in any order in a command before or after the file specification. Separate multiple options with a blank space.

There are two groups of options that you can use with commands: Client options (set in your options file), or client command options (used on the command line).

- **Client options:** The group of options that are set in your client options file. You can override an option in the client options file when you enter the option with a command on the command line.
- **Client command options:** Use a client command option *only* when you enter the option with a command on the command line. You cannot set these options in an options file.

Related concepts:

“Client options reference” on page 316

Options in interactive mode

In interactive mode, options that you enter on the initial command line override the value that you specified in your options file.

This value remains in effect for the entire interactive session unless overridden by a different value on a given interactive command.

For example, if you set the `subdir` option to *yes* in your `dsm.opt` or `dsm.sys` file, and you specify `subdir=no` on the initial command line, the `subdir=no` setting remains in effect for the entire interactive session unless overridden by the `subdir=yes` value on a given interactive command. However, the `subdir=yes` value specified within the interactive session only affects the command on which it is entered. When that command completes, the value reverts back to `subdir=no`, the value at the beginning of the interactive session.

Parameters

Commands can have required parameters, optional parameters, or no parameters at all.

Required parameters provide information to perform a task. The most commonly required parameter is a file specification.

For example, if you want to archive a file named `budget.fin` from the project directory, you would enter the following:

```
dsmc archive /project/budget.fin
```

Some commands have optional parameters. If you do not enter a value for an optional parameter, the backup-archive client uses the default value. For example, the **restore** command includes a required parameter, **sourcefilespec**, that specifies the path and file name in storage that you want to restore. The optional parameter, **destinationfilespec**, specifies the path where you want to place the restored files. If you do not specify the **destinationfilespec**, by default, the client restores the files to the original source path. If you want to restore the files to a *different* directory, enter a value for **destinationfilespec**.

Example: Restore the file /project/budget.fin to the new path /newproj/newbudg.fin

```
dsmc restore /project/budget.fin /newproj/
```

Enter parameters in the order indicated in the command syntax diagram.

File specification syntax

There are some syntax rules that you need to know about when entering file specification parameters such as **filespec**, **sourcefilespec**, and **destinationfilespec**.

The following are the syntax rules:

- Do not use wildcards as part of the file space name or anywhere in the **destinationfilespec**. The one exception to this rule is the **set access** command where wildcards are permitted in the two lowest levels of the file spec.

Example: Allow access to all files in all directories in and subordinate to the /home directory:

```
set access backup /home/* * *  
set access backup /home/*/* * *
```

With UNIX clients, do not use wildcards in a directory path name, for example:

```
/home/j*asler/file1.c
```

- There is a maximum number of file specifications per command:
 - The **Query** commands can accept only one file specification.
 - The **restore** and **retrieve** commands can accept a source file specification and a destination file specification.
 - There is a limit of 20 operands on some commands. This limit is to prevent excessive sessions that are caused when wildcards are expanded by the UNIX shell command processor.

You can prevent shell expansion from causing you to go over the 20-operand limit by placing quotation marks around your source filespec expansion characters for restore commands.

Note: Using quotation marks has the side affect of causing a no-query restore.

You can use the `removeoperandlimit` option to specify that the backup-archive client removes the 20-operand limit. If you specify the `removeoperandlimit` option with the **incremental**, **selective**, or **archive** commands, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits.

- The length of a file specification is limited.
 - On AIX, HP-UX, Solaris, and Mac: The maximum number of characters for a file name is 255. The maximum combined length of the file name and path name is 1024 characters. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.
 - On Linux: The maximum length for a file name is 255 bytes. The maximum combined length of both the file name and path name is 4096 bytes. This length matches the PATH_MAX that is supported by the operating system. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that comprises a path and file name can vary. The actual limitation is the number of bytes in the path and file components, which might correspond to an equal number of characters.
- When you enter the **sourcefilespec**, if the directory name ends with /, then /* is implied.

When you enter a **destinationfilespec**, if the name ends with /, then it is considered a directory, otherwise it is considered a file.

The following example illustrates these two rules. Even though mydir and yourdir are directories, the command will fail because /* is implied after mydir, and yourdir is considered a file.

```
restore /home/mydir/ /away/yourdir
restore c:\home\mydir\ c:\away\yourdir
```

- If a file specification does not begin with a directory delimiter, the file specification is assumed to be a subdirectory of the current working directory. The client appends the file specification to the working directory to build the complete path.
- The only command that accepts a simple file space name is the **incremental** command. The following example is valid:

```
dsmc i /Users
```

The following example is not valid, because the command is the **selective** command:

```
dsmc sel /Users
```

Related reference:

“Filelist” on page 402

“Removeoperandlimit” on page 488

Wildcard characters

Use wildcard characters when you want to specify multiple files with similar names in *one* command. Without wildcard characters, you must repeat the command for each file.

In a command, you can use wildcard characters in the file name or file extension *only*. You cannot use them to specify destination files, file systems, or server names. You cannot specify a directory whose name contains an asterisk (*) or a question mark (?).

Valid wildcard characters that you can use include:

- * Asterisk. Matches zero or more characters.
- ? Question mark. Matches any single character at the present position.

The following table shows examples of each wildcard.

Table 79. Wildcard characters

Pattern	Matches	Does not match
Asterisk (*)		
ab*	ab, abb, abxxx	a, b, aa, bb
ab*rs	abrs, abtrs, abrsrs	ars, aabrs, abrss
ab*ef*rs	abefrs, abefghrs	abefr, abers
abcd.*	abcd.c, abcd.txt	abcd, abcdc, abcdtxt
Question Mark (?)		
ab?	abc	ab, abab, abzzz
ab?rs	abrs	abrs, abllrs
ab?ef?rs	abdefjrs	abefrs, abdefrs, abefjrs
ab??rs	abcdrs, abzzrs	abrs, abjrs, abkkkrs

Important: Use an asterisk (*) instead of a question mark (?) as a wildcard character when trying to match a pattern on a multibyte code page, to avoid unexpected results.

Note: In batch mode, enclose values containing wildcards in quotation marks. Otherwise, UNIX shells expand unquoted wildcards, and it is easy to exceed the 20 operand limit. It is more efficient to let the client process wildcard file specifications because many fewer server interactions are needed to complete the task. For example:

```
dsmc selective "/home/me/*.c"
```

Client commands reference

The following sections contain detailed information about each of the backup-archive client commands.

Information for each command includes the following information:

- A description of the command.
- A syntax diagram of the command.
- Detailed descriptions of the command parameters. If the parameter is a constant (a value that does not change), the minimum abbreviation appears in uppercase letters.
- Examples of using the command.

Archive

The **archive** command archives a single file, selected files, or all files in a directory and its subdirectories on a server.

Archive files that you want to preserve in their current condition. To release storage space on your workstation, delete files as you archive them using the `deletefiles` option. Retrieve the archived files to your workstation whenever you need them again.


Use the `snapshotroot` option with the **archive** command along with an independent software vendor application that provides a snapshot of a logical volume to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. The `snapshotroot` option does not provide any facilities to take a volume snapshot, only to manage data that is created by a volume snapshot.

AIX only: You can enable snapshot-based file archive by using the option `snapshotproviderfs=JFS2`.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

filespec

Specifies the path and name of the file you want to archive. Use wildcard characters to include a group of files or to include all files in a directory.

To include multiple file specifications, separate each *filespec* parameter with a space character. If multiple file specifications are included, and two or more of the specifications have common parent directories, then it is possible for the common directory objects to be archived more than once. The conditions under which this behavior occurs are runtime-dependent, but the behavior itself has no adverse effects.

For example, if the *filespec* is `/home/amr/ice.doc /home/amr/fire.doc`, then `/home` and `/home/amr` might be archived twice. The file objects `ice.doc`, and `fire.doc`, are archived only once.

If you want to avoid including the shared parent directory more than once, use separate, non-overlapping **archive** commands to archive each file specification.

If you archive a file system, include a trailing slash (`/home/`).

There is a limit of 20 operands. This limit prevents excessive sessions that are caused when wildcards are expanded by the UNIX shell command processor. You can prevent shell expansion from causing you to go over the 20-operand limit by placing quotation marks around file specifications that contain wildcards (`"home/docs/*"`).

You can use the **removeoperandlimit** option to specify that the 20-operand limit is removed. If you specify the **removeoperandlimit** option, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits. For example, remove the 20 operand limit to archive 21 file specifications:

selective -removeoperandlimit filespec1 filespec2 ... filespec21

You can use the **filelist** option, instead of file specifications, to identify which files to include in this operation. However, these two methods are mutually exclusive. You cannot include file specification parameters and use the **filelist** option. If the **filelist** option is specified, any file specifications that are included are ignored.

Table 80. Archive command: Related options

Option	Where to use
archmc	Command line only.
archsymb linkasfile	Client user-options file (dsm.opt) or command line.
autofsrename	Client options file (dsm.opt) only.
changingretries	Client system options file or command line.
compressalways	Client user-options file (dsm.opt) or command line.
compression	dsm.sys file within a server stanza or command line.
deletefiles	Command line only.
description	Command line only.
dirsonly	Command line only.
encryptiontype	dsm.sys file within a server stanza.
encryptkey	dsm.sys file within a server stanza.
filelist	Command line only.
filesonly	Command line only.
preserve lastaccessdate	Client user-options file (dsm.opt) or command line.
removeoperandlimit	Command line only.
snapshotcachesize	Client options file (dsm.opt) or include.fs option.
snapshotroot	Command line only.
subdir	Client options file (dsm.opt) or command line.
tapeprompt	Client options file (dsm.opt) or command line.
v2archive	Command line only.

Examples

Task Archive a single file that is named budget in the /home/proj1 directory.

Command: archive /home/proj1/budget

Task Archive all files in the /home/proj1 directory with a file extension of .txt.

Command: archive "/home/proj1/*.txt"

Task Archive all files in the directory tree that is headed by the /home directory.

Command: archive -subdir=yes "/home/*"

Task Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, archive the /usr/dir1/sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name /usr.

Command: dsmc archive /usr/dir1/sub1/ -subdir=yes
-snapshotroot=/snapshot/day1

Related concepts:

“Snapshotproviderfs” on page 526

614 IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients: Installation and User's Guide

3. When you specify a FBClientName value, there must be only one value for FBPolicyName.
4. You can specify up to 10 values for FBClientName if only one PolicyName is specified, and no values for FBVolumeName are specified.
5. When you specify the FBVolumeName option, you can have only one FBPolicy, and only one FBClientName specified.
6. You can specify multiple FBVolumeNames if condition 5 is satisfied.
7. You must always specify the FBReposLocation option for Linux.

Parameters

Table 81. Archive FastBack command: Related options

Option	Where to use
fbpolicyname "Fbpolicyname" on page 398	Command line and scheduler.
fbserver "Fbserver" on page 400	Command line and scheduler.
fbclientname "Fbclientname" on page 396	Command line and scheduler.
fbvolumename "Fbvolumename" on page 401	Command line and scheduler.
fbreposlocation "Fbreposlocation" on page 399	Command line and scheduler.
fbbranch "Fbbranch" on page 396	Command line and scheduler.
archmc "Archmc" on page 318	Command line and scheduler.

Examples

Command line:

The backup-archive client is installed on a Linux proxy client machine. Use this command to archive all FastBack volumes for all Linux FastBack clients that are defined for FastBack policy1:

```
dsmc archive fastback -fbpolicyname=Policy1
-fbserver=myfbserver -fbreposlocation=myfbserver@WORKGROUP
```

The FastBack server name, -myFbDrHub is the short host name of the FastBack disaster recovery hub server where the repository is located.

The -fbreposlocation parameter specifies the location of the repository. The repository location is required. If you do not provide the repository location, the command fails.

FBServer should point to the short host name of the FastBack DR hub in this case.

Command line:

The repository, rep_server1, is located on a FastBack disaster hub, myFbDrHub

```
dsmc archive fastback -fbpolicyname="Policy 1"
-fbserver=myFbDrHub -fbreposlocation=\\myFbDrHub\rep_server1
```

The repository location is required. If you do not provide the repository location, the command fails.

The FastBack server name, -myFbDrHub, is the short host name of the FastBack disaster recovery hub where the repository is located.

FBServer should point to the short host name of the FastBack DR hub in this case.

Command line:

Archive all volumes protected by FastBack policy named policy1 from the FastBack server named basil:

```
dsmc archive fastback -Fbpolicyname=policy1
-FBServer=basil -ARCHMC="my_tsm_mgmt_class"
-fbreposlocation=basil@WORKGROUP
```

Related concepts:

"Configuring the client to back up and archive Tivoli Storage Manager FastBack data" on page 94

Backup FastBack

Use the **backup fastback** command to back up Tivoli Storage Manager FastBack volumes specified by the fbpolicyname, fbclientname and fbvolumename options for long-term retention.

Before using this command, configure the client to back up and archive Tivoli Storage Manager FastBack data. Also, before you issue this command, at least one snapshot should exist in the Tivoli Storage Manager FastBack repository for the Tivoli Storage Manager FastBack policy being archived or backed up.

If a policy specification contains both Windows and Linux FastBack clients, only the Linux volumes will be backed up or archived to the Tivoli Storage Manager server by the Linux backup-archive client.

Tivoli Storage Manager FastBack options are supported for the incremental backup of the latest snapshots, depending on the option specified:


- All clients and volumes associated with a Tivoli Storage Manager FastBack policy or a list of Tivoli Storage Manager FastBack policies.
- All volumes associated with a specific FastBack client or a list of FastBack clients for a given Tivoli Storage Manager FastBack policy.
- A specific volume or volumes associated with a FastBack client for a given Tivoli Storage Manager FastBack policy.

Supported Clients

This command is valid for Linux x86_64 clients that are configured as Tivoli Storage Manager FastBack dedicated proxies.

Syntax

```
►►—BACKUP FASTBack—FBPolicyname—=—name—FBServer—=—name—►
```



Important:

1. At least one FBpolicyName is always required.
2. You can specify up to 10 values for FBPolicyName, if no values are specified for both FBClientName and FBVolumeName.
3. When you specify a FBClientName value, there must be only one value for FBPolicyName.
4. You can specify up to 10 values for FBClientName if only one PolicyName is specified, and no values for FBVolumeName are specified.
5. When you specify the FBVolumeName option, you can have only one FBPolicy, and only one FBClientName specified.
6. You can specify multiple FBVolumeNames if condition 5 is satisfied.
7. You must specify the FBReposLocation option.

Table 82. Backup FastBack command: Related options

Option	Where to use
fbpolicyname "Fbpolicyname" on page 398	Command line and scheduler.
fbserver "Fbserver" on page 400	Command line and scheduler.
fbclientname "Fbclientname" on page 396	Command line and scheduler.
fbvolumename "Fbvolumename" on page 401	Command line and scheduler.
fbreposlocation "Fbreposlocation" on page 399	Command line and scheduler.
fbbranch "Fbbranch" on page 396	Command line and scheduler.
backmc "Backmc" on page 329	Command line and scheduler.

Examples

Command line:

The backup-archive client is installed on a Linux proxy client machine. Use this command to back up all FastBack volumes for all Linux FastBack clients that are defined for FastBack policy1:

```
dsmc backup fastback -fbpolicyname=Policy1  
-fbserver=myfbserver  
-fbreposlocation=myfbserver@WORKGROUP
```

The repository location is required. If you do not provide the repository location, the command will fail.

The FastBack server name, -myfbserver, is the short host name of the FastBack server where the repository is located.

Command line:

The repository, rep_server1, is located on a FastBack disaster hub, myFbDrHub

```
dsmc backup fastback -fbpolicyname="Policy 1"  
-fbserver=myFbDrHub -fbreposlocation=\\myFbDrHub\\rep_server1
```

The FastBack server name, -myFbDrHub, is the short host name of the FastBack disaster recovery hub server where the repository is located.

The -fbreposlocation option specifies the location of the repository. The repository location is required. If you do not provide the repository location, the command fails.

The FBServer option should point to the short host name of the FastBack DR hub in this case.

Command line:

Back up all volumes protected by FastBack policy named policy1 from the FastBack server named basil:

```
dsmc backup fastback -Fbpolicyname=policy1  
-FBServer=basil -BACKMC="my_tsm_mgmt_class"  
-fbreposlocation=basil@WORKGROUP
```

Related concepts:

"Configuring the client to back up and archive Tivoli Storage Manager FastBack data" on page 94

Backup Group

Use the **backup group** command to create and back up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.

AIX only: You can enable snapshot-based group backup by using the option snapshotproviderfs=JFS2.

A group backup allows you to create a consistent point-in-time backup of a group of files that is managed as a single logical entity. Objects in the group are subject to the following processing rules:

- Management class rebinding for grouped objects:
 - During full backups, all objects in a backup group are assigned to the same management class.
 - During differential backups, if a new management class is specified on an include statement for an existing backup group, the following behavior occurs:
 - Any new and changed objects in the backup group are bound to the new management class.

- Any member objects of the group that are not changed appear as though they have not been bound to the new management class. These unchanged objects are not included in the **Total number of objects rebound** statistics that are displayed when the **Backup Group** command completes.
- The unchanged objects are reassigned to a newly created backup group, and the new backup group is bound to the new management class. However, the original management class name is still displayed for the unchanged group objects.

Even though the original management class name is still displayed for the unchanged objects, they are effectively bound to the new management class of the backup group.

- Existing **exclude** statements for any files in the group are ignored.
- All objects in the group are exported together.
- All objects in the group are expired together as specified in the management class. No objects in a group are expired until all other objects in the group are expired, even when another group they belong to gets expired.
- If you are performing full and differential group backups to a sequential device, during a restore the data is in no more than two locations. To optimize restore time, perform periodic full backups to back up the data to one location on the sequential media.
- During a full group backup, all objects in the filelist are sent to the server. During a differential group backup, only data that has changed since the last full backup is sent to the server. Objects in the filelist that have not changed since the last full backups are assigned as members of the differential group backup. This data is not resent to the server, reducing backup time.

The **backup group** command requires the following options:

filelist

Specifies a list of files to add to a new group.

groupname

Specifies the fully qualified name of the group containing a list of files.

virtualfsname

Specifies the name of the virtual file space for the group on which you want to perform the operation. The **virtualfsname** option cannot be the same as an existing file space name.

mode

Specifies whether you want to back up all of the files in the filelist or only files that have changed since the last full backup.

Note:

1. If any file in the group backup fails, the entire group backup fails.
2. Use the **query group** command to query members of a group backup on the Tivoli Storage Manager server.
3. Use the **restore group** command to restore specific members or all members of a group backup on the server.
4. Unless you are running Mac OS X, use the **delete group** command to delete a specific group backup from the server.
5. Use the **query filespace** command to display virtual file space names for your node that are stored on the server.
6. A group backup can be added to a backup set.

Supported Clients

This command is valid for all UNIX and Linux clients except Mac OS X.

Syntax

►►Backup Group— *options* —►►

Parameters

Table 83. Backup Group command: Related options

Option	Where to use
filelist "Filelist" on page 402	Command line only.
groupname "Groupname" on page 412	Command line only.
mode "Mode" on page 454	Command line only.
snapshotproviderfs "Snapshotproviderfs" on page 526	System-options file (dsm.sys) within a server stanza or with the include.fs option.
virtualfsname "Virtualfsname" on page 565	Command line only.

Examples

Task Perform a full backup of all the files in the /home/dir1/filelist1 file to the virtual file space name accounting containing the group leader /home/group1 file.

Command:

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

Related information

"Include options" on page 421

"Query Group" on page 676

"Restore Group" on page 705

"Delete Group" on page 649

"Query Filespace" on page 674

Backup Image

The **backup image** command creates an image backup of one or more volumes on your system.

You can use the **backup image** command to back up NTFS or ReFS, or unformatted RAW volumes. If a volume is NTFS-formatted, only those blocks that are used by the file system are backed up. On ReFS volumes, all blocks are backed up.

If you set the **imagegapsize** option to 0, all blocks, including unused blocks at the end of the volume, are backed up.

If you specify an AIX JFS2 file system for image backup, only those blocks that are used by the file system are backed up. If you set the **imagegapsize** option to zero, all blocks, including blocks at the end of the volume, are backed up.

Note:

1. AIX only: By default, snapshot-based image backup is enabled for JFS2 volumes. To turn off snapshot-based image backups, set **-snapshotproviderimage=NONE** on this command.
2. For the Linux clients, image backup is only supported on partitions with id 0x83 or logical volumes that are created with the Linux Logical Volume Manager. Backing up other partitions, such as extended partitions that contain mounted file systems or database data, can produce inconsistent backup data if the data changes during the image backup operation.
3. For the Linux client, image backup of DASD devices with raw-track access mode on Linux on z Systems is not supported.
4. Backup image is not supported on any GPFS file system.
5. The Tivoli Storage Manager API must be installed to use the **backup image** command.
6. When you change the attribute of a JFS2 file system to an HSM-managed file system, an image backup is not done for that file system.

Important: The last incremental backup time refers to the server time and the file modification time refers to the client time. If the client and server time are not synchronized, or the client and server are in different time zones, this affects incremental-by-date backup and image backup where **mode=incremental**.

The client backs up the files that have modification dates and times (on the client) that are later than the date and time of the last incremental backup of the file system on which the file is stored (on the server).

If the server time is ahead of the client time, incremental-by-date backups, or image backup with **mode=incremental**, skip the files, which had been created or modified after the last incremental or image backup with a modification date earlier than the last incremental backup time stamp.

If the client time is ahead of the server time, all files that had been created or modified before the last incremental or image backup and have a modification time stamp later than the last incremental backup time stamp, are backed up again. Typically, these files would not get backed up because they had already been backed up.

The backup date can be checked by the **query filespace** command.

The backup-archive client must support the raw device type on the specific platform to perform an image backup of a raw device. You can perform an image backup only on local devices. Clustered devices or file systems as well as devices

or file systems that are shared between two or more systems are not supported. If you want to perform an image backup for a file system that is mounted on a raw device, the raw device must be supported.


Use the **include.image** option to include a file system or logical volume for image backup, or to specify volume-specific options for image backup.

The **backup image** command uses the **compression** option.

Supported Clients

This command is valid for AIX, HP-UX, all Linux clients, and Solaris.

Syntax



Parameters

filespec

Specifies the name of one or more logical volumes. If you want to back up more than one file system, separate their names with spaces. Do not use pattern matching characters. If you do not specify a volume name, the logical volumes that are specified with the **domain.image** option are processed. If you do not use the **domain.image** option to specify file systems to process, an error message is displayed and no image backup occurs.

Specify the file space over which the logical volume is mounted or the logical volume name. If there is a file system that is configured in the system for a given volume, you cannot back up the volume with the device name.

For example, if the /dev/lv01 file space is mounted on the /home volume, you can issue backup image /home, but backup image /dev/lv01 fails with an error:

ANS1063E Invalid path specified

Note: For Sun systems, specify either a file system name or a raw device name (block device type).

Table 84. Backup Image command: Related options

Option	Where to use
asnodename	Client options file (dsm.opt) or command line.
asnodename "Asnodename" on page 319	Client system options file (dsm.sys) or command line.
compressalways "Compressalways" on page 338	Client system options file (dsm.sys) or command line.
compression "Compression" on page 339	Client options file or command line.
dynamicimage "Dynamicimage" on page 375	Use with the backup image command or the include.image option in the options file.

Table 84. Backup Image command: Related options (continued)

Option	Where to use
imagegapsize "Imagegapsize" on page 417	Use with the backup image command, the include.image option, or in the options file.
mode "Mode" on page 454	Command line only.
postsnapshotcmd "Postsnapshotcmd" on page 476	Use with the backup image command, the include.image option, or in the options file.
presnapshotcmd "Presnapshotcmd" on page 483	Use with the backup image command, the include.image option, or in the options file.
snapshotcachesize "Snapshotcachesize" on page 525	Use with the backup image command, the include.image option, or in the options file.
snapshotproviderimage "Snapshotproviderimage" on page 527	Client options file or with include.image option.

Examples

Task Back up the /home/test file space over which the logical volume is mounted and perform an image incremental backup that backs up only new and changed files after the last full image backup.

```
dsmc backup image /home/test -mode=incremental
```

Task Perform a static image backup of the logical volume that is mounted at the /home directory.

```
dsmc backup image /home -snapshotproviderimage=none
```

Task Perform a dynamic image backup of the logical volume that is mounted at the /home directory.

Command: dsmc backup image /home -dynamicimage=yes

Task Perform a snapshot image backup of the /home directory.

```
AIX client: dsmc backup image /home
             -snapshotproviderimage=JFS2
Linux client: dsmc backup image /home
              -snapshotproviderimage=LINUX_LVM
```

Task Back up the /dev/lv01 raw logical volume.

```
dsmc backup image /dev/lv01
```

Related information

"Imagegapsize" on page 417

"Snapshotproviderimage" on page 527

"Snapshotcachesize" on page 525

"Mode" on page 454

“Comparing methods 1 and 2” on page 186 To decide which method is appropriate for your environment.

Static, dynamic, and snapshot image backup

The traditional image backup prevents write access to the volume by other system applications during the operation.

Use the `dynamicimage` option to back up the volume as is without remounting it read-only. Corruption of the backup can occur if applications write to the volume while the backup is in progress. In this case, run `fsck` after a restore.

The `dynamicimage` option is not supported for JFS2 volumes.

For Linux x86_64 clients only: By default, Tivoli Storage Manager performs a snapshot image backup of file systems residing on a logical volume created by the Linux Logical Volume Manager during which the volume is available to other system applications. Snapshot image backup requires a Version 5.1 Tivoli Storage Manager server.

For AIX clients only: By default, Tivoli Storage Manager performs a snapshot image backup of JFS2 volumes during which the volume is available to other system applications. AIX allows the creation of a snapshot of a JFS2 volume while it is still online. The snapshot is created inside the same volume group as the source volume. You must ensure that the volume group provides enough free disk space to create the snapshot. The snapshot contains the old data blocks while the modified data is stored in the source volume. Use the `snapshotcachesize` option with the `backup image` command, in the `dsm.sys` file, or with the `include.image` option to specify an appropriate snapshot size so that all old data blocks can be stored while the image backup occurs.

The Linux Logical Volume Manager allows the creation of a snapshot of a logical volume while the logical volume itself is still online. The snapshot is created inside the same volume group as the source logical volume. You must ensure that the volume group provides enough free disk space to create the snapshot. The snapshot contains the old data blocks while the modified data is stored in the source logical volume. Use the `snapshotcachesize` option with the **backup image** command, in the `dsm.sys` file, or with the `include.image` option to specify an appropriate snapshot size so that all old data blocks can be stored while the image backup occurs. A snapshot size of 100 percent will ensure a valid snapshot.

Utilizing image backup to perform file system incremental backup

There are two methods of utilizing image backups to perform efficient incremental backups of your file system. These backup methods allow you to perform point-in-time restore of your file systems and improve backup and restore performance.

You can perform the backup only on formatted volumes; not on raw logical volumes. You can either use *image backup with file system incremental* or you can use *image backup with image incremental mode* to perform image backups of volumes with mounted file systems.

The following are some examples of using *image backup with file system incremental*.

- To perform a full incremental backup of the file system: `dsmc incremental /myfilesystem`
- To perform an image backup of the same file system: `dsmc backup image /myfilesystem`
- To periodically perform incremental backups: `dsmc incremental /myfilesystem`

You must follow the next steps in the order shown to ensure that the server records additions and deletions accurately.

Use this command to restore the file system to its exact state as of the last incremental backup: `dsmc restore image /myfilesystem -incremental -deletefiles`.

During the restore, the client does the following:

- Restores the most recent image on the server.
- Deletes all of the files restored in the previous step which are inactive on the server. These are files which existed at the time of the image backup, but were subsequently deleted and recorded by a later incremental backup.
- Restores new and changed files from the incremental backups.

If you do not follow the steps exactly, two things can occur:

1. After the original image is restored, all files backed up with the **incremental** command are restored individually.
2. If you perform a **backup image** before performing an **incremental**, files deleted from the original image are *not* deleted from the final restored file system.

The following are some examples of using *image backup with image incremental mode*.

- To perform an image backup of the same file system: `dsmc backup image /myfilesystem`
- To perform an incremental image backup of the file system: `dsmc backup image /myfilesystem -mode=incremental`

This sends only those files that were added or changed since the last image backup to the server.

- To periodically perform full image backups: `dsmc backup image /myfilesystem`
- To restore the image: `dsmc restore image /myfilesystem -incremental`

On restore, the backup-archive client ignores the `deletefiles` option when the `image+image incremental` technique of backing up has been used. The restore will include files that were deleted after the last full image backup plus the latest versions of files added or changed after the last image backup.

Note: You should perform full image backups periodically in the following cases. This will improve restore time because fewer changes are applied from incrementals.

- When a file system changes substantially (more than 40%).
- Once each month.
- As appropriate for your environment.

The following restrictions apply when using the image backup with image incremental mode:

- The file system can have no previous full incremental backups produced by the **incremental** command.
- Incremental-by-date image backup does not inactivate files on the server; therefore, when files are restored, none can be deleted.

- If this is the first image backup for the file system, a full image backup is performed.
- Using `mode=incremental` backs up only files with a changed date, not files with changed permissions.
- If file systems are running at or near capacity, an out-of-space condition could result during the restore.

Backup NAS

The **backup nas** command creates an image backup of one or more file systems that belong to a Network Attached Storage (NAS) file server, otherwise known as NDMP Backup. You are prompted for the Tivoli Storage Manager administrator ID.

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. The NAS node name identifies the NAS file server to the Tivoli Storage Manager server; the NAS node name must be registered at the server. Place the `nasnodename` option in your client options file (`dsm.opt`). The value in the client options file is the default, but can be overridden on the command line.

Use the `toc` option with the **backup nas** command or the `include.fs.nas` option to specify whether the Tivoli Storage Manager server saves Table of Contents (TOC) information for each file system backup. If you save TOC information, you can use the **QUERY TOC** server command to determine the contents of a file system backup with the **RESTORE NODE** server command to restore individual files or directory trees.

You can also use the Tivoli Storage Manager web client to examine the entire file system tree and select files and directories to restore. Creation of a TOC requires that you define the **tocdestination** attribute in the backup copy group for the management class to which this backup image is bound. TOC creation requires more processing, network resources, storage pool space, and possibly a mount point during the backup operation. If you do not save TOC information, you can still restore individual files or directory trees using the **RESTORE NODE** server command, if you know the fully qualified name of each file or directory and the image in which that object was backed up.

The `toc` option is only supported for images that are backed up by Version 5.2 or later client and server.

Specifying `mode=differential` on the **BACKUP NODE** server command or the **backup nas** command where no full image exists, shows that a full backup was started. Using the **QUERY PROCESS** server command shows that a full backup is in process.

Use the `mode` option to specify whether to perform a full or differential NAS image backup. A full image backup backs up the entire file system. The default is a differential NAS image backup on files that change after the last full image backup. If an eligible full image backup does not exist, a full image backup is performed. If a full image exists, whether it is restorable, or expired and being maintained because of dependent differential images, specifying `mode=differential` sends a differential image backup. If a full image is sent during a differential backup, it is reflected as a full image using the **QUERY NASBACKUP** server

command. The **QUERY NASBACKUP** server command also displays NAS images that are restorable and displays full image or differential image as the object type.

Use the **monitor** option to specify whether you want to monitor a NAS file system image backup and display processing information on your screen.

Use the **monitor process** command to display a list of all processes for which an administrative user ID has authority. The administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web.


Use the **cancel process** command to stop NAS backup processing.

Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.

Supported Clients

This command is valid for AIX, and Solaris clients only.

Syntax



Parameters

filespec
Specifies the name of one or more file systems on the NAS file server. If you do not specify this parameter, the backup-archive client processes all of the file systems that are defined by the domain.nas option.

If you do not specify the *filespec* or the domain.nas option, the default **all-nas** value is used for domain.nas and all file systems on the NAS file server are backed up.

Table 85. Backup NAS command: Related options

Option	Where to use
mode "Mode" on page 454	Command line only.
monitor "Monitor" on page 457	Command line only.
nasnodename "Nasnodename" on page 460	Client options file (dsm.sys) or command line.
toc "Toc" on page 552	Command line or with the include.fs.nas option in your client options file (dsm.sys).

Examples

Task Perform the NAS image backup of the entire file system.

Command: backup nas -mode=full -nasnodename=nas1 /vol/vol0 /vol/vol2

Task Perform the NAS image backup of the entire file server.

Command: backup nas -nasnodename=nas1

Task Perform the NAS image backup of the entire file system and save Table of Contents (TOC) information for the file system backup.

Command: backup nas -mode=full -nasnodename=netappsj /vol/vol0
-toc=yes

Related information

"Nasnodename" on page 460

"Toc" on page 552

"Mode" on page 454

"Monitor" on page 457

"Cancel Process" on page 641

"Domain.nas" on page 366

Backup VAPP

Use the **backup vapp** command to back up a single VMware vApp, a list of vApps, or all vApps within an organization virtual data center.

To use this command, you must have a license to use Tivoli Storage Manager for Virtual Environments V7.1, or later. Before you can use this command, the Tivoli Storage Manager nodes must be configured for the vCloud environment. In the Data Protection for VMware vCloud GUI, go to the Configuration window and click **Run Configuration Wizard**.

Supported Clients

This command is valid for Linux clients that are configured to perform vCloud Director vApp backups.

Syntax

```
➤--Backup VAPP-- --vapp_specification-- --ASNODENAME=--provider_vdc_node--➤
└--VCDHost=--vCloud_server-- --VCDUser=--vCloud_administrator-- --VCDPw=--password--┐
└--VMCHost=--hostname-- --VMCUser=--username-- --VMCPw=--password--┐└--PREView-- --options--┐➤
```

Parameters

vapp_specification

Required parameter. Specifies the vApps that you want to back up. You can specify a single vApp, a list of vApps, or back up all vApps that are in an organization virtual data center. If you do not specify any vApps on this command, the vApps are selected from the DOMAIN.VCD client option.

The vApp specification must include an organization name and an organization virtual data center name. The following parameters identify the vApps to back up:

ORG=*organization_name*

Specifies the name of the organization that the organization virtual data center belongs to.

ORGVDC=*org_vdc_name*

Specifies the name of the organization virtual data center that contains the vApps to protect.

VAPP=*vapp_name*

Specifies one or more vApps to protect. If you omit vApp names, all vApps in the specified organization and organization virtual data center are processed.

Wildcards and some special characters cannot be included in vApp name on this command. Specifically, you cannot include any of the following characters in a vApp name: " ' : ; * ? , < > / \ |.

You must specify these keyword parameters in the following order: **ORG**=, **ORGVDC**=, **VAPP**=. The following syntax rules also apply:

- If any value contains a space character, enclose the entire specification in quotation marks.
- To specify more than one organization virtual data center name that is under the same organization, separate the organization virtual data center names with commas and do not insert a space character after the commas.

"org=organization_name,orgvdc=org_vdc_name1,org_vdc_name2"

When you specify multiple organization virtual data centers on the same statement, you cannot specify vApp names.

- To specify multiple vApp names within the same organization virtual data center, separate the vApp names with commas. Do not insert a space character after the commas.

"org=organization_name_name,orgvdc=org_vdc_name,
vapp=vapp_name1,vapp_name2,vapp_name3"

- To specify organization virtual data centers from multiple organizations, separate the **org**= and **orgvdc**= pairs with semicolons. Do not insert a space character after the semicolons.

"org=organization_name1,orgvdc=org_vdc_name1;
organization=organization_name2,orgvdc=org_vdc_name2"

-ASNODENAME=*provider_vdc_node*

Required parameter. Specifies the node name of the client node that is associated with a provider virtual data center. This parameter can be set on the this command or in the client options file (dsm.opt). However, if you have backup schedules that are associated with the data move node, you must specify this option on the command or in the schedule definition, and not in the client options file.

-VCDHost=*vCloud_server*

-VCDUser=*vCloud_administrator*

-VCDPw=*password*

Optional parameters for this command. These three parameters specify the following information:

- The host name or IP address of the vCloud Director server (VCDHost).

- The account name for a user on that server (VCDUser).
- The password that is associated with the user account (VCDPw).

These options are typically not specified on the command line. Instead, the server address is specified on the VCDHost option in the client options file. The administrator name and password are then associated with that host by a **set password** command, where you specify that host, the administrator name, and the administrator password. You also specify TYPE=VCD on the **set password** command to indicate that these credentials are for logging on to a vCloud server.

You can specify these parameters on the command line. However, passing these parameters on the command line does supply your login credentials in plain text. Adding these parameters on the command line should be done only as a temporary measure when you are troubleshooting logins.

-VMCHost=hostname
-VMCUser=username
-VMCPw=password

Optional parameters for this command. These three parameters specify the following information:

- The host name or IP address of the vCenter or ESXi server that hosts the vApps that you want to protect (VMCHost).
- The account name for a user on that server (VMCUser).
- The password that is associated with the user account (VMCPw).

These options are typically not specified on the command line. Instead, the server address is specified on the VMCHost option in the client options file. The administrator name and password are then associated with that host by a **set password** command, where you specify that host, the administrator name, and the administrator password. You also specify TYPE=VM on the **set password** command to indicate that these credentials are for logging on to a vCenter or ESXi server.

You can specify these parameters on the command line. However, passing these parameters on the command line does supply your login credentials in plain text. Adding these parameters on the command line should be done only as a temporary measure when you are troubleshooting logins.

-PREView

Optional. This parameter can be specified only on the command line for the **Backup VAPP** command. It creates a list of vApps, including their VMs, that are included in the backup operation when you run the **Backup VAPP** command without this parameter. Add this parameter to determine whether the vApp specification or the DOMAIN.VCD statement includes the vApps that you want protected.

Related options

Table 86. Backup VAPP command: Related options

Option	Where to use
asnodename	Client user-options file (dsm.opt) or client systems option file (dsm.sys), on the command line, or on the General tab in the Preferences editor.
include.vapp	Client options file.

Table 86. Backup VAPP command: Related options (continued)

Option	Where to use
mode	Client user-options file (dsm.opt) or client systems option file (dsm.sys), or on the command line. Specify IFFull or IFIncremental. The default for vApp backups is IFIncremental.
vappmc	Client user-options file (dsm.opt) or client systems option file (dsm.sys), or on the command line. This option can also be included on the server in a client option set.
vcdpw	Required parameter. Specifies the password for the vCloud administrator that is specified on -VCDUser. Specify this option on the set password command and specify TYPE=VCD.
vmbackuptype	Client user-options file (dsm.opt) or client systems option file (dsm.sys), or on the command line. This option can also be included on the server in a client option set. This option defaults to the correct value (FULLVM). If you specify any other value, your setting is ignored and the default is assumed.
vmchost	Client user-options file (dsm.opt) or client systems option file (dsm.sys), or on the command line. This option can also be included on the server in a client option set.
vmcuser	Client user-options file (dsm.opt) or client systems option file (dsm.sys), or on the command line. You can also set this option on a set password command where TYPE=VM is specified. This option can also be included on the server in a client option set.
vmcpw	Required parameter. Specify this option on a set password command and specify TYPE=VM This option can also be included on the server in a client option set.
vmlimitperdatastore	Client user-options file (dsm.opt) or client systems option file (dsm.sys), or on the command line. This option can also be included on the server in a client option set.
vmlimitperhost	Client user-options file (dsm.opt) or client systems option file (dsm.sys), or on the command line. This option can also be included on the server in a client option set.
vmmaxparallel	Client user-options file (dsm.opt) or client systems option file (dsm.sys), or on the command line. This option can also be included on the server in a client option set.

Table 86. Backup VAPP command: Related options (continued)

Option	Where to use
vmctlmc	Client user-options file (dsm.opt) or client systems option file (dsm.sys), or on the command line. This option can also be included on the server in a client option set.
vmvstortransport	Client user-options file (dsm.opt) or client systems option file (dsm.sys), or on the command line. This option can also be included on the server in a client option set.

Example commands for backing up vApps

Back up a vApp that is named master_vApp2. The vApp has two virtual machines: VM1 and VM2. The vApp is in the Orion organization, in the virtual data center that is named Orion_Silver. This is the first time that the vApp is backed up, therefore all used disk blocks are included in the backup.

```
dsmc backup vApp "org=Orion,orgvdc=Orion_Silver,
  vapps=master_vApp2" -asnodename=PVDC_NODE_For_Orion_Silver
IBM Tivoli Storage Manager
Command Line Backup-Archive Client Interface
  Client Version 7, Release 1, Level 0.0
  Client date/time: 11/14/2013 13:22:52
(c) Copyright by IBM Corporation and other(s) 1990, 2013. All Rights Reserved.
```

```
Node Name: DM_pvdc1_DM1
Session established with server VM-08R2-TSM62-6_SERVER2: Windows
  Server Version 6, Release 3, Level 0.0
  Server date/time: 11/14/2013 14:22:52  Last access: 11/14/2013 14:22:39
```

```
Accessing as node: OVDC_ORG1_ORGVDC2
Full BACKUP VAPP of vApp 'master_vApp2'.
Organization: Haifa
Organization vDC: Haifa_Silver
```

```
Backup VAPP command started.  Total number of vApps to process: 1
Total number of virtual machines to process: 2
```

```
Starting Full VAPP backup of VMware vApp 'master_vApp2'
```

```
Putting vApp into maintenance mode.
Backing up vApp configuration data.
```

```
Backing up VM1...
<normal output suppressed for this example>
```

```
Backing up VM2 ...
<normal output suppressed for this example>
```

```
... Removing maintenance mode from vApp.
Successful Full VAPP backup of VMware vApp 'master_vApp2'
  mode: 'Periodic Full - Full'
  target node name: 'OVDC_ORG1_ORGVDC2'
  data mover node name: 'DM_pvdc1_DM1'
  VMs: VM1, VM2
```

```
Statistics for vApp 'master_vApp2'.
```

```
Total number of objects inspected:      1
Total number of objects backed up:      1
```

```

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 inspected:          4.50 GB
Total number of bytes transferred:        4.50 GB
Data transfer time:                       681.06 sec
Network data transfer rate:               6,928.30 KB/sec
Aggregate data transfer rate:             6,298.05 KB/sec
Objects compressed by:                   0%
Total data reduction ratio:               0.00%
Subfile objects reduced by:               0%
Elapsed processing time:                   00:12:29

```

```

Backup VAPP command complete
Total number of vApps backed up successfully: 1
Total number of virtual machines backed up successfully: 2
Total number of vApps failed: 0
Total number of vApps processed: 1

```

The following example shows the **-PREVIEW** option output. The output shows that the vApp that is named `golden_vApp3` has three virtual machines, which are named VM1, VM2, and VM3.

```

dsmc backup vApp "org=Orion,orgvdc=Orion_Silver,
vapps=golden_vApp3" -preview

```

```

IBM Tivoli Storage Manager
Command Line Backup-Archive Client Interface
  Client Version 7, Release 1, Level 0.0
  Client date/time: 11/14/2013 13:52:00
(c) Copyright by IBM Corporation and other(s) 1990, 2013. All Rights Reserved.

```

```

Node Name: DM_pvdc1_DM1
Session established with server VM-08R2-TSM62-6_SERVER2: Windows
  Server Version 6, Release 3, Level 0.0
  Server date/time: 11/14/2013 14:52:00  Last access: 11/14/2013 14:23:06

```

```

Accessing as node: OVDC_ORG1_ORGVDC2
Full BACKUP VAP of vApp 'golden_vApp3'.
Organization: Orion
Organization vDC: Orion_Silver

```

```

1. vAppName: golden_vApp3
   DomainKeyword: vapp= golden_vApp3
   VMNAME[1]: VM1
   VMNAME[2]: VM2
   VMNAME[3]: VM3

```

```

Total number of vApps processed: 1
Accessing as node: OVDC_Orion_Orion_Silver

```

To back up vApps from Org1 and Org1_vdc, use the following command:

```

dsmc backup vapp -domain.vdc="org=Org1,orgvdc=Org1_vdc"

```

To back up all vApps in multiple organizations and organization virtual data centers, specify each vApp specification, separated by a semicolon.

```

dsmc backup vapp -domain.vcd="org=Haifa,orgvdc=vdc1;
org=Mainz,orgvdc=vdc2;org=SanJose,orgvdc=vdc4"

```

Related reference:

“Asnodename” on page 319

“Mode” on page 454

"Vmbackuptype" on page 570
"Vmchost" on page 570
"Vmcpw" on page 571
"Vmcuser" on page 573
"Vmlimitperdatastore" on page 579
"Vmlimitperhost" on page 580
"Vmmaxparallel" on page 581
"Vmmc" on page 584
"Vmctlmc" on page 572
"Vmvstortransport" on page 597

Backup VM

Use the **backup vm** command to start a full backup of a virtual machine.

Backing up VMware virtual machines

Use the **backup vm** command to back up VMware virtual machines.

One or more virtual machines are backed up by the Tivoli Storage Manager data mover node. *Data mover node* is the name that is given to a configuration where the backup-archive client runs on a vStorage backup server and is configured to protect the virtual machines in a Virtual Center or ESX/ESXi server. You must configure the VMware virtual machine before you use this command. For information about configuring the VMware virtual machine, see "Preparing the environment for full backups of VMware virtual machines" on page 203.

A full VM backup stores a backup copy of all virtual disk images and configuration information for a virtual machine. Full VM backups enable a complete restore of a virtual machine, but they take more time and more server space than a file-level or incremental backup.

If you set `vmenabletemplatebackups` option to **yes**, a **backup vm** operation includes the template VMs, but only if the vStorage backup server is connected to a vCenter Server, and not to an ESX or ESXi host.

If a snapshot fails during backup processing, the client attempts to back up the VMware virtual machine one more time. To control the number of total snapshot attempts, set the `INCLUDE.VMSNAPSHOTATTEMPTS` option in the client options file.

Virtual machines that are deployed in vApps are included in **backup vm** operations.

Backup management tags are used to exclude virtual machines from backups and to specify the retention policy of backups. The tags and categories are created when you use one of the following methods:

- Enable tagging support on the data mover node with the `vmtagdatamover` option and run the **backup vm** command.
- Use the IBM Data Protection extension to manage Tivoli Storage Manager backups.
- Run the **set vmtags** command on any data mover node.

When the `vmtagdatamover` option is set to *yes*, all tags that are assigned to a virtual machine are included in **backup vm** operations

A Full VM backup uses VMware Changed Block Tracking (CBT) to create content-aware (used-block only) backups. The client enables changed block tracking (CBT) on an ESX or ESXi server when a backup begins. VMware CBT requires an ESX 4.1 (or later) host, with virtual hardware 7 (or later). You cannot perform incremental or full VM content-aware backups on virtual machines that do not support CBT.

When CBT is enabled, it tracks disk changes when I/O operations are processed by the ESX or ESXi server storage stack on the following disks:

- A virtual disk that is stored on VMFS; the disk can be an iSCSI disk, a local disk, or a disk that is on a SAN.
- A virtual disk that is stored on NFS.
- An RDM that is in virtual compatibility mode.

When I/O operations are not processed by the ESX or ESXi storage stack, changed block tracking cannot be used to track disk changes. The following disks cannot use CBT:

- An RDM that is in physical compatibility mode.
- A disk that is accessed directly from inside a VM. For example, vSphere cannot track changes that are made to an iSCSI LUN that is accessed by an iSCSI initiator in the virtual machine.

Complete information about changed block tracking requirements is described in the *VMware Virtual Disk API Programming Guide* in the VMware product documentation. In the guide, search for “Low Level Backup Procedures” and read the “Changed Block Tracking on Virtual Disks” section.

For VMware servers that do not support CBT, both the used and the unused areas of the disk are backed up and an informational message is logged in the `dsmerror.log` file. Use the `-preview` option on the **backup vm** command to view the current CBT status. CBT status has three values:

Off Indicates the CBT configuration parameter (**ctkEnabled**) is not enabled in the virtual machine's configuration parameters. **Off** is the default state.

Not Supported

Indicates that the virtual machine does not support CBT. Changed-block only backups are not possible.

On Indicates the virtual machine supports CBT and that CBT is enabled in the virtual machine's configuration parameters (`ctkEnabled=true`).

The client turns on CBT (it sets `ctkEnable=true`) with each backup attempt. After the client turns on CBT, it remains on, even if the virtual machine is deleted from the Tivoli Storage Manager server. With CBT enabled, after the first full VM backup is performed, only the changed blocks on the disk are backed up or restored.

If you are no longer performing Tivoli Storage Manager backups of a virtual machine, you can turn off CBT. To turn off CBT, right-click the virtual machine that you want to turn off CBT for in the vSphere client. Click **Edit Settings > Options > General > Configuration Parameters**. Then, set the **ctkEnabled** configuration parameter to false.

Tip: You can use the compression option with backups only if the backup is being saved to a storage pool that was enabled for client-side deduplication.

You specify the `-vmbackuptype` and `-mode` options to indicate how the backups are to be performed. For full VM backups, use `-vmbackuptype=fullvm`, and specify any of the following mode options:

Full Full mode. In this mode, an image backup is created of all objects on a virtual machine's disks. Contrast mode=`full` with mode=`incremental`, which creates a snapshot that contains only the data that changed since the last backup. The incremental-forever backup methods provide the most efficient backup operations. If you are not licensed to use Tivoli Storage Manager for Virtual Environments, you must specify a full VM backup.

Incremental

Incremental mode. In this mode, an image backup is created of the objects that changed since the last backup. To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

IFFull Incremental-forever-full mode. In this mode, a snapshot of all used blocks on a virtual machine's disks are backed up to the server. You must be licensed to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware, or Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V to use this option.

IFIncremental


Incremental-forever-incremental. In this mode, a snapshot is created of the blocks that changed since the last backup. You must be licensed to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware, or Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V to use this option.

For information about the incremental-forever backup strategy, see IBM Tivoli Storage Manager for Virtual Environments, Data Protection for VMware: Backup and restore types .

Supported Clients

This command is valid only on supported Linux clients that are installed on a vStorage backup server that protects VMware virtual machines.

Syntax



Parameters

vmname

Specify the name of one or more virtual machines that you want to back up. The name is the virtual machine display name. Separate multiple virtual machine names with commas. If you set the `vmenabletemplatebackups` option to **yes**, *vmname* can specify the name of a template VM to back up.

VMware vCenter allows for two or more virtual machines to have the same display name. However, the backup-archive client requires that all virtual machine names in a vCenter server configuration be unique. To prevent errors during processing, ensure all virtual machines have a unique display name.

Wildcard characters can be used in virtual machine names that are specified as this parameter. However, wildcard processing differs, depending on which backup mode is used.

- For backups that use `mode=iffull` or `mode=ifincremental`, wildcards can be used to match VM name patterns. For example:
 - `backup vm VM_TEST*` includes all virtual machines that have names that begin with `VM_TEST`
 - `backup vm VM??` includes any virtual machine that has a name that begins with the letters “VM”, followed by 2 characters
- For backups that use `mode=full`, wildcard characters cannot be included in a VM name. However, if you omit the VM name from the command and instead, you specify an asterisk (*), that syntax begins a full image backup of all Hyper-V virtual machines that are defined to the Hyper-V host.

If you do not specify *vmname*, you can identify the virtual machine with one of the following options:

- `domain.vmfull`

:vmdk=disk_label

This keyword is an extension to the *vmname*. It specifies the label (name) of the virtual machine disk to include in the backup operation. You can exclude a disk by preceding the keyword with the exclusion operator (-). For more ways to include or exclude disks from processing, see “Domain.vmfull” on page 369, “Exclude.vmdisk” on page 394, “Include.vmdisk” on page 431.

-VMBACKUPUPDATEGUID

To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

This option updates the globally unique identifier (GUID) for the virtual machine that you are backing up. This parameter is intended for use only in the following scenario:

You want to restore a previously backed up virtual machine named ORION. But, before you shut down and replace the copy of ORION that is running in your production environment, you want to verify the configuration of the restored virtual machine before you use it to replace the existing ORION.

1. You restore the ORION virtual machine and give it a new name: `dsmc restore vm Orion -vmname=Orion2`

2. You update and verify the ORION2 virtual machine and determine that it is ready to replace the existing virtual machine that is named ORION.
3. You power down and delete ORION.
4. You rename ORION2 so that it is now named ORION.
5. The next time that you backup ORION, by using either an incremental-forever full, or incremental-forever-incremental backup, you add the **-VMBACKUPUPDATEGUID** parameter to the **backup vm** command. This option updates the GUID, on the Tivoli Storage Manager server, so the new GUID is associated with the stored backups for the ORION virtual machine. The chain of incremental backups is preserved; there is no need to delete existing backups and replace them with new backups.

-PREVIEW

This option displays information about a virtual machine, including the labels of the hard disks in the virtual machine, and the management class information for a virtual machine.

You can use the disk labels with the `:vmdk=` or `:-vmdk=` keywords to include or exclude disks from a backup operation. The following text is sample output from the **-preview** parameter:

```
backup vm vm1 -preview
Full BACKUP VM of virtual machines 'VM1'

vmName:vm1
VMDK[1]Label: Hard disk 1
VMDK[1]Name: [ds5k_svt_1] tsmcetlnx14/tsmcetlnx14.vmdk
VMDK[1]Status: Included
VMDK[2]Label: Hard disk 2
VMDK[2]Name: [ds5k_svt_1] tsmcetlnx14/tsmcetlnx14_1.vmdk
VMDK[2]Status: Excluded - user,Independent,pRDM
```

This example output from **-preview** shows that VMDK 2 was excluded by the previous backup. Disks that were included in a backup have a status of **Included**. Disks that were excluded from the backup have a status of **Excluded**, followed by a reason code. The reason codes can be any of the following:

user Indicates that the disk was skipped because it was excluded on a `domain.vmfull` statement, on the command line, or in the client options file.

Independent

Indicates that the disk is an independent disk. Independent disks cannot be part of a snapshot, so they are excluded from **backup vm** operations. Ensure that the `vmprocessvmwithindependent` option is set to yes or the entire virtual machine is bypassed by a backup operation if it contains one or more independent disks.

pRDM

Indicates that the disk is a physical Raw Device Mapped (pRDM) disk. pRDM disks cannot be part of a snapshot, so they are excluded from **backup vm** operations. Ensure that the `vmprocessvmwithprdm` option is set to yes or the entire virtual machine is bypassed by a backup operation if it contains one or more raw device mapping (RDM) volumes that are provisioned in physical-compatibility mode (pRDM).

The output from the **-preview** parameter also shows the management class name that is associated with the virtual machine, along with information

about where the management class was set. This information can help you verify whether the domain and tag values are set correctly for the management class. For example:

```
backup vm -preview
```

Full BACKUP VM of virtual machines specified in DOMAIN.VMFULL option.

```
1. vmName: tag_vm_2
   DomainKeyword: all-vm
   toolsRunningStatus: guestToolsNotRunning
   toolsVersionStatus: guestToolsNotInstalled
   consolidationNeeded: No
   Change Block Tracking: On
   managementClassName: STANDARD
   managementClassLocation: Node Default

   VMDK[1]Label: 'Hard disk 1' (Hard Disk 1)
   VMDK[1]Name: '[Raid1-lannds2] tag_vm_2/tag_vm_2.vmdk'
   VMDK[1]Status: Included
...

12. vmName: vm-jean
   DomainKeyword: all-vm
   toolsRunningStatus: guestToolsNotRunning
   toolsVersionStatus: guestToolsNotInstalled
   consolidationNeeded: No
   Change Block Tracking: On
   managementClassName: MGMTCLASS1 (invalid)
   managementClassLocation: VM Tag Management Class (IBM Spectrum Protect)

   VMDK[1]Label: 'Hard disk 1' (Hard Disk 1)
   VMDK[1]Name: '[Raid1-lannds2] vm-jean/vm-jean.vmdk'
   VMDK[1]Status: Included
```

where:

managementClassName

Displays the name of the management class that the virtual machine is bound to.

If the "(invalid)" label is shown next to the management class name, either the name was incorrectly specified, the management class was removed on the Tivoli Storage Manager server, or no backup copy group was found in the management class on the server. When the management class name is invalid, the virtual machine backup operation fails.

managementClassLocation

Displays where the management class was set. The following locations are possible:

Node Default

The management class is set on the default domain of the VMware datacenter node.

VMMC option

The management class is set with the `vmmc` option.

VMCTLMC option

The management class is set with the `vmctlmc` option.

VAPPMC option

The management class is set with the `vappmc` option.

INCLUDE.VM option

The management class is set with the `include.vm` option.

VM Tag Management Class (IBM Spectrum Protect)

The management class is set as a tag value of the

Management Class (IBM Spectrum Protect) tag category. Tag values can be set with backup management settings in the IBM Data Protection extension in the vSphere Web Client, or by using tools such as VMware vSphere PowerCLI version 5.5 R2 or later.

Important: In order to display the management class information that is set by tags, you must set the `vmtagdatamover yes` option in the client options file, or you must include the `-vmtagdatamover=yes` parameter when you run the `dsmc backup vm` command. If you did not set the `vmtagdatamover` option or if it is set to no, the client ignores any management class tag values, and displays the management class definition that is set in the default domain of the datacenter node, the `vmmc` option, or the `include.vm` option.

Return codes for virtual machine backup operations

Backup operations for virtual machines can complete with the return codes that are shown in the following table.

Return code	Description
0	A command to back up one or more virtual machines completed successfully.
8	A command to back up multiple virtual machines succeeded for only some of the virtual machines that were targeted by the command. Examine the log file to determine the processing status for each of the targeted virtual machines.
12	Indicates that either of the following error conditions occurred: - The backup command could not back up any of the virtual machines that were targets of the backup operation. - The backup command failed and it stopped before all virtual machines that were specified were inspected. Examine the log file to determine the reason for the failure.

vStorage API for data protection example commands

Perform an IFIncremental backup of two VMs named vm3 and vm4.

```
dsmc backup vm vm3,vm4 -vmbackuptype=fullvm -mode=ifincremental
```

Perform an IFFull backup of a VM named vm1.

```
dsmc backup vm vm1 -vmbackuptype=fullvm -mode=iffull
```

Perform an IFFull VM backup of a VM named vm1, but include only Hard Disk 1 in the backup operation.

```
dsmc backup vm "vm1:vmdk=Hard Disk 1" -vmbackuptype=fullvm -mode=iffull
```

Perform an incremental-forever backup of a virtual machine that is named vm1, but exclude Hard Disk 1 and Hard Disk 4 from the backup operation.

```
dsmc backup vm "vm1:-vmdk=Hard Disk 1:-vmdk=Hard Disk 4"
-vmbackuptype=fullvm -mode=iffull
```

Perform an incremental-forever-full backup of two virtual machines that are named vm1 and vm2. On vm1, back up only Hard Disk 2 and Hard Disk 3. On vm2, back up all virtual disks.

```
dsmc backup vm "vm1:vmdk=Hard Disk 2:vmdk=Hard Disk 3",  
vm2 -vmbackuptype=fullvm -mode=iffull
```

Perform parallel backups of the VMware virtual machines that are selected for backup by using the selection criteria (domain parameters) on the `domain.vmfull` statement. Limit the number of parallel backups to 5 virtual machines, and prevent more than 2 of them, on any ESX or ESXi server, from being included in the parallel backup operation.

```
dsmc backup vm -vmmaxparallel=5 -vmlimitperhost=2 -vmlimitperdatastore=0
```

Related links for backing up VMware virtual machines

- “**Query VM**” on page 691
- “**Restore VM**” on page 715
- “`Domain.vmfull`” on page 369
- “`Include.vm`” on page 430
- “`Mbobjrefreshthresh`” on page 450
- “`Mbpctrefreshthresh`” on page 451
- “`Mode`” on page 454
- “`Vmbackdir`” on page 568
- “`Vappmc`” on page 561
- “`Vmbackuptype`” on page 570
- “`Vmchost`” on page 570
- “`Vmctlmc`” on page 572
- “`Vmcpw`” on page 571
- “`Vmcuser`” on page 573
- “`Vmlimitperdatastore`” on page 579
- “`Vmlimitperhost`” on page 580
- “`Vmmc`” on page 584
- “`Vmmaxparallel`” on page 581
- “`Vmtagdatamover`” on page 591
- “**Set Vmtags**” on page 740
- “Virtual machine exclude options” on page 394
- “Virtual machine include options” on page 428

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. You are prompted for the Tivoli Storage Manager administrator ID.

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.

Supported Clients

This command is valid for AIX, Linux, and Solaris clients only.

Syntax

►►—Cancel Process—◄◄

Parameters

There are no parameters for this command.

Examples

Task Cancel current NAS image backup or restore processes.

Command: `cancel process`

Cancel Restore

The **cancel restore** command displays a list of your restartable restore sessions in the server database.

You can cancel only one restartable restore session at a time. Run the **cancel restore** command again to cancel more restores. To restart restartable restore sessions, use the **restart restore** command.

Use the **cancel restore** command under the following circumstances:

- You cannot back up files that are affected by the restartable restore.
- Restartable restore sessions lock the file space so that files cannot be moved off of the sequential volumes of the server.

Supported Clients

This command is valid for all clients.

Syntax

►►—Cancel Restore—◄◄

Parameters

There are no parameters for this command.

Examples

Task Cancel a restore operation.

`cancel restore`

Delete Access

The **delete access** command deletes authorization rules for files that are stored on the server.

When you delete an authorization rule, you revoke user access to any files or images that are specified by that rule.

Supported Clients

This command is valid for all clients.

Syntax

►►Delete—ACcess—►►

Parameters

There are no parameters for this command.

Examples

Task Display a list of current authorization rules and select the rules that you want to delete.

```
delete access
```

See the following screen example:

Index	Type	Node	Owner	Path
1	Backup	NODE1	USER1	home/dev/proja/list/
2	Archive	NODE3	LUIE	home/fin/budg/depta/
3	Backup	NODE4	USER2	home/plan/exp/deptc/
4	Archive	NODE5	USER2S	home/mfg/invn/parta/

Enter Index of rule(s) to delete, or quit to cancel:

To delete the authorization rules that allow luie and user2s to access your files or images, type 2 4 or 2,4 and press Enter.

Delete Archive

The **delete archive** command deletes archived files from Tivoli Storage Manager server storage. Your administrator must give you the authority to delete archived files.

Important: When you delete archived files, you cannot retrieve them. Verify that the files are obsolete before you delete them.

Supported Clients

This command is valid for all clients.

Syntax

►►Delete ARchive—
└─options┐ └─filespec┐
└─{—filespace—name—}—filespec┐

Parameters

filespec

Specifies the path and file name that you want to delete from storage. Use wildcard characters to specify a group of files or all files in a directory. You can also use the **filelist** option to process a list of files. The backup-archive client opens the file that you specify with this option and processes the list of files within according to the specific command.

Note: If you indicate *filespace*name, do not include a drive letter in the file specification.

{filespacename}

Specifies the file space (enclosed in braces) on the server that contains the file you want to delete. This is the name on the workstation drive from which the file was archived.

Use the *filespace*name if the name was changed, or if you are deleting files that are archived from another node with drive labels that are different from yours.

Table 87. Delete Archive command: Related options

Option	Where to use
dateformat "Dateformat" on page 344	Client options file (dsm.opt) or command line.
description "Description" on page 351	Command line only.
filelist "Filelist" on page 402	Command line only.
noprompt "Noprompt" on page 464	Command line only.
numberformat "Numberformat" on page 465	Client options file (dsm.opt) or command line.
pick "Pick" on page 472	Command line only.
subdir "Subdir" on page 539	Client options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 541	Client options file (dsm.opt) or command line.
timeformat "Timeformat" on page 549	Client options file (dsm.opt) or command line.

Examples

Task Delete a file that is named budget.

```
dsmc delete archive /user/home/proj1/budget
```

Task Delete all files that are archived from the /user/home/proj1 directory with a file extension of .txt.

```
dsmc del arch "/user/home/proj1/*.txt"
```

Task Delete files that are archived from the /user/project directory by 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.

```
dsmc delete archive "/user/project/*" -pick
```

Task Delete selected files from the group of files that are archived with the description "Monthly Budgets 2010" located in /user/projects and its subdirectories.

```
dsmc delete ar "/user/projects/*" -description="Monthly Budgets 2010" -pick -subdir=yes
```

Related information

Delete Backup

The **delete backup** command deletes files, images, and virtual machines that were backed up to Tivoli Storage Manager server storage. Your administrator must give you authority to delete objects.


When you delete files, the Tivoli Storage Manager server takes all of the backed up files that meet the *filespec* and *deltype* options that are specified and deactivates them. The server also assigns a deactivation date of *infinite-minus* so that the files are no longer available for restore and are purged, immediately on the subsequent run of file expiration. The file is not physically removed until the expiration process runs.

Important: After you delete backup files, you cannot restore them; verify that the backup files are no longer needed before you delete them. You are prompted to choose whether you want to continue with the delete. If you specify **yes**, the specified backup files are scheduled for deletion and removed from server storage.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

filespace/filespec

filespec

Specifies the path and file name that you want to delete from storage. To specify a file in another file space, precede the file name with the file space name. Use wildcard characters to specify a group of files or all files in a directory. Separate file specifications with a space. You can also use the `filelist` option to process a list of files. The backup-archive client opens the file that is specified with this option and processes the list of files within according to the specific command.

Note: If you indicate *filespacename*, do not include a drive letter in the file specification.

When you use `-deltype=inactive` or `-deltype=active`, use wildcard characters to specify a group of files or all files in a directory.

When you use `-deltype=all`, specify a fully wildcarded directory.

objtype

Specifies the type of object that you want to delete. You can specify either of the following values:

FILE

Specifies that you want to delete directories and files. This value is the default object type.

IMAGE

Specifies that you want to delete an image backup. Specifies that you want to delete an image backup. Objtype=image is not supported on Mac OS X.

VM *vmname*

Specifies that you want to delete one or more versions of a virtual machine backup; the virtual machine is identified by the *vmname* variable parameter. The virtual machine name cannot contain wildcard characters.

When objtype=VM is specified, the *filelist* option cannot be used. Specifying objtype=VM changes the behavior of the *-deltype* option. When objtype=vm is specified, you can use either *-deltype=active* or *-deltype=inactive*. You cannot use *-deltype=all*. Specifying *-deltype=inactive* displays a list of both inactive and active backups. You can use this list to specify which virtual machine backups that you want to delete. To delete only active virtual machine backups, use *-deltype=active*.

When you specify *-objtype=VM*, this command deletes only virtual machine backups that were created with any of the following modes: FULL, IFINCR, and IFFULL. Individual incremental backups (backups that were created by using *MODE=INCR*) that were created after a full backup was run cannot be deleted with this command. However, if you delete a full virtual machine image backup (created by using *MODE=FULL*), and if the server has any incremental backups (*MODE=INCR*) that were created for this VM after the full backup, then deleting the full VM backup also deletes the files that were created by a *MODE=INCR* backup.

If you delete an active backup for a virtual machine, the most recent inactive copy becomes the active backup. If you specify the *-pick* or *-inactive* option, only the backup that you specify is deleted. If you select a backup that is created by *MODE=IFINCR*, only the selected incremental backup is deleted; other incremental backups for the virtual machine are not deleted.

deltype

Specifies the deletion type. Specify one of the following values:

ACTIVE

Delete only active file objects. Directory objects are not deleted. This value is the default deletion type.

Note: If there are any inactive objects, then after the active object is deleted, the most current inactive object is changed from inactive to active.

To delete all versions of a file, first issue the **delete backup** command with *-deltype=inactive*, then enter 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 particular directory, including all subdirectories and their files.

Note: The parent directory of the deleted files and subdirectories is not deleted. If you specify `deltypes=ALL`, you cannot use the `pick` option because `deltypes=ALL` and the `pick` option are mutually exclusive.

Table 88. Delete Backup command: Related options

Option	Where to use
description "Description" on page 351	Command line only.
filelist "Filelist" on page 402	Command line only.
fromdate "Fromdate" on page 409	Command line, and in the GUI find function.
fromtime "Fromtime" on page 412	Command line, and in the GUI find function.
noprompt "Noprompt" on page 464	Command line only.
pick "Pick" on page 472	Command line only.
pitdate "Pitdate" on page 473	Command line, and in the GUI find function.
pittime "Pittime" on page 474	Command line, and in the GUI find function.
subdir "Subdir" on page 539	Client options file (<code>dsm.opt</code>) or command line.
tapeprompt "Tapeprompt" on page 541	Client options file (<code>dsm.opt</code>) or command line.
timeformat "Timeformat" on page 549	Client options file (<code>dsm.opt</code>) or command line.
today "Today" on page 554	Command line, and in the GUI find function.
totime "Totime" on page 554	Command line, and in the GUI find function.

Examples

Task Delete all active and inactive file objects that are named `budget` in directory `/data/plan/proj1`.

Commands:

```
delete backup /data/plan/proj1/budget.jan
  -deltypes=inactive
delete backup /data/plan/proj1/budget.jan
  -deltypes=active
```

Task Delete all inactive files that have a `.txt` extension that were backed up from the `/data/plan/proj1` directory and its subdirectories.

Command: `delete backup "/data/plan/proj1/*.txt" -deltypes=inactive -subdir=yes`

Task Delete selected active files that are backed up from the

/home/marymb/project directory. Use the -pick option to display a list of backup copies that match the file specification. From the list, you can select which versions to delete.

Command: delete backup "/home/marymb/project/*" -pick

Task Delete all active and inactive versions of files and subdirectories in the /home/storman/myproject directory. Then, delete all active and inactive versions of the /user/myproject directory.

Command:

```
delete backup "/home/storman/myproject*"
-deltype=all
```

Related reference:

"Filelist" on page 402

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 you backed up or archived.

You must be an authorized user to use this command.

Tivoli Storage Manager assigns a separate file space on the server for each workstation file system from which you back up or archive files. The file space name is the same as the file system name.

When you enter the **delete filesystem** command, a list of your file spaces is displayed. From this list, select the file space that you want to delete.

Your Tivoli Storage Manager administrator must give you authority to delete a file space. You need BACKDEL authority if the file space you want to delete contains backup versions, or ARCHDEL authority if the file space contains archive copies. If the file space contains both backup versions and archive copies, you need both types of authority.

Important: When you delete a file space, you delete all backup versions and archive copies within that file space. When you delete a file space, *you cannot restore the files*. Verify that the files are obsolete before you delete them.

You can use the **delete filesystem** command to interactively delete NAS file spaces from server storage. Use the nasnodename option to identify the NAS file server. Use the class option to specify the class of the file space to delete.

Supported Clients

This command is valid for all clients.

Syntax

►►Delete Filespace —options—►►

Parameters

Table 89. Delete Filespace command: Related options

Option	Where to use
class "Class" on page 333	Command line only.
detail "Detail" on page 353	Command line only.
nasnodename "Nasnodename" on page 460	Client system options file or command line.
scrolllines "Scrolllines" on page 511	Client system options file or command line.
scrollprompt "Scrollprompt" on page 512	Client system options file or command line.

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

Related information

"Nasnodename" on page 460

"Class" on page 333

Delete Group

Use the **delete group** command to delete a group backup on the Tivoli Storage Manager server.

After you delete 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 filesystem** command. No files are 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 filesystem** if you want to remove the file space and all the data it contains.

Note:

1. Use the **inactive** option to display both active and inactive group backup versions. By default, the client displays active versions.
2. Use the **pick** option to select a specific group to delete from the Tivoli Storage Manager server.
3. Use the **noprompt** option if you want to suppress the confirmation prompt that normally appears before you delete a group backup version. By default, the client prompts you for confirmation before you delete the group backup. Using this option can speed up the delete procedure. However, it also increases the danger of accidentally deleting a group backup version that you want to save. Use this option with caution.
4. Use the **query filesystem** command to display virtual file space names for your node that are stored on the server.

Supported Clients

This command is valid for all UNIX and Linux clients, except for Mac OS X.

Syntax

►►Delete GRoup— *filespec* — *options* —►►

Parameters

filespec

Specifies the virtual file space name and the group name that you want to delete from the server storage.

Table 90. Delete Group command: Related options

Option	Where to use
inactive “Inactive” on page 419	Command line only.
noprompt “Noprompt” on page 464	Command line only.
pick “Pick” on page 472	Command line only.
pitdate “Pitdate” on page 473	Command line only.
pittime “Pittime” on page 474	Command line only.

Examples

Task Delete the current active version of the /virtfs/group1 group.

Command:

```
delete group /virtfs/group1
```

Task Delete a backup version of the /virtfs/group1 group from a list of active and inactive versions.

Command:

```
delete group /virtfs/group1 -inactive -pick
```

Related information

“Inactive” on page 419

“Pick” on page 472

“Noprompt” on page 464

“Query Filespace” on page 674

Expire

The **expire** command deactivates the backup objects that you specify in the file specification or with the `filelist` option. You can specify an individual file to expire, or a file that contains a list of files to expire. If `OBJTYPE=VM`, this command deactivates the current backup for a virtual machine.

When you are 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 are not displayed in a subsequent query from the GUI. However, these files are displayed on the command line, if you specify the correct query with a wildcard character for the directory.

Note: Because the **expire** command changes the server picture of the client file system without changing the client file system, the **expire** command is not allowed on files that are on a file system that is monitored by the Tivoli Storage Manager journal daemon.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

OBJTYPE=FILE filespec

Specifies a path and a file name that you want to expire. You can enter only one file specification on this command. However, you can use wildcards to select a group of files or all the files in a directory. If you specify the `filelist` option, the `filespec` designation is ignored.

OBJTYPE=VM vmname

`vmname` specifies the name of a virtual machine. The active backup for the specified virtual machine is expired. The virtual machine name cannot contain wildcard characters.

When `objtype=VM` is specified, the expire command expires only full virtual machine backups (`MODE=FULL` or `MODE=IFFULL`) for the virtual machine that is specified on the `vmname` parameter.

Table 91. Expire command: Related options

Option	Where to use
dateformat "Dateformat" on page 344	Client user-options file (dsm.opt) or command line.
filelist "Filelist" on page 402	Command line only.
noprompt "Noprompt" on page 464	Command line only.
numberformat "Numberformat" on page 465	Client user-options file (dsm.opt) or command line.
pick "Pick" on page 472	Command line only.
timeformat "Timeformat" on page 549	Client user-options file (dsm.opt) or command line.

Examples

Task Deactivate the letter1.txt file in the home directory.

Command: expire "/home/letter1.txt"

Task Deactivate all files in the /admin/mydir directory.

Command: expire /admin/mydir/*

Task Deactivate all files that are named in the /home/avi/filelist.txt file.

Command: expire -filelist=/home/avi/filelist.txt

Help


Use the **help** command to display information about commands, options, and messages.

Tip: If you use the **help** command on the initial command line, no server contact is made and no password is needed.

Supported Clients

This command is valid for all clients.

Syntax



Entering the **help** command with no arguments causes help to display the complete table of contents. Either with the initial command or when HELP displays a prompt, you can enter the following parameters.

Parameters

command-name [*subcommand-name*]

Specifies a command name and, optionally, a subcommand name or their abbreviation, for example: **backup image**, or **b i**. In that case, the combination must be unique. Non-unique abbreviations result in the display of the first section of the entire help file that matches the abbreviation. This parameter is optional.

option-name

Specifies the name of an option, for example: **domain** or **do**. This parameter is optional.

TOC-section-number

Specifies a table of contents section number, for example: 1.5.3. This parameter is optional.

[*ANS*]*message-number*

Specifies a message number with or without its prefix, for example: **ans1036** or **1036**. This parameter is optional. The severity code is never necessary. Entering **ans1036E** results in a not-found response.

Important: If you enter arguments that do not fit these descriptions, you may get unexpected results (or no results) to be displayed. If you enter more than two arguments, your help request is rejected. Where a command name and an option name are the same, for example: **incremental** (command) and **incremental** (option), you can get help on the option by entering its table-of-contents section number.

The requested help text is displayed in one or more sections, depending on the number of display lines that are available in your command window. When enough lines are displayed to fill the display space, or when the end of the requested help text is displayed, you see a prompt along with instructions for what can be entered at that prompt. To continue displaying text for your current selection, press enter or type the 'd' key to scroll down. To scroll up in the current selection, press the 'u' key and press Enter. Other choices might be presented, so read the prompt instructions.

Proper display of the help text requires a usable display width of 72 characters. A display width fewer than 72 characters causes sentences that are 72 characters wide to wrap to the next line. This can cause the displayed help text to begin somewhere within the section rather than at the beginning. The undisplayed lines can be viewed by using the scrolling function of the terminal to move up.

Examples

Task Display the table of contents of the help topics.

Command: `dsmc help`

Task Display the information in help topic 2.1.2

Command: `dsmc help 2.1.2`

Task Display help information on the **archive** command.

Command: `dsmc help archive`

Task Display help information on message ANS1036.

Command: `dsmc help 1036`

Incremental

The **incremental** command backs up all new or changed data in the locations that you specify, unless you exclude them from backup services.

You can back up all new or changed files or directories in the default client domain or from file systems, directories, or files.

To incrementally back up selected files or directories, enter a file specification in the command. If you do not enter a file specification, the default is to back up files or directories in the default domain.

AIX only: You can enable snapshot-based incremental backup by using the option `snapshotproviderfs=JFS2`.

The following attributes in the management class that is assigned to the file or directory affect whether the data is backed up:

Frequency

The number of days that must elapse between successive backups of the object. The **frequency** attribute applies only to a full incremental backup.

This management class attribute is ignored during a journal-based backup.

Mode Specifies whether changes since the last backup operation affect the processing. If `mode=modified`, only objects that changed since the last backup operation are processed. If `mode=absolute`, every object is processed, regardless of whether the object changed since the last backup operation.

If the copy group mode is set to `modified`, it can be overridden by using the client **absolute** option. For more information about the **absolute** option, see "Absolute" on page 316.

Serialization

Permits or denies backup of files or directories according to the following values:

- **static**: To be backed up, data must not be modified during backup or archive.
- **shared static**: If data in the file or directory changes during each of the allowed attempts to back up or archive it, it is not backed up or archived. The value of the `changingretries` option determines how many attempts are made. The default is 4.
- **dynamic**: The object is backed up or archived on the first attempt whether or not data changes during the process.
- **shared dynamic**: The object is backed up or archived on the last attempt, even if data changes during the process.

Using the **include** option in an include-exclude list, you can override the default management class for a file or group of files.

You can perform either a full incremental backup or an incremental-by-date backup. The default is a full incremental backup.

If you are journaling a file system and the journal is valid, the full incremental backup performs a journal-based backup. More than one journal-based backup session can be started, but only one journal-based backup session can proceed. All other journal-based backup sessions that need access to the same file space must wait until the current journal-based backup session completes before the next session can proceed. You can perform a full incremental backup without the journal by using the **nojournal** option.

You can also use the **selective** command to perform a backup that backs up only the files, directories, or empty directories that you specify regardless of whether they were changed.


A full incremental backs up all files and directories that are new or were changed since the last incremental backup. During a full incremental backup, the client queries the server. Tivoli Storage Manager uses this information when it performs the following actions:

- Backing up new files or directories.
- Backing up files or directories whose contents were changed since the previous backup.
- Marking inactive backup versions on the server for files or directories that are deleted from the workstation.
- Rebinding backup versions to management classes if the management class assignments change.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

filespec

Specifies the path and file name that you want to back up. Use wildcard characters to select a group of files or all the files in a directory. If you do not specify a file specification, the **domain** option determine what to back up.

If you specify a file system, all new and changed files are backed up. In addition, the last incremental date for the file space is updated on the server. If you specify a file or directory, the last incremental date is not updated. This means that the file or directory might be backed up again if a later backup is performed by using the **incrbydate** option. If you specify a file system, specify the file system without a trailing slash.

Table 92. Incremental command: Related options

Option	Where to use
absolute "Absolute" on page 316	Command line only.
changingretries "Changingretries" on page 333	dsm.sys file or command line.

Table 92. Incremental command: Related options (continued)

Option	Where to use
compressalways "Compressalways" on page 338	Client user-options file (dsm.opt) or command line.
compression "Compression" on page 339	dsm.sys file within a server stanza or command line.
detail "Detail" on page 353	Command line only.
diffsnapshot "Diffsnapshot" on page 354	Command line only.
dirsonly "Dirsonly" on page 357	Command line only.
domain "Domain" on page 361	dsm.sys file or the client user-options file (dsm.opt) or command line.
encryptiontype "Encryptiontype" on page 383	System-options file (dsm.sys) within a server stanza.
encryptkey "Encryptkey" on page 384	System-options file (dsm.sys) within a server stanza.
filelist "Filelist" on page 402	Command line only.
filesonly "Filesonly" on page 406	Command line only.
incrbydate "Incrbydate" on page 436	Command line only.
memoryefficientbackup "Memoryefficientbackup" on page 452	Client user-options file (dsm.opt), client system-options file (dsm.sys), server, or command line.
nojournal "Nojournal" on page 463	Command line only.
preservelastaccessdate "Preservelastaccessdate" on page 479	Client user-options file (dsm.opt) or command line.
removeoperandlimit "Removeoperandlimit" on page 488	Command line only.
snapdiff "Snapdiff" on page 519	Command line only.
snapshotcachesize "Snapshotcachesize" on page 525	Client options file (dsm.opt) or with the include.fs option.
snapshotproviderfs "Snapshotproviderfs" on page 526	System-options file (dsm.sys) within a server stanza or with the include.fs option.
snapshotroot "Snapshotroot" on page 528	Command line only.
subdir "Subdir" on page 539	Client user-options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 541	Client user-options file (dsm.opt) or command line.

Examples

Task Run an incremental backup of the client domain that is specified in your client user-options file (dsm.opt).

Incremental

Run an incremental backup that backs up all files in the domain regardless of whether they were changed since the last backup.

Incremental -absolute

Task Run an incremental backup for the /home, /usr, and /proj file systems.

```
Incremental /home /usr /proj
```

Task Run an incremental backup for the /proj/test directory.

```
Incremental /proj/test/
```

Task Run an incremental-by-date backup for the /home file system.

```
Incremental -incrbydate /home
```

Task Run an incremental backup of the abc file in the /fs/dir1 directory.

```
Incremental -subdir=yes /fs/dir1/abc
```

Task Run an incremental backup of the directory object /fs/dir1, but not any of the files in the /fs/dir1 directory.

```
Incremental /fs/dir1
```

Task Run an incremental backup of the directory object /fs/dir1, all of the files in the fs/dir1 directory, and all files and subdirectories under /fs/dir1.

```
Incremental -subdir=yes /fs/dir1/
```

Task Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, run an incremental backup of all files and directories under the local snapshot and manage them on the Tivoli Storage Manager server under the file space name /usr.

```
dsmc inc /usr -snapshotroot=/snapshot/day1
```

Task Run an incremental backup for the /home file system by using the **snapdiff** option and specify the option to create the difference snapshot. In the following example, /home is the NFS mount point for a NAS/N-Series file server volume.

```
incremental /home -snapdiff -diffsnapshot=create
```

Task Run an incremental backup of the /proj file system by using the **snapdiff** option. Specify the option to use the latest snapshot on the file server as the difference snapshot. In the following example, /proj is the NFS mount point for a NAS/N-Series file server volume.

```
incremental /proj -snapdiff -diffsnapshot=latest
```

Related information

“Absolute” on page 316

“Journal-based backup”

“Selective” on page 725

“Include options” on page 421

Journal-based backup

A backup for a particular file system is journal-based when the Tivoli Storage Manager journal daemon is installed and configured to journal the file system, and a valid journal has been established.

Journal-based backup is supported on the AIX Backup-Archive Client, on JFS and JFS2 file systems.

Journal-based backup is supported on the Linux Backup-Archive client on Ext2, Ext3, Ext4; XFS, ReiserFS, JFS, VxFS, and NSS. GPFS is not supported for journal-based backups on Linux.

If the journal daemon is installed and running, then by default the **incremental** command performs a journal-based backup on file systems which are being monitored by the journal engine daemon. The following conditions must be met in order to successfully perform a journal-based backup:

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

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

The backup-archive 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.

The journal daemon does not record changes in UNIX special files.

The journal daemon excludes specific system files from having changes recorded in the journal. Because changes to these files are not journaled, the client does not back up these files. See the journal daemon configuration file `tsmjbbd.ini` located in the backup-archive client installation directory for specific system files that are excluded.

Note:

1. When using antivirus software, there are limitations to journal-based backup. Some antivirus software can incorrectly generate change notifications to the TSM journal service, causing files that have not changed to be incorrectly backed up during journal based backup. To avoid these problems, use Norton Anti-Virus Corporate Edition 8.0 and higher.
2. A journal-based backup might not fall back to the traditional incremental backup if the policy domain of your node is changed on the server. This depends on when the policy set within the domain was last updated and the date of the last incremental backup. In this case, you must force a full traditional incremental backup to rebind the files to the new domain. Use the `nojournal` option with the **incremental** command to specify that you want to perform a traditional full incremental backup, instead of the default journal-based backup.

Add an exclude snapshot statement to the `tsmjbbd.ini` file for AIX 6.1 (or later) to prevent JFS2 internal snapshot directories from being monitored by the

journal-based backup daemon. If you do not exclude the snapshot directories, the files in them are backed up. Backing up the snapshot directories is redundant and wastes server space.

Under the following conditions, the journal database is considered invalid and the client reverts to the traditional full incremental backup:

- A journaled file space name has changed.
- The client node name has changed.
- The client contacts a different server to do the backup.
- Policy changes have occurred (new policy set activation).
- The journal is corrupt (out of space conditions, disk error).
- The journal is not running.

Journal-based backup differs from the traditional full incremental backup in the following ways:

- Tivoli Storage Manager does not enforce non-default copy frequencies (other than 0).
- Changes to UNIX special files are not detected.

You can use the `nojournal` option with the **incremental** command to perform a traditional full incremental backup instead of the default journal-based backup.

Incremental-by-Date

An incremental-by-date backup backs up new and changed files with a modification date later than the date of the last incremental backup stored at the server, unless the files are excluded from backup by an **exclude** statement.

If an incremental-by-date is performed on only part of a file system, the date of the last full incremental is not updated, and the next incremental-by-date will back up these files again. Changes to the access control lists (ACL) or Extended Attributes do not cause the files to be backed up during an incremental-by-date. Use the **query filespace** command to determine the date and time of the last incremental backup of the entire file system.

To perform an incremental-by-date backup, use the `incrbydate` option with the **incremental** command.

Unlike a full incremental, an incremental-by-date does not maintain current server storage of *all* your workstation files for the following reasons:

- It does not expire backup versions of files that are deleted from the workstation.
- It does not rebind backup versions to a new management class if the management class has changed.
- It does not back up files with attributes that have changed, unless the modification dates and times have also changed.
- It ignores the copy group frequency attribute of management classes.

For these reasons, if you have limited time during the week to perform backups, but extra time on the weekends, you can perform an incremental-by-date backup on weekdays and a full incremental backup on weekends to maintain current server storage of your workstation files.

If the **incremental** command is retried because of a communication failure or session loss, the transfer statistics will display the number of bytes that the client

attempted to transfer during all command attempts. Therefore, the statistics for bytes transferred might not match the file statistics, such as those for file size.

Associate a local snapshot with a server file space

Use the `snapshotroot` option with the **incremental** command in conjunction with a vendor-supplied application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server.

The `snapshotroot` option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot.

Loop

The **loop** command starts an interactive command line session that is maintained until you enter `quit`.

If you are required to enter a password, you are prompted for it before the loop mode prompt appears.

Note: It is not possible to enter loop mode without a valid server contact. One of the consequences is that certain commands, such as `restore backupset -location=file`, are only accepted on the initial command line when a valid server is not available.

In an interactive command line session, it is unnecessary to precede each command name with **dsmc** and your password, if one is required.

In interactive mode, options that you enter on the initial command line override the value that you specified in your client user-options file (`dsm.opt`) or `dsm.sys` file. This value remains in effect for the entire interactive session unless overridden by a different value on a given interactive command. For example, if you set the `subdir` option to `yes` in your client user-options file (`dsm.opt`), and you specify `subdir=no` on the initial command line, the `subdir=no` setting remains in effect for the entire interactive session unless overridden by the `subdir=yes` value on a given interactive command. However, the `subdir=yes` value only affects the command it is entered on. When that command completes, the value reverts back to `subdir=no`, the value at the beginning of the interactive session.

You can enter all valid commands in interactive mode *except* the **schedule** and **loop** commands.

There are some options that you cannot use in the interactive session created by the **loop** command and are identified in the option description by this statement: *This option is valid only on the initial command line. It is not valid in interactive mode.*

Note:

1. In loop mode, following a restore operation directly from tape, the mount point is not released in case additional restore requests are made to that volume. If you request a backup operation in the same session and that mount point is the only one available, the backup operation stops with the following message:

Waiting for mount of offline media

In this case, the mount point is not released until one of the following conditions is met:

- The device class MOUNTRETENTION limit is satisfied.
 - The client idletimeout period is satisfied.
 - The dsmc loop session is closed after the restore operation completes, allowing you to start a subsequent loop mode session to perform the backup operation.
2. In interactive mode, you cannot enter a file specification that contains national language characters. If a command contains national characters, process the command in batch mode by preceding the command with the executable program name, **dsmc**.

Supported Clients

This command is valid for all clients.

Syntax

»—LOOP—«

Parameters

There are no parameters for this command.

Examples

Task Start an interactive command line session.

Command: dsmc

At the **tsm>** prompt, enter a command.

There are two methods for ending an interactive session:

- Enter quit
- If you set editor=*yes*, you can do the following:
 1. Press the Escape key (Esc).
 2. Type Q and press the Enter key.

Note: The default setting is editor=*yes*.

Note: To interrupt a **dsmc** command before the client has finished processing, enter **QQ** on the Tivoli Storage Manager console. In many cases, but not all, this interrupts the command.

Related information

Chapter 10, “Processing options,” on page 287 for options that you cannot use in interactive mode.

Macro

The **macro** command runs 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 10 levels of commands.

Comment lines are not supported within the macro file that you specify for the **macro** command.

Supported Clients

This command is valid for all clients.

Syntax

►►—Macro— *macroname*—————►►

Parameters

macroname

Specifies the fully qualified name of the file that contains the commands.

Examples

The following is an example of how to use the **macro** command.

Task Selectively back up files in the following directories:

- /devel/project/proja
- /devel/project/projb
- /devel/project/projc

Command: macro backabc.mac

Where backabc.mac contains the following statements:

```
Selective /devel/project/proja/  
Selective /devel/project/projb/  
Selective /devel/project/projc/
```

Monitor Process

The **monitor process** command displays a list of current NAS (if NDMP support is enabled) image backup and restore processes for which the administrative user has authority. You are prompted for the Tivoli Storage Manager administrator ID.

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.

Supported Clients

This command is valid for AIX, Linux, and Solaris clients only.

Syntax

►►—MONitor Process—————►►

Parameters


There are no parameters for this command.

Task	Monitor current NAS image backup or restore processes.
Command:	monitor process

The **preview archive** command simulates an archive command without sending data to the server.

Supported Clients

Syntax



filespec

-filter Specifies the output to display. You can display included objects, excluded objects, or both.

INCLuded

EXCLuded

-FILENAME=

-CONsole

-TRAverse

Yes Preview the current directories and subdirectories. This is the default.

No Preview only the current directory, not subdirectories.

Important: Specifying **-traverse** does not preview directories excluded using the `exclude.dir` option.

Preview Backup


The **preview backup** command simulates a backup command without sending data to the server.

The **preview backup** command generates a tab-delineated text file that can be imported into a spreadsheet program. The preview contains information such as whether the file is excluded or included. If the file is excluded, the pattern, or reason, that the file is excluded is listed, along with the source for the pattern.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

filespec

Specifies the path and file name that you want to back up. Use wildcard characters to select a group of files or all the files in a directory.

-filter Specifies the output to display. You can display included objects, excluded objects, or both.

ALL Display output for included and excluded objects. This is the default.

INCLuded

Display output for included objects only.

EXCLuded

Display output for excluded objects only.

-FILENAME=

Specifies the filename in which to write the tab-delineated output. The default is `dsmprev.txt`.

-CONsole

Output is written to the console, and the file.

-TRAverse

Preview the current directory and subdirectories.

Yes Preview the current directories and subdirectories. This is the default.

No Preview only the current directory, not subdirectories.

Important: Specifying **-traverse** does not preview directories excluded using the `exclude.dir` option.

Query Access

The **query access** command shows who was given access to backup versions or archive copies of specific files.

The backup-archive client displays a list of authorization rules that you defined with the **set access** command or with the **Utilities > Node Access List** menu in the backup-archive client graphical user interface (GUI).

The following information is included.

- Authority that you gave a user to restore backup versions or retrieve archive copies.
- The node name of the user to whom you gave authorization.
- The ID of the user at that node to whom you gave authorization.
- The files to which the user has access.
- The vApps to which a user or node has access.

Supported Clients

This command is valid for all clients.

Syntax

►►—Query Access—◄◄

Parameters

The **-asnodename** option must be supplied when you issue this command to determine which users or nodes have access to VMware vApp backups.

Examples

Task Display a list of nodes that have access to backed up vApps.

Command: `dsmc query access -asnodename=pvdcnode`

where: the **-asnodename** option specifies the name of the provider virtual data center node that manages vApps that you have access to. Specify the name of the provider virtual datacenter node that manages the vApps, on the **-asnodename** option.

Task Display a list of users who have access to your files.

Command: `query access`

Query Archive

The **query archive** command displays a list of your archived files and the following information about each file: file size, archive date, file specification, expiration date, and archive description.

If you use the `detail` option with the **query archive** command, the client displays the following additional information:

- Last modification date
- Last access date
- Last file attributes (inode) change date
- Compression type
- Encryption type
- Client-side data deduplication
- Retention initiation
- Whether the file is on hold

The following example shows sample output when the **query archive** command is issued with the `detail` option:

```
Size Archive Date - Time File - Expires on - Description
-----
219 B 03/03/2016 09:32:13 /Volumes/Data/info.txt 03/03/2016
Archive Date: 03/03/2016
RetInit:STARTED Obj
Held:NO
Modified: 03/02/2016 19:43:00 Accessed: 03/03/2016 09:31:23 Inode changed: 03/02/2016 19:43:00
Compression Type: LZ4 Encryption Type: None Client-deduplicated: YES
```

For more information about the compression type, see “Compression” on page 339.

Supported Clients

This command is valid for all clients.

Syntax

```
➤➤ Query Archive [ _options_ ] [ _filespec_ ] _____ ➤➤
                  [ _"filespec"_ ]
```

Parameters

filespec

Specifies the path and file name that you want to query. Use wildcard characters to specify a group of files or all the files in a directory. If you use wildcard characters, enclose the file specification in double quotation marks. Specify an asterisk (*) to query all archived files in the current directory.

Table 93. Query Archive command: Related options

Option	Where to use
dateformat “Dateformat” on page 344	Client user-options file (dsm.opt) or command line.
description “Description” on page 351	Command line only.
detail “Detail” on page 353	Command line only.
dirsonly “Dirsonly” on page 357	Command line only.
filelist “Filelist” on page 402	Command line only.
filesonly “Filesonly” on page 406	Command line only.

Table 93. Query Archive command: Related options (continued)

Option	Where to use
fromdate "Fromdate" on page 409	Command line only.
fromnode "Fromnode" on page 410	Command line only.
fromowner "Fromowner" on page 411	Command line only.
fromtime "Fromtime" on page 412	Command line only.
numberformat "Numberformat" on page 465	Client user-options file (dsm.opt) or command line.
querysummary "Querysummary" on page 485	Command line only.
scrolllines "Scrolllines" on page 511	Client user-options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 512	Client user-options file (dsm.opt) or command line.
subdir "Subdir" on page 539	Client user-options file (dsm.opt) or command line.
timeformat "Timeformat" on page 549	Client user-options file (dsm.opt) or command line.
today "Today" on page 554	Command line only.
totime "Totime" on page 554	Command line only.

Examples

Task Display a list of all your archived files in the current working directory.

Command: `q archive "*"`

Task Display a list of all your archived files in the /devel directory and all of its subdirectories.

Command: `query archive "/devel/*" -subdir=yes`

Task Display a list of all your archived files in the current directory. Use the dateformat and timeformat options to reformat the dates and times.

Command: `q ar -date=5 -time=1 "*"`

Task Display a list of all your archived files in the current directory. Use the detail option to display the last modification date and the last access date of each file.

Command: `q ar -detail "*"`

Task Display a list of archived files in the /home/proj directory whose first four characters of the file name begin with proj.

Command: `q ar "/home/proj/proj*"`

Query Backup

The **query backup** command displays a list of backup versions of your files that are stored on the Tivoli Storage Manager server, or that are inside a backup set from the server when the `backupsetname` option is specified.

The command displays the following file information:

- File specification
- File size
- Backup date
- Whether the file is active or inactive
- The management class that is assigned to the file. Only the first 10 characters of the management class name are displayed.

If you use the `detail` option with the **query backup** command, the client displays the following extra information:

- Last modification date
- Last access date
- Last file attributes (inode) change date
- Compression type
- Encryption type
- Client-side data deduplication
- Whether the file is migrated or premigrated. A value of Yes means that the file is migrated or premigrated. A value of No means that the file is not migrated or premigrated.

The following example shows sample output when the **query backup** command is issued with the `detail` option:

Size		Backup Date	Mgmt Class	A/I	File
----		-----	-----	----	----
1,500,000	B	03/15/2016 16:01:25	DEFAULT	A	/home/test/mydir/myfile1.txt
Modified: 03/15/2016 16:00:10 Accessed: 03/16/2016 15:31:23 Inode changed: 03/15/2016 16:00:10					
Compression Type: LZ4 Encryption Type: None Client-deduplicated: YES					

For more information about the compression type, see “Compression” on page 339.

Supported Clients

This command is valid for all clients.

Syntax

```
►► Query Backup [—options] [—filespec] [—"filespec"]
```

Parameters

filespec

Specifies the path and file name that you want to query. Use wildcard characters to specify a group of files or all the files in a directory. If you use wildcard characters, enclose the file specification in double quotation marks. Specify an asterisk (*) to display information about backup versions for all of your files in the current directory. Do not use wildcard characters when you query NAS file system images with `-class=nas` option setting.

Table 94. Query Backup command: Related options

Option	Where to use
backupsetname "Backupsetname" on page 329	Command line only.
class "Class" on page 333	Command line only.
dateformat "Dateformat" on page 344	Client system options file (dsm.sys) or command line.
detail "Detail" on page 353	Command line only.
dironly "Dironly" on page 357	Command line only.
filelist "Filelist" on page 402	Command line only.
filesonly "Filesonly" on page 406	Command line only.
fromdate "Fromdate" on page 409	Command line only.
fromowner "Fromnode" on page 410	Command line only.
fromowner "Fromowner" on page 411	Command line only.
fromtime "Fromtime" on page 412	Command line only.
inactive "Inactive" on page 419	Command line only.
nasnodename "Nasnodename" on page 460	Client system options file (dsm.sys) or command line.
numberformat "Numberformat" on page 465	Client user-options file (dsm.opt) or command line.
pitdate "Pitdate" on page 473	Command line only.
pittime "Pittime" on page 474	Command line only.
querysummary "Querysummary" on page 485	Command line only.
scrolllines "Scrolllines" on page 511	Client user-options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 512	Client user-options file (dsm.opt) or command line.
subdir "Subdir" on page 539	Client user-options file (dsm.opt) or command line.
timeformat "Timeformat" on page 549	Client user-options file (dsm.opt) or command line.

Table 94. Query Backup command: Related options (continued)

Option	Where to use
todate "Todate" on page 554	Command line only.
totime "Totime" on page 554	Command line only.

Examples

Task Display a list of all active and inactive backup versions of your files in the current directory.

```
dsmc query backup -inactive "*"
```

Task Display a list of all your backups in the current directory. Use the detail option to display the last modification date and the last access date of each file.

```
dsmc q backup -detail "*"
```

Task Display a list of files that were backed up from the /home/proj directory with file names that begin with proj.

```
dsmc q b "/home/proj/proj*"
```

Task Display a list of active and inactive backup file versions in the /home file system.

```
dsmc q b -ina -su=yes /home/
```

Task Query file system images from the nas2 NAS file server.

```
dsmc query backup -nasnodename=nas2 -class=nas
```

Related information

"Restore data from a backup set" on page 224

Query NAS file system images

You can use the **query backup** command to display information about file system images backed up for a NAS file server. The client prompts you for an administrator ID.

Where supported, use the nasnodename option to identify the NAS file server to query. Place the nasnodename option in your client system-options file (dsm.sys). The value in the client system-options file is the default, but this value can be overridden on the command line.

Use the class option to specify the class of the file space to query. To display a list of images belonging to a NAS node, use the -class=nas option.

Related reference:

"Class" on page 333

"Nasnodename" on page 460

Query Backupset

The **query backupset** command queries a backup set from a local file, tape device (if applicable), or the Tivoli Storage Manager server.


This command displays the backup set name, generation date, retention (for a backup set on the Tivoli Storage Manager server), and user-supplied description.

Supported Clients

This command is valid for all clients.

Tape support is only available on AIX, Solaris, and HP-UX.

Syntax



Parameters

BACKUPSETName=

Specifies the name of a backup set you want to query. You can use wildcards to specify the backup set name. If you use wildcards or do not specify a backup set name, all backup sets that you own are displayed. This parameter is required.

When a backup set is created, the server assigns root as the owner of the backup set. When querying a backup set on the server, a non-root user does not see the backup set listed, even if they know the backup set name and use it in the query.

The value of **backupsetname** depends on the location of the backup set, and corresponds to one of these three choices:

backupsetname

Specifies the name of the backup set from the server. If the **location** parameter is specified, you must set `-location=server`.

localfilename

Specifies the file name of the first backup set volume. You must set `-location=file`.

tapedevice

Specifies the name of the tape device that contains the backup set volume. You must use a Windows native device driver, not the device driver that is provided by Tivoli. You must set `-location=tape`.

LOCation=

Specifies where the backup-archive client searches for the backup set. If you do not specify the location parameter, the client searches for backup sets on the Tivoli Storage Manager server.

server Specifies that the client searches for the backup set from the server. This location is the default.

file Specifies that the client searches for the backup set from a local file.

tape Specifies that the client searches for the backup set from a local tape device.

Table 95. Query Backupset command: Related options

Option	Where to use
description "Description" on page 351	Command line only.
scrolllines "Scrolllines" on page 511	Client user-options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 512	Client user-options file (dsm.opt) or command line.

Examples

Task Query all backup sets from the Tivoli Storage Manager server.

Command: query backupset -backupsetname=*

Task Query a backup set that is called monthly_financial_data from the Tivoli Storage Manager server.

Command: query backupset
-backupsetname=monthly_financial_data.12345678

Task Query the backup set in the file /home/budget/weekly_budget_data.ost.

Command: dsmc query backupset -backupsetname="/home/budget/
weekly_budget_data.ost" -loc=file

Task Query the backup set from the /dev/rmt0 tape device.

Command: dsmc query backupset -backupsetname=/dev/rmt0 -loc=tape

Related information

"Restore data from a backup set" on page 224

Query Backupset without the backupsetname parameter

The **query backupset** command can be used without the **backupsetname** parameter.


The preferred syntax for **query backupset** command requires the **backupsetname** parameter. Prior to the introduction of the **backupsetname** parameter, the backup-archive client queried backup sets with a different syntax.

While you can use syntax from previous releases for this command, do not do so unless you have a specific need and cannot replace the old syntax with the syntax in Tivoli Storage Manager Version 6.1. For best results, use the **backupsetname** parameter.

Supported Clients

This command is valid for all clients. Tape support is only available on AIX, Solaris, and HP-UX.

Syntax



Parameters

backupsetname

Specifies the name of the backup set from the Tivoli Storage Manager server. If the **location** parameter is specified, you must set **-location=server**.

localfilename

Specifies the file name of the first backup set volume. You must set -location=file.

tapedevice

Specifies the name of the tape device containing the backup set volume. You must use a Windows native device driver, not the device driver provided by Tivoli. You must set `-location=tape`.

LOCation=

Specifies where the client searches for the backup set. If you do not specify the location parameter, the client searches for backup sets on the Tivoli Storage Manager server.

server Specifies that the client searches for the backup set from the server. This is the default.

file Specifies that the client searches for the backup set from a local file.

tape Specifies that the client searches for the backup set from a local tape device.

Table 96. Query Backupset command: Related options

Option	Where to use
description “Description” on page 351	Command line only.
scrolllines “Scrolllines” on page 511	Client user-options file (dsm.opt) or command line.
scrollprompt “Scrollprompt” on page 512	Client user-options file (dsm.opt) or command line.

Examples

Task Query all backup sets from the Tivoli Storage Manager server.

Command: query backupset

Task Query a backup set called `monthly_financial_data` from the Tivoli Storage Manager server.

Command: query backupset monthly financial data.12345678

Task Query the backup set in the file /home/budget/weekly budget data.ost.

Command: dsmc query backupset /home/budget/weekly_budget_data.ost
-loc=file

Task Query the backup set from the /dev/rmt0 tape device.

Command: dsmc query backupset /dev/rmt0 -loc=tape

Related information

“Restore data from a backup set” on page 224

Query Filespace

The **query filesystem** command displays a list of file spaces for a node. The file spaces are stored on the Tivoli Storage Manager server, or inside a backup set from the server when the backupsetname option is specified. You can also specify a single file space name to query.

A *file space* is a logical space on the server that contains files you backed up or archived. A separate file space is assigned on the server for each node at your workstation from which you back up or archive files.

A separate file space is assigned on the server for each file system at your workstation from which you back up or archive files. The file space name is the same as the file system name.

Supported Clients

This command is valid for all clients.

Syntax

►► Query Filespace —filesystemname— —options— ►►

Parameters

filesystemname

Specifies an optional character string that can include wildcards. Use this argument to specify a subset of file spaces. The default is to display all file spaces.

Table 97. Query Filespace command: Related options

Option	Where to use
backupsetname “Backupsetname” on page 329	Command line only.
class “Class” on page 333	Command line only.
dateformat “Dateformat” on page 344	Client user-options file (dsm.opt) or command line.
detail “Detail” on page 353	Command line only.
fromnode “Fromnode” on page 410	Command line only.

Table 97. Query Filespace command: Related options (continued)

Option	Where to use
fromowner "Fromowner" on page 411	Command line only.
nasnodename "Nasnodename" on page 460	Client system options file (dsm.sys) or command line.
scrolllines "Scrolllines" on page 511	Client user-options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 512	Client user-options file (dsm.opt) or command line.
timeformat "Timeformat" on page 549	Client user-options file (dsm.opt) or command line.

Examples

Display your file spaces. Use the dateformat and timeformat options to reformat the dates and times.

```
query filesystem -date=5 -time=4
```

Display the /home file space.

```
query filesystem /home
```

Display file space names that include the pattern smith.

```
query filesystem "*smith"
```

Query a file space from the nas2 NAS file server.

```
query filesystem -nasnodename=nas2 -class=nas
```

Display detailed file space information that shows the replication status during a failover.

Command:

```
query filesystem -detail
```

Output:

#	Last Incr Date	Type	fsID	Unicode	Replication	File Space Name
1	00/00/0000 00:00:00	HFS	3	Yes	Current	/
	Last Store Date	Server			Local	
	Backup Data :	04/29/2013 16:49:55			04/29/2013 16:49:55	
	Archive Data :	No Date Available			No Date Available	

Related concepts:

"Restore data from a backup set" on page 224

"Automated client failover overview" on page 87

Related tasks:

"Determining the status of replicated client data" on page 92

Related reference:

"Nasnodename" on page 460

"Class" on page 333

"Nrtablepath" on page 464

Query NAS file spaces

Use the `nasnodename` option to identify the NAS file server to query. When using an interactive command-line session with a non-administrative ID, the client prompts you for an administrator ID.

Place the `nasnodename` option in your client system-options file (`dsm.sys`). The value in the client system-options file is the default, but this value can be overridden on the command line. If the `nasnodename` option is not specified in the client system-options file, it must be specified on the command line when processing NAS file systems.

Use the `class` option to specify the class of the object to query. To display a list of file spaces belonging to a NAS node, use the `-class=nas` option.

Query Group

Use the **query group** command to display information about a group backup and its members.

Note:

1. Use the `showmembers` option to display and select individual group members that you want to query. The `showmembers` option is not valid with the `inactive` option. If you want to display members of a group that are not currently active, use the `pitdate` and `pittime` options to specify the backup date and time of the member you want to query.
2. Use the **query filesystem** command to display virtual file space names for your node that are stored on the Tivoli Storage Manager server.
3. If you perform a full and differential group backup, a query of this group using the `-inactive` option displays two active backups of the same name, one of type FULL and one of type DIFF.

```
tsm> q group {/fs}/v1 -inactive
```

Size	Backup	Date	Mgmt	Class	A/I	Group
978	B	06/02/2007 11:57:04	DEFAULT	A	FULL	/fs/v1
32	B	06/05/2007 13:52:04	DEFAULT	A	DIFF	/fs/v1

If you query a group backup without the `-inactive` option, the query displays only the latest group backup, whether it is type FULL or type DIFF:

```
tsm> q group {/fs}/v1
```

Size	Backup	Date	Mgmt	Class	A/I	Group
32	B	06/05/2007 13:52:04	DEFAULT	A	DIFF	/fs/v1

Supported Clients

This command is valid for all clients, except for Mac OS X.

Syntax

►►—Query Group— —*filespec*— —*options*—►►

Parameters

filespec

Specifies the virtual file space name and the group name on the server that you want to query.

Table 98. Query Group command: Related options

Option	Where to use
fromnode “Fromnode” on page 410	Command line only.
fromowner “Fromowner” on page 411	Command line only.
inactive “Inactive” on page 419	Command line only.
pitdate “Pitdate” on page 473	Command line only.
pittime “Pittime” on page 474	Command line only.
showmembers “Showmembers” on page 517 (does not apply to Mac OS X)	Command line only.

Examples

Task Display all the groups in the /virtfs file space.

Command:

```
query group /virtfs/*
```

Task Display active and inactive versions of the /virtfs/group1 file space.

Command:

```
query group /virtfs/group1 -inactive
```

Task Display the /virtfs/group1 file space. Use the showmembers option to display a list of group members from which you can select one or more to query.

Command:

```
query group /virtfs/group1 -showmembers
```

Related information

“Query Filespace” on page 674

Query Image

The **query image** command displays information about file system images that are stored on the Tivoli Storage Manager server, or that are inside a backup set from the Tivoli Storage Manager server, when the `backupsetname` option is specified.

The following information about file system images is displayed:

- Image Size - The volume size which was backed up.
- Stored Size - The actual image size that is stored on the server. The stored image on the Tivoli Storage Manager server is the same size as the volume capacity. For online snapshot-based image backups, the stored image can be larger than the file system based on the size of the cache files. The stored image on the server is the same size as the volume capacity.
- File system type
- Backup date and time
- Management class that is assigned to image backup
- Whether the image backup is an active or inactive copy
- The image name

Note: The Tivoli Storage Manager API must be installed to use the **query image** command.

Supported Clients

This command is valid for AIX, HP-UX, all Linux clients, and Solaris.

Syntax

```
➤➤ query Image [—options] [—logicalvolumename] [—filespacename] ➤➤
```

Parameters

logicalvolumename

The name of a logical volume you want to query. You must specify the exact name of the image. You cannot use wildcards. The default is all active images (unless restricted by one or more options).

filespace_{name}

Specifies the file system name that you want to query.

Omitting *logicalvolumename* and *filespace_{name}* causes all images to be displayed.

Table 99. Query Image command: Related options

Option	Where to use
<code>backupsetname</code> "Backupsetname" on page 329	Command line only.
<code>dateformat</code> "Dateformat" on page 344	Client user option file (dsm.opt) or command line.
<code>fromnode</code> "Fromnode" on page 410	Command line only.

Table 99. Query Image command: Related options (continued)

Option	Where to use
fromowner "Fromowner" on page 411	Command line only.
inactive "Inactive" on page 419	Command line only.
numberformat "Numberformat" on page 465	Client user option file (dsm.opt) or command line.
pitdate "Pitdate" on page 473	Command line only.
pittime "Pittime" on page 474	Command line only.
scrolllines "Scrolllines" on page 511	Client user options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 512	Client user options file (dsm.opt) or command line.
timeformat "Timeformat" on page 549	Client user option file (dsm.opt) or command line.

Examples

Task Display all backed up images.

Command: q image

Task Display all backed up images that are owned by kutras at node avalon .

Command: query image -fromnode=avalon -fromowner=kutras

Task Display active and inactive version of the /usr image.

Command: q i /usr -inactive

Task Display all images that are contained within the backup set weekly_backup_data.32145678.

Command: query image -backupsetname=weekly_backup_data.32145678

Related information

"Restore data from a backup set" on page 224

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, and so on), and the name of the source file.

Tivoli Storage Manager excludes some files from file system backup and restore operations. You can use the **query inclexcl** command to display a list of these files. In the output of the command, these files have Operating System next to the path.

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.

Use the *detail* option to display the management class that is associated with an include-exclude statement.

Supported Clients

This command is valid for all clients.

Syntax

```
►► Query INCLexcl — [test pattern] [-DETail] ►►
```

Parameters

test pattern

Use for testing the validity of patterns you want to use in your include-exclude list. When you use a test pattern with this command, the following occurs:

- The internal include-exclude list is not displayed
- The pattern is processed as if it came from an include-exclude statement, including all the usual error checking
- The pattern is displayed as it would appear in the include-exclude list

If the test pattern has no errors, the compiled pattern result is the same as the test pattern.

-DETail

Displays the management class that is associated with the include-exclude statement.

Examples

Task Exclude a file from deduplication by excluding it in the client options file:

```
Exclude Dedup *\\...\\file2
```

Task Display a basic list of include-exclude statements. Command:

```
query incl excl
```

Task Display a list of include-exclude statements. Display the management class that is associated with each statement.

```
query incl excl -detail
```

Task Test the validity of this pattern: `/.../?x?/*.log`

```
query incl excl /.../?x?/*.log
```

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 which control 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 extra management classes. You can assign specific management classes to files using `include` options that are located in the client user-options file (`dsm.opt`). If you do not assign a management class to a file, the default management class is used.

When you archive files, you can override the assigned management class by using the `archmc` option.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

Table 100. Query Mgmtclass command: Related options

Option	Where to use
detail “Detail” on page 353	Command line only.
fromnode “Fromnode” on page 410	Command line only.

Examples

Task Display default and available management classes.
Command: `query mgmtclass`

Query Node

The **query node** command displays all the nodes for which an administrative user ID has authority to perform operations. You are prompted for the Tivoli Storage Manager administrator ID.

Ideally, the administrative user ID has at least client owner authority over the client workstation node they are using either from the command line or from the web.

Use the `type` option to specify the type of node to filter for. The following are the valid values:

- `nas`
- `client`
- `server`
- `any`

The default is **any**.

Note: When the Tivoli Storage Manager for Virtual Environments: Data Protection for VMware license file is installed on a vStorage backup server, the platform string that is stored on the Tivoli Storage Manager server is set to “TDP VMware”

for every nodename that is used on that machine. The platform string can be used in the context of PVU calculations. If a nodename is being used to back up the machine with standard Backup-Archive client functions (for example, file-level or image backup), then this platform string would be interpreted as a “client” for the purposes of PVU calculations.

For more information about processor value units, see *Estimating processor value units* in the Tivoli Storage Manager server documentation.

Supported Clients

This command is valid for all clients.

Syntax

```
►► Query Node [—options] ►►
```

Parameters

Table 101. Query Node command: Related options

Option	Where to use
type “Type” on page 556	Command line only.
scrolllines “Scrolllines” on page 511	Client user options file (dsm.opt) or command line.
scrollprompt “Scrollprompt” on page 512	Client user options file (dsm.opt) or command line.

Examples

Task Display all NAS nodes.

Command: query node -type=nas

Task Display all client nodes that are backup-archive clients.

Command: query node -type=client

Related information

“Type” on page 556


Query Options

Use the **query options** command to display all or part of your options and their current settings that are relevant to the command-line client.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

pattern

An optional character string that can include wildcards. Use this argument to specify a subset of options. The default is to display all options.

Table 102. Query Options command: Related options

Option	Where to use
scrolllines “Scrolllines” on page 511	Client user options file (dsm.opt) or command line.
scrollprompt “Scrollprompt” on page 512	Client user options file (dsm.opt) or command line.

Examples

Task Display all options and their values.

```
query options
```

Task Display only options that begin with *comm*.

```
query options comm*
```

Task Display the value of the **replace** option.

```
query options replace
```

Task Issue the command to display all options and their values. The failover status information is displayed.

```
query options
```

Output:

```
MYPRIMARYSERVERNAME: SERVER1
MYREPLICATIONSERVER: TARGET
  REPLSERVERNAME: TARGET
    Address: 192.0.2.9
      Port: 1501
        SSLPort: 1502
          GUID: 39.5a.da.d1.ae.92.11.e2.82.d3.00.0c.29.2f.07.d3
            Used: yes
```

Related concepts:

“Automated client failover configuration and use” on page 87

Related tasks:

“Determining the status of replicated client data” on page 92

Query Restore

The **query restore** command displays a list of your restartable restore sessions in the server database. The list contains these fields: owner, replace, subdir, preservepath, source, and destination.

A restartable restore session is created when a wildcard restore command fails because of network outage, client failure, server outage, or a similar problem. When such a failure occurs, the file space is locked on the server and its files cannot be moved off the sequential volumes of the server. To unlock the file space, either restart the restore and allow it to complete (**query restore** command), or cancel the restore (**cancel restore** command). Use **query restore** to determine if you have any restartable restore sessions and which file spaces are affected.

Supported Clients

This command is valid for all clients.

Syntax

►►—Query Restore—►►

Parameters

There are no parameters for this command.

Examples

Task Display your restartable restore session in the server database.

Command: query restore

Query Schedule

The **query schedule** command displays the events that are 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.

Supported Clients

This command is valid for all clients.

Syntax

►►—Query Schedule—►►

Parameters

There are no parameters for this command.

Examples

Task Display your scheduled events.

Command: query schedule

Query Session

The **query session** command displays information about your session, including the current node name, when the session was established, server information, and server connection information.

Supported Clients

This command is valid for all clients.

Syntax

►—Query SSession—◄

Parameters

There are no parameters for this command.

Examples

Task Display your session information.

Command: query session

A sample **query session** display follows:

```
Server Name.....: HALLEY_SERVER1
Server Type.....: Windows
Archive Retain Protect..: "No"
Server Version.....: Ver. 6, Rel. 2, Lev. 0.0
Last Access Date.....: 09/03/2009 09:08:13
Delete Backup Files.....: "No"
Delete Archive Files....: "Yes"
Deduplication.....: "Server Only"
```

```
Node Name.....: HALLEY
User Name.....:
```

Possible client-side deduplication values:

- None
 - Displayed when connected to a pre V6.1 Tivoli Storage Manager server
- Server Only
- Client Or Server

Task

A sample **query session** display with LAN-free enabled follows:

TSM Server Connection Information

```
Server Name.....: TEMPLAR
Server Type.....: AIX
Archive Retain Protect..: "No"
Server Version.....: Ver. 6, Rel. 1, Lev. 4.0
Last Access Date.....: 08/12/10 22:10:15
Delete Backup Files.....: "No"
Delete Archive Files....: "Yes"
```

```
Node Name.....: LAN2
User Name.....: root
```

```
Storage Agent Name.....: TEMPLAR_STA
Storage Agent Type.....: AIX
Storage Agent Version...: Ver. 6, Rel. 1, Lev. 3.3
```

Query Systeminfo

Use the **query systeminfo** command to gather information and output this information to a file or the console.


This command is intended primarily as an aid for IBM support to help diagnosing problems. However, users who are familiar with the concepts addressed by this information might also find it useful.

If you use the console option, no special formatting of the output is performed to accommodate screen height or width. Therefore, the console output can be difficult to read due to length and line-wrapping. If the console output is difficult to read, use the filename option with the **query systeminfo** command. This combination allows the output to be written to a file that can be submitted to IBM support.

Supported Clients

This command is valid for all 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. The default is to gather all items.

You can gather information on one or more of the following items:

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

Note:

1. Use the `filename` option to specify a file name in which to store the information that is gathered from the items you specify. If you do not specify a

file name, by default the information is stored in the /Library/Application Support/tivoli/tsm/client/ba/bin/dsminfo.txt file (for Mac OS X) or the dsminfo.txt file (for other UNIX and Linux).

2. Use the console option if you want to output the information to the console.

Table 103. Query Systeminfo command: Related options

Option	Where to use
console "Console" on page 341	Command line only.
filename "Filename" on page 405	Command line only.

Examples

Task Gather and store the contents of the dsm.opt file and the Tivoli Storage Manager error log file in the tsminfo.txt file.

Command: query systeminfo dsmdptfile errorlog
-filename=tsminfo.txt

Related information

"Filename" on page 405

"Console" on page 341

Query VAPP

Use the **query vapp** command to determine which VMware vApps are backed up.

To use this command, you must have a license to use Tivoli Storage Manager for Virtual Environments V7.1, or later version. Before you can use this command, the Tivoli Storage Manager nodes must be configured for the vCloud environment. In the Data Protection for VMware vCloud GUI, go to the Configuration window and click **Run Configuration Wizard**.

Supported Clients

This command is valid for Linux clients that are configured to perform vCloud Director vApp backups.

Syntax

►►—Query VAPP— —vapp_specification— —-ASNODENAME=—provider_vdc_node—►
└─options─┘

Parameters

vapp_specification

Required parameter. Specifies the vApps that you want to query. Specify an asterisk (*) to query all vApps that belong to all organization virtual data centers that the data mover node has access to.

The vApp specification must include an organization name and an organization virtual data center name. The following parameters identify the vApp to query:

ORG=*organization_name*

Specifies the name of the organization that the organization virtual data center belongs to.

ORGVDC=*org_vdc_name*

Specifies the name of the organization virtual data center that contains the vApps to query.

VAPP=*vapp_name*

Optional parameter that specifies the name of a single vApp to query. If you do not specify a vApp name, the command displays all vApps for the specified organization and organization virtual data center.

You must specify these keyword parameters in the following order: **ORG=**, **ORGVDC=**, **VAPP=**. The following syntax rules also apply:

- If any value contains a space character, enclose the entire specification in quotation marks.
- To specify more than one organization virtual data center name, separate the organization virtual data center names with commas and do not insert a space character after the commas.
"org=organization_name,orgvdc=org_vdc_name1,org_vdc_name2"
- To specify more than one vApp name, separate the vApp names with commas and do not insert a space character after the commas.
"org=organization_name_name,orgvdc=org_vdc_name,
vapp=vapp_name1,vapp_name2,vapp_name3"
- To specify organization virtual data centers from multiple organizations, separate the **org=** and **orgvdc=** pairs with semicolons. Do not insert a space character after the semicolons.
"org=organization_name1,orgvdc=org_vdc_name1;
organization=organization_name2,orgvdc=org_vdc_name2"
- A wildcard can be used in a vApp name to match one or more characters. For example, to query all vAPPs that have names that begin with TEST, use the following syntax:
"org=org1,orgvdc=vdc1,vapp=TEST*"
- To query all vApp snapshots that are restorable by the data mover that protects the organization virtual data center, specify an asterisk (*) as the vApp specification.

dsmc query vapp * -asnodename=providerVDC

-ASNODENAME=*provider_vdc_node*

Required parameter. Specifies the node name of the client node that is associated with a provider virtual data center. This parameter can be set on the command line or in the client options file (dsm.opt).

Related options

Table 104. Query VAPP command: Related options

Option	Where to use
asnodename	Client user-options file (dsm.opt) or client systems option file (dsm.sys), on the command line, or on the General tab in the Preferences editor.
detail	Command line only. Shows more detail about each vApp.

Table 104. Query VAPP command: Related options (continued)

Option	Where to use
inactive	Command line only. Show information about inactive and active vApp snapshots.

Example commands for querying VMware vApps

The following example command queries all backed up active versions of all vApps from all combinations of organizations and organization virtual data centers to which the data mover node has access:

```
dsmc query vapp *
IBM Tivoli Storage Manager
Command Line Backup-Archive Client Interface
  Client Version 7, Release 1, Level 0.0
  Client date/time: 11/14/2013 14:02:26
(c) Copyright by IBM Corporation and other(s) 1990, 2013. All Rights Reserved.
```

```
Node Name: DM_pvdc1_DM1
Session established with server TEST_VM: Windows
  Server Version 6, Release 3, Level 0.0
  Server date/time: 11/14/2013 15:02:26 Last access: 11/14/2013 14:52:00
```

#	Backup Date	Mgmt Class	Size	Type	A/I	vApp
1	09/07/2013 13:47:35	MG1	21.95 GB	IFINCR	A	vApp: VAPP_1 Org : Marketing OvDC: Marketing_Gold
2	09/07/2013 15:14:35	MG2	13.50 GB	IFFULL	A	vApp: WIN27 Org : Development OvDC: Development_fast
3	09/07/2013 16:12:36	MG1	1.95 GB	IFFULL	A	vApp: SUSE10 Org : Sales OvDC: Sales_vDC2
4	09/12/2013 08:40:48	MG17	9.00 GB	IFINCR	A	vApp: RHEL2 Org : Support OvDC: Suport_vDC2

The following example shows detailed output of all backed up active versions of vApps from the organization that is named MilkyWay and the organization virtual data center that is named Orion_vDC2.

```
dsmc query vapp org=MilkyWay,orgvdc=Orion_vDC2 -detail
IBM Tivoli Storage Manager
Command Line Backup-Archive Client Interface
  Client Version 7, Release 1, Level 0.0
  Client date/time: 11/14/2013 14:02:26
(c) Copyright by IBM Corporation and other(s) 1990, 2013. All Rights Reserved.
```

```
Node Name: DM_pvdc1_DM1
Session established with server TEST_VM: Windows
  Server Version 6, Release 3, Level 0.0
  Server date/time: 11/14/2013 15:02:26 Last access: 11/14/2013 14:52:00
```

#	Backup Date	Mgmt Class	Size	Type	A/I	vApp
1	09/07/2013 13:47:35	MG1	21.95 GB	IFINCR	A	vApp: VAPP_1 Org : MilkyWay OvDC: Orion_vDC2
Incremental backup size: 50 KB Owner : Catalin VM[1] VCD Name : vm1 VM[1] vSphere Name : vm1 (3419c577-fd0e-4389-ac9a-6b483e2fb0f6) VM[1] size : 5.66 GB						

```

VM[1] Status      : backup successful
VM[2] VCD Name    : vm2
VM[2] vSphere Name : vm2 (35a1415c-0da7-4b22-b4a9-35d0df26518f)
VM[2] size        : 16.27 GB
VM[2] Status      : backup failed

2 09/07/2013 15:14:35 MG2          13.50 GB IFFULL A vApp: WIN27
                                   Org : MilkyWay
                                   OvDC: Orion_vDC2

Incremental backup size: 100 KB
Owner : Stefan
VM[1] VCD Name : vm1
VM[1] vSphere Name : vm1 (3419c577-fd0e-4389-ac9a-6b483e2fb0f6)
VM[1] size : 4.32 GB
VM[1] Status : backup successful
VM[2] VCD Name : vm2
VM[2] vSphere Name : vm2 (35a1415c-0da7-4b22-b4a9-35d0df26518f)
VM[2] size : 9.17 GB
VM[2] Status : backup failed

```

The following example shows all backed up active and inactive versions of vApps from all organizations and organization virtual data centers.

```

dsmc query vapp * -ina
IBM Tivoli Storage Manager
Command Line Backup-Archive Client Interface
  Client Version 7, Release 1, Level 0.0
  Client date/time: 11/14/2013 14:02:26
(c) Copyright by IBM Corporation and other(s) 1990, 2013. All Rights Reserved.

```

```

Node Name: DM_pvdc1_DM1
Session established with server TEST_VM: Windows
  Server Version 6, Release 3, Level 0.0
  Server date/time: 11/14/2013 15:02:26 Last access: 11/14/2013 14:52:00

```

#	Backup Date	Mgmt Class	Size	Type	A/I	vApp
1	09/07/2013 13:47:35	MG1	21.95 GB	IFINCR	A	vApp: VAPP_1 Org : HQ OvDC: HQ_Gold
2	09/07/2013 10:47:35	MG1	21.95 GB	IFINCR	I	vApp: VAPP_1 Org : HQ OvDC: HQ_Gold
3	09/07/2013 15:14:35	MG2	13.50 GB	IFFULL	A	vApp: WIN27 Org : Manufacturing OvDC: Manufacturing_fast
4	09/07/2013 09:14:35	MG2	13.50 GB	IFFULL	I	vApp: WIN27 Org : Manufacturing OvDC: Manufacturing_fast
5	09/07/2013 16:12:36	MG1	1.95 GB	IFFULL	A	vApp: SUSE10 Org : Sales OvDC: Sales_vDC2

More examples:

Display all vApps in all organization virtual data centers of the organization that is named Bucharest and the organization virtual data centers that are named Production and Development:

```
dsmc q vapp org=Bucharest,orgvdc=Production,Development
```

Display all vApps whose names begin with "Production" from the organization that is named Haifa and the organization virtual data center that is named Development:

```
dsmc q vapp org=Haifa,orgvdc=Development,vapp=Production*
```

Related reference:

“Asnodename” on page 319

“Detail” on page 353

“Inactive” on page 419

Query VM

Use the **query VM** command to list and verify the successful backups of virtual machines.

Query VM for VMware virtual machines

Use the **query vm** command to determine which VMware virtual machines were backed up.

Supported Clients

This command is valid on Linux clients that are installed on a vStorage backup server.

Syntax

►► Query VM — *vmname* — *options* ►►

Parameters

vmname

Specifies the virtual machine host name that you want to query. If you omit the virtual machine name, the command displays all VM backups on the Tivoli Storage Manager server.

Table 105. Query VM command: Related options for VMware virtual machine queries.

Option	Where to use
detail “Detail” on page 353 Valid for vmbackuptype=fullvm	Command line.
inactive “Inactive” on page 419 Valid for vmbackuptype=fullvm	Command line.
pitdate “Pitdate” on page 473 Valid for vmbackuptype=fullvm	Command line.
pittime “Pittime” on page 474 Valid for vmbackuptype=fullvm	Command line.
vmbackuptype “Vmbackuptype” on page 570	Command line or client options file.
vmchost “Vmchost” on page 570	Command line or client options file.
vmcpw “Vmcpw” on page 571	Command line or client options file.
vmcuser “Vmcuser” on page 573	Command line or client options file.

Query VM examples (VMware)

The following are samples of using the **query VM** command with full VM and file-level VM, and with full VM and the **-detail** option.

Full VM

```
QUERY
VM -INA -vmbackuptype=FULLVM
# Backup Date Mgmt Class Type A/I Virtual
Machine
1 09/26/2010 11:34:14 DEFAULT VMFULL A vm1
2 09/26/2010 11:34:37 DEFAULT VMFULL A vm2
3 09/26/2010 11:34:49 DEFAULT VMFULL A vm3
4 09/26/2010 12:35:08 DEFAULT VMFULL I vm2
5 09/26/2010 12:35:18 DEFAULT VMFULL I vm3
```

Full VM with -detail option

```
q vm Small-SQL -detail
IBM Tivoli Storage Manager
Command Line Backup-Archive Client Interface
Client Version 7, Release 1
Client date/time: 06/26/2014 14:08:44
(c) Copyright by IBM Corporation and other(s) 1990, 2014. All Rights
Reserved.

Node Name: BAClient-DM
Session established with server TSM_SERVER1: Windows
Server Version 7, Release 1, Level 0.0
Server date/time: 06/26/2012 14:37:39 Last access: 06/26/2014 14:22:55

Query Virtual Machine for Full VM backup

# Backup Date Mgmt Class Size Type A/I Virtual
Machine
- - - - -
1 06/26/2012 14:07:13 STANDARD 19.53 GB FULL A Small-SQL
The size of this incremental backup: n/a
The number of incremental backups since last full: n/a
The amount of extra data: n/a
The TSM objects fragmentation: n/a
Backup is represented by: n/a
Application protection type: TSM VSS
Application(s) protected: MS SQL 2008
DISK[1]Label: Hard disk 1
DISK[1]Name: [datastore1] Cloned-Small-SQL/Cloned-Small-SQL-000001.vmdk
DISK[1]Status: Protected
```

Query all VMware virtual machines that were backed up using
-vmbacktype=fullvm:

```
q vm * -vmbackuptype=fullvm
```

Related tasks:

“Preparing the environment for full backups of VMware virtual machines” on page 203

Restart Restore

The **restart restore** command displays a list of your restartable restore sessions in the server database.

You can restart only one restartable restore session at a time. Run the **restart restore** command again to restart further restores.

The restarted restore uses the same options that you used in the failed restore. The restarted restore continues from the point at which the restore previously failed.

To cancel restartable restore sessions, use the **cancel restore** command. Use the **restart restore** command when:

- Restartable restore sessions lock the file space at the server so that files cannot be moved off the sequential volumes of the server.
- You cannot back up files that are affected by the restartable restore.

Options from the failed session supersede new or changed options for the restarted session.

Supported Clients

This command is valid for all clients.

Syntax

►►—REStArT Restore—◄◄

Parameters

There are no parameters for this command.

Examples

Task Restart a restore.

Command: restart restore

Restore

The **restore** command obtains copies of backup versions of your files from a Tivoli Storage Manager server, or inside a backup set.

To restore files, specify the directories or selected files, or select the files from a list. Restore files to the directory from which you backed them up or to a different directory. The backup-archive client uses the **preservepath** option with the subtree value as the default for restoring files.

Note:

1. On UNIX and Linux systems when a symbolic link is created its modification time is set to the current system time and cannot be changed. So, when restoring a symbolic link its modification date and time is set to the date and time of the restore, not to the date and time the link had when it was backed up. As a result, the client backs up the symbolic link during the next incremental backup because its modification time changed since the last backup.

If you set the **subdir** option to yes when you restore a specific path and file, the client recursively restores all subdirectories under that path, and any instances of the specified file that exist under any of those subdirectories.

When you restore an entire directory or directory tree, and you do not specify the **inactive**, **latest**, **pick**, **today**, and **fromdate** options on the **restore** command, the client tracks which objects are restored. If the restore process is interrupted for any reason, you can restart the restore at the point of interruption by entering the **restart restore** command. It is possible to create more than one restartable restore session. Restores are only restartable if the file specification is fully wildcarded. For example, for a restore that is restartable, enter:

```
dsmc rest /home/* -sub=yes
```

For a restore that is not restartable, enter:

```
dsmc rest "/Users/user1/file?.c" -sub=yes
```


Use the **query restore** command to display a list of your restartable restore sessions in the server database. Further backups of the file system cannot be performed unless the restartable restore completes by using the **restart restore** command or is canceled by using the **cancel restore** command.

```
dsmc rest "/Users/user1/file?.c" -sub=yes
```

Supported Clients

This command is valid for all 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**.

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.

{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.

Specify the file space name if the drive label name has changed or if you are restoring files that were backed up from another node that had drive labels that are different from yours.

destinationfilespec

Specifies the path and file name where you want to place the restored files. If you do not specify a destination, the client restores the files to the original source path.

- If the *sourcefilespec* names a single file, the *destinationfilespec* can be a file or a directory. If you are restoring a single file, you can optionally end the specification with a file name if you want to give the restored file a new name.
- If the *sourcefilespec* is wildcarded or `subdir=yes` is specified, the *destinationfilespec* must be a directory and end with a directory delimiter (`\`).

Note: If the destination path or any part of it does not exist, the client creates it.

Note: If you do not specify a destination, the client determines whether the original file system can be reached. If the original file system cannot be reached, the client will not restore the file. In this case, you can specify a different destination and try the command again.

BACKUPSETName=

Specifies the name of a backup set. This parameter is optional. If you specify the **backupsetname** parameter with the **restore** command, you cannot use the **pick** option.

The value of **backupsetname** depends on the location of the backup set, and corresponds to one of the following options:

backupsetname

Specifies the name of the backup set from the Tivoli Storage Manager server. If the **location** parameter is specified, you must set `-location=server`. If the backup set resides in Tivoli Storage Manager server storage, the backup set must have a TOC.

localfilename

Specifies the file name of the first backup set volume. You must set `-location=file`.

tapedevice

Specifies the name of the tape device that contains the backup set volume. You must use a Windows-provided device driver, not the device driver that is provided by Tivoli. You must set `-location=tape`.

LOCation=

Specifies where the client searches for the backup set. If you do not specify the **location** parameter, the client searches for backup sets on the Tivoli Storage Manager server.

server Specifies that the client searches for the backup set from the server. This is the default location.

file Specifies that the client searches for the backup set from a local file.

tape Specifies that the client searches for the backup set from a local tape device.

Table 106. Restore command: Related options

Option	Where to use
dateformat "Dateformat" on page 344	Client user options file (dsm.opt) or command line.
dironly "Dironly" on page 357	Command line only.
filelist "Filelist" on page 402	Command line only.
filesonly "Filesonly" on page 406	Command line only.
followsymbolic "Followsymbolic" on page 407	Client user options file (dsm.opt) or command line.
fromdate "Fromdate" on page 409	Command line only.
fromnode "Fromnode" on page 410	Command line only.
fromowner "Fromowner" on page 411	Command line only.
fromtime "Fromtime" on page 412	Command line only.
ifnewer "Ifnewer" on page 417	Command line only.
inactive "Inactive" on page 419	Command line only.
latest "Latest" on page 445	Command line only.
numberformat "Numberformat" on page 465	Client user options file (dsm.opt) or command line.
pick Note: If you specify the backupsetname parameter with the restore command, you cannot use the pick option. "Pick" on page 472	Command line only.
pitdate "Pitdate" on page 473	Command line only.
pittime "Pittime" on page 474	Command line only.
preservepath "Preservepath" on page 480	Command line only.
replace "Replace" on page 489	Client user options file (dsm.opt) or command line.
subdir "Subdir" on page 539	Client user options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 541	Client user options file (dsm.opt) or command line.

Table 106. Restore command: Related options (continued)

Option	Where to use
timeformat "Timeformat" on page 549	Client user options file (dsm.opt) or command line.
today "Today" on page 554	Command line only.
totime "Totime" on page 554	Command line only.

Examples

Task Restore a single file named budget in the /Users/user1/Documents directory.

```
restore /home/devel/projecta/budget
```

Task Restore a single file named budget, which exists in the current directory.

```
restore file budget
```

Task Restore all files with a file extension of .c from the /home/devel/projecta directory.

```
restore "/home/devel/projecta/*.c"
```

Task Restore files in the /user/project directory. Use the pick and inactive options to select active and inactive backup versions.

```
restore "/user/project/*" -pick -inactive
```

Task Restore all files from the /home/devel/projecta directory that end with the character .c to the /home/newdevel/projectn/projecta directory. If the projectn or the projectn/projecta directory does not exist, it is created.

```
restore "/home/devel/projecta/*.c" /home/newdevel/projectn/
```

Task Restore all files in the /home/mydir directory to their state as of 1:00 PM on August 17, 2002.

```
restore -pitt=8/17/2002 -pitt=13:00:00 /home/mydir/
```

Task Restore all objects in the /home/myid/ directory. Since this restore operation is fully wildcarded, if the restore process is interrupted, a restartable restore session is created.

```
res "/home/myid/*"
```

Task Restore all files in the /home/mydir directory to their state as of 1:00 PM on August 17, 2002.

```
restore -pitt=8/17/2002 -pitt=13:00:00 /home/mydir/
```

Related information

"Restore data from a backup set" on page 224

"Preservepath" on page 480

"File system and ACL support" on page 153

Restore from file spaces that are not Unicode-enabled

If you want to restore from file spaces that are not Unicode-enabled, you must specify the source on the server and a destination on the client, prior to installing the Unicode-enabled client.

Note: This Unicode section applies only to Mac OS X.

For example, assume that Jaguar is the name of your startup disk and you back up all of the .log files in the /Users/user5/Documents directory. Before the backup takes place, the server renames the file space to Jaguar_OLD. The backup places the data specified in the current operation into the Unicode-enabled file space named /. The new Unicode-enabled file space now contains only the /Users/user5/Documents directory and the *.log files specified in the operation.

If you want to restore a file from the *renamed* (old) file space to its original location, you must enter both the source and destination as follows:

```
restore Jaguar_OLD/Users/user5/Documents  
/mylog.log /Users/user5/Documents/
```


Restore Backupset

The **restore backupset** command restores a backup set from the Tivoli Storage Manager server, a local file, or a local tape device. You can restore the entire backup set, or, in some cases, specific files within the backup set.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

{filespace name}

Specifies the file space (enclosed in braces) on the server that contains the files you want to restore. This is the name on the workstation drive from which the files were backed up, or the virtual file space name for a group.

Specify a file space name when you restore a backup set containing a group.

Specify a file space name when the *sourcefilespec* does not exist on the target computer. This can occur if the drive label name has changed or if you are restoring files that were backed up from another node that had drive labels that are different from yours.

sourcefilespec

Specifies the source path of a portion of the backup set. The default is to restore the entire backup set.

SYSTEMSTATE

Specifies that you want to restore the entire backup of the system state from the backup set. You cannot restore individual system state components from the backup set. This parameter is valid only for systems that support system state. To restore individual system state components from a backup set, use the **restore systemstate** command.

destinationfilespec

Specifies the destination path for the restored files. If you do not specify a *sourcefilespec*, you cannot specify a *destinationfilespec*. If you do not specify a destination, the backup-archive client restores the files to the original source path. If you are restoring more than one file, you must end the file specification with a directory delimiter (/), otherwise, the client assumes that the last name is a file name and reports an error. If you are restoring a single file, you can optionally end the destination file specification with a file name if you want to give the restored file a new name. When the *sourcefilespec* does not exist on the target workstation, you must specify *destinationfilespec*.

-BACKUPSETName=

Specifies the name of the backup set from which to perform a restore operation. You cannot use wildcard characters to specify the backup set name. The value of *backupsetname* depends on the location of the backup set, and corresponds to one of the following three choices:

backupsetname

Specifies the name of the backup set on the server from which to perform a restore operation. If **location** option is specified, you must set **-location=server**.

localfilename

Specifies the file name of the first backup set volume. You must set **-location=file**.

tapedevice

Specifies the name of the tape device containing the backup set volume. You must use a Windows-provided device driver, not the device driver that is provided by Tivoli. You must set **-location=tape**.

-LOCation=

Specifies the location of the backup set. If you do not specify the location parameter, the client searches for backup sets on the Tivoli Storage Manager server. If you specify the location parameter, the value must be one of the following three choices:

server Specifies that the backup set is on the Tivoli Storage Manager server. Server is the default location.

file Specifies that the backup set is on an available file system.

tape Specifies that the backup set is on an available tape device.

Table 107. Restore Backupset command: Related options

Option	Where to use
dironly "Dironly" on page 357	Command line only.
filesonly "Filesonly" on page 406	Command line only.
ifnewer "Ifnewer" on page 417	Command line only.
preservepath "Preservepath" on page 480	Command line only.
quiet "Quiet" on page 486	Client user options file (dsm.opt) or command line.
replace "Replace" on page 489	Client user options file (dsm.opt) or command line.
subdir "Subdir" on page 539	Client user options file (dsm.opt) or command line.

Examples

Task Restore the entire backup set called monthly_financial_data.87654321 from the server.

```
dsmc restore backupset
-backupsetname=monthly_financial_data.87654321
-loc=server
```

Task Restore the entire backup set contained in the file: /home/budget/weekly_budget_data.ost.

```
dsmc restore backupset
-backupsetname="/home/budget/weekly_budget_data.ost"
-loc=file
```

Task Restore the entire backup set from the /dev/rmt0 device.

```
dsmc restore backupset
-backupsetname=/dev/rmt0 -loc=tape
```

Task Restore a single file named /home/jones/budget.dev from the /dev/rmt0 tape device, to the original source path.

```
dsmc restore backupset
-backupsetname=/dev/rmt0 "/home/jones/budget.dev"
-loc=tape
```

Task Restore all files in the budget directory that contain a file extension of .txt from the tapes on the /dev/rmt0 device, to the original source path.

```
dsmc restore backupset "/home/budget/*.txt"
-backupsetname=/dev/rmt0 -loc=tape
```

Task Restore the entire backup set contained in local file named /home/jones/bset01.file"

```
dsmc restore backupset
-backupsetname="/home/jones/bset01.file"
-loc=file
```

Task Restore groups from the backup set mybackupset.12345678 on the Tivoli Storage Manager server to the /home/devel/projectb directory. The groups' virtual file space is accounting.

```
dsmc restore backupset {/accounting}/*
/home/devel/projectb/
-backupsetname=mybackupset.12345678 -loc=server
-subdir=yes
```

Task Restore groups from the local backup set mybackupset.ost to the /home/devel/projectb/ directory. The groups' virtual file space is accounting.

```
dsmc restore backupset {/accounting}/*
/home/devel/projectb/
-backupsetname=mybackupset.ost
-loc=server -subdir=yes
```

Related information

"Restore data from a backup set" on page 224

Restore backup sets: considerations and restrictions

This topic lists some considerations and restrictions that you must be aware of when restoring backup sets.

Backup set restore considerations

Consider the following when restoring backup sets:

- If the object you want to restore was generated from a client node whose name is different from your current node, specify the original node name with the **filespace** parameter on any of the restore commands.
- If you are unable to restore a backup set from portable media, check with your Tivoli Storage Manager administrator to ensure that the portable media was created on a device using a compatible format.
- If you use the **restore backupset** command on the initial command line with the parameter **-location=tape** or **-location=file**, the client does not attempt to contact the Tivoli Storage Manager server.
- When restoring a group from a backup set:
 - The entire group, or all groups, in the virtual file space are restored. You cannot restore a single group by specifying the group name, if there are several groups in the same virtual file space. You cannot restore a part of a group by specifying a file path.
 - Specify a group by using the following values:
 - Specify the virtual file space name with the **filespace** parameter.
 - Use the **subdir** option to include subdirectories.
- Limited support is provided for restoring backup sets from tape devices attached to the client system. A native device driver provided by the device manufacturer

must always be used. The device driver provided by Tivoli to be used with the Tivoli Storage Manager server cannot be used on the client system for restoring local backup sets.

- If a backup set contains files from several owners, the backup set itself is owned by the root user ID, and non-root user IDs cannot see the backup set. In this case, non-root user IDs can restore their files by obtaining the backup set name from the Tivoli Storage Manager administrator. Non-root users can restore only their own files.
- To enable the client GUI to restore a backup set from a local device, without requiring a server connection, use the `localbackupset` option.

Backup set restore restrictions

Be aware of the following restrictions when restoring backup sets:

- A backup set data that was backed up with the API cannot be restored or used.
- You cannot restore image data from a backup set using the **restore backupset** command. You can restore image data from a backup set only with the **restore image** command.
- You cannot restore image data from a local backup set (`location=tape` or `location=file`). You can restore image data from a backup set only from the Tivoli Storage Manager server.

Related reference:

“`Localbackupset`” on page 446

“**Restore**” on page 693

“**Restore Image**” on page 707

“**Restore Backupset**” on page 698

Restore backup sets in a SAN environment

You can restore backup sets in a storage area network (SAN) in the following ways:

- If the backup set is on a SAN-attached storage device, specify the device using the *filename* parameter and use the `location=tape` option, where applicable. The backup-archive client restores the backup set directly from the SAN-attached storage device, gaining high-speed restore performance.

Note: You must ensure that the correct tape is mounted in the SAN-attached tape drive prior to issuing the **restore** command. The backup-archive client will not initiate a SCSI autochanger to mount the tape automatically.

- If the backup set is not on local media or a SAN-attached storage device, you can specify the backup set using the `backupsetname` option. Use the `location=server` option to restore the backup set directly from the server using the LAN.

Restore Backupset without the backupsetname parameter

The **restore backupset** command can be used without the **backupsetname** parameter.


The preferred syntax for **restore backupset** command requires the **backupsetname** parameter. Before the introduction of the **backupsetname** parameter, the backup-archive client restored backup sets with a different syntax. The previous syntax is supported, but whenever possible, follow the syntax that requires the

backupsetname parameter. The previous syntax is documented for those cases when it cannot be replaced by the preferred syntax.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

options

All options that are valid with the preferred syntax of **restore backupset** are valid with the previous syntax of **restore backupset**.

{filespacename}

Specifies the file space (enclosed in braces) on the server that contains the files you want to restore. This is the name on the workstation drive from which the files were backed up, or the virtual file space name for a group.

Specify a file space name when you restore a backup set containing a group.

Specify a file space name when the *sourcefilespec* does not exist on the target computer. This can occur if the drive label name has changed or if you are restoring files that were backed up from another node that had drive labels that are different from yours.

sourcefilespec

Specifies the source path of a portion of the backup set. The default is to restore the entire backup set.

SYSTEMSTATE

Specifies that you want to restore the entire backup of the system state from the backup set. You cannot restore individual system state components from the backup set. This parameter is valid only for systems that support system state. To restore individual system state components from a backup set, use the **restore systemstate** command.

destinationfilespec

Specifies the destination path for the restored files. If you do not specify a *sourcefilespec*, you cannot specify a *destinationfilespec*. If you do not specify a destination, the client restores the files to the original source path. If you are restoring more than one file, you must end the file specification with a directory delimiter (/), otherwise, the client assumes that the last name is a file name and reports an error. If you are restoring a single file, you can optionally

end the destination file specification with a file name if you want to give the restored file a new name. When the *sourcefilespec* does not exist on the target workstation, you must specify the *destinationfilespec*.

backupsetname

Specifies the name of the backup set from the Tivoli Storage Manager server. If the **location** parameter is specified, you must set `-location=server`.

localfilename

Specifies the file name of the first backup set volume. You must set `-location=file`.

tapedevice

Specifies the name of the tape device containing the backup set volume. You must use a Windows-provided device driver, not the device driver that is provided by Tivoli. You must set `-location=tape`.

LOCation=

Specifies the location of the backup set. If you do not specify the location parameter, the client searches for backup sets on the Tivoli Storage Manager server. If you specify the location parameter, the value must be one of the following three choices:

server Specifies that the backup set is on the server. Server is the default location.

file Specifies that the backup set is on an available file system.

tape Specifies that the backup set is on an available tape device.

Examples

Task Restore the entire backup set called `monthly_financial_data.87654321` from the server.

```
dsmc restore backupset monthly_financial_data.87654321 -loc=server
```

Task Restore the entire backup set contained in the file: `/home/budget/weekly_budget_data.ost`.

```
dsmc restore backupset "/home/budget/weekly_budget_data.ost"
-loc=file
```

Task Restore the entire backup set from the `/dev/rmt0` device.

```
dsmc restore backupset "/dev/rmt0" -loc=tape
```

Task Restore a single file named `/home/jones/budget.dev` from the `/dev/rmt0` tape device, to the original source path.

```
dsmc restore backupset /dev/rmt0 "/home/jones/budget.dev" -loc=tape
```

Task Restore all files in the budget directory that contain a file extension of `.txt` from the tape(s) on the `/dev/rmt0` device, to the original source path.

```
dsmc restore backupset /dev/rmt0 "/home/budget/*.txt" -loc=tape
```

Task Restore the entire backup set contained in local file `/home/jones/bset01.file`

```
dsmc restore backupset "/home/jones/bset01.file" -loc=file
```

Task Restore groups from the backup set `mybackupset.12345678` on the Tivoli Storage Manager server to the `/home/devel/projectb` directory. The groups' virtual file space is accounting.

```
dsmc restore backupset mybackupset.12345678 {/accounting}/*  
/home/devel/projectb/ -loc=server -subdir=yes
```

Task Restore groups from the local backup set mybackupset.ost to the /home/devel/projectb/ directory. The groups' virtual file space is accounting.

```
dsmc restore backupset mybackupset.ost {/accounting}/*  
/home/devel/projectb/ -loc=server -subdir=yes
```

Related information

“Restore data from a backup set” on page 224

Restore Group

Use the **restore group** command to restore specific members or all members of a group backup.


Note:

- 1. Use the pick option to display a list of groups from which you can select one group to restore.
- 2. Use the showmembers option with the pick option to display and restore one or more members of a group. In this case, you first select the group from which you want to restore specific members, then you select one or more group members to restore.
- 3. You can restore a group from a backup set.

Supported Clients

This command is valid for all clients, except Mac OS X.

Syntax



Parameters

source
Specifies the virtual file space name and the group name on the server that you want to restore.

destination
Specifies the path where you want to place the group or one or more group members. If you do not specify a destination, the client restores the files to their original location.

Table 108. Restore Group command: Related options

Option	Where to use
backupsetname “Backupsetname” on page 329	Command line only.
followsymbolic “Followsymbolic” on page 407	Client options file (dsm.opt) or command line.

Table 108. Restore Group command: Related options (continued)

Option	Where to use
fromdate "Fromdate" on page 409	Command line only.
fromnode "Fromnode" on page 410	Command line only.
fromowner "Fromowner" on page 411	Command line only.
fromtime "Fromtime" on page 412	Command line only.
ifnewer "Ifnewer" on page 417	Command line only.
inactive "Inactive" on page 419	Command line only.
latest "Latest" on page 445	Command line only.
pick "Pick" on page 472	Command line only.
pitdate "Pitdate" on page 473	Command line only.
pittime "Pittime" on page 474	Command line only.
preservepath "Preservepath" on page 480	Command line only.
replace "Replace" on page 489	Client options file (dsm.opt) or command line.
showmembers "Showmembers" on page 517 (does not apply to Mac OS X)	Command line only.
subdir "Subdir" on page 539	Client user options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 541	Client user options file (dsm.opt) or command line.
todate "Todate" on page 554	Command line only.
totime "Totime" on page 554	Command line only.

Examples

Task Restore all members in the /virtfs/group1 group backup to their original location on the client system.

Command:

```
restore group /virtfs/group1
```

Task Display all groups within the /virtfs virtual file space. Use the showmembers option to display a list of group members from which you can select one or more to restore.

Command:

```
restore group /virtfs/  
* -pick -showmembers
```

Task Display a list of groups within the /virtfs virtual file space from which you can select one or more groups to restore.

Command:

```
restore group /virtfs/* -pick
```

Related information

“Restore Backupset” on page 698

Restore Image

The **restore image** command restores a file system or raw volume image that was backed up using the **backup image** command.

The restore obtains the backup image from a Tivoli Storage Manager server, or inside a backup set from the Tivoli Storage Manager server, when the **backupsetname** option is specified. This command can restore an active base image, or a point-in-time base image, with associated incremental updates.

Note:

1. Using the **incremental** option with the **restore image** command to perform a dynamic image backup is not supported.
2. If you use Tivoli Storage Manager HSM for Windows or Tivoli Storage Manager for Space Management, and you restore a file system image backup and plan to run reconciliation, you must restore the files that were backed up after the image backup. Otherwise, migrated files that were 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, the client issues a warning message on the console and in the error log.

If bad sectors are present on the target volume, you can use the **imagetofile** option with the **restore image** command to specify that you want to restore the source image to a file. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume.

Considerations:


- The API must be installed to use the **restore image** command.
- Image restore is not supported for the Sun QFS file system.
- Image restore is not supported for GPFS file systems on Linux x86_64, Linux on POWER and Linux on System z.
- On Linux systems, some file systems such as ext2, ext3, ext4, btrfs, and xfs use a universally unique identifier (UUID) to identify themselves to the operating system. If you create an image backup of such a volume and you restore it to a different location, you might have two volumes with the same UUID. If you use UUID to define your file systems in `/etc/fstab`, be aware that the backup-archive client might be unable to correctly mount the restored file system because the UUIDs conflict. To avoid this situation, restore the image to its original location. If you must restore it to a different location, change the UUID of either the original or restored volume before you mount the restored file system. Refer to the Linux documentation for instructions on how to change a UUID. You might also need to manually edit the `/etc/fstab` file so the original volume, the restored volume, or both volumes can be mounted.

- If you use the **pick** option, the following information is displayed for file system images that were backed up by the client:
 - Image Size
 - Stored Size - This value is the actual image size that is stored on the Tivoli Storage Manager server. The stored image on the server is the same size as the volume capacity.
 - File system type
 - Backup date and time
 - Management class that is assigned to image backups
 - Whether the image backup is an active or inactive copy
 - The image name
- If for some reason a restored image is corrupted, you can use the **fsck** tool to attempt to repair the image.

Supported Clients

This command is valid for AIX, HP-UX, all Linux clients, and Solaris.

Syntax



Parameters

sourcefilespec

Specifies the name of a source image file system to be restored. Only a single source image can be specified; you cannot use wildcard characters.

destinationfilespec

Specifies the name of an existing mounted file system or the path and file name to which the source file system is restored. The default is the original location of the file system.

Table 109. Restore Image command: Related options

Option	Where to use
backupsetname “Backupsetname” on page 329	Command line only.
dateformat “Dateformat” on page 344	Client user option file (dsm.opt) or command line.
deletefiles “Deletefiles” on page 351	Command line only.
fromnode “Fromnode” on page 410	Command line only.
fromowner “Fromowner” on page 411	Command line only.
imagnetofile “Imagnetofile” on page 418	Command line only.

Table 109. Restore Image command: Related options (continued)

Option	Where to use
inactive “Inactive” on page 419	Command line only.
incremental “Incremental” on page 438	Command line only.
noprompt “Noprompt” on page 464	Command line only.
pick “Pick” on page 472	Command line only.
pitdate “Pitdate” on page 473	Command line only.
pittime “Pittime” on page 474	Command line only.
timeformat “Timeformat” on page 549	Client user option file (dsm.opt) or command line.
verifyimage “Verifyimage” on page 565	Command line only.

The **restore image** command does not define or mount the destination file space. The destination volume must exist, must be large enough to hold the source, and if it contains a file system, must be mounted. If an image backup contains a file system, and you restore them to a different location, be aware of the following points:

- If the destination volume is smaller than the source volume, the operation fails.
- If the destination volume is larger than the source, after the restore operation you lose the difference between the sizes. The lost space can be recovered by increasing the size of the volume, which also increases the size of the restored volume.

Examples

Task Restore the /home/test directory over which the logical volume is mounted, to its original location.

Command: `dsmc rest image /home/test`

Task Restore the /home/proj directory over which the logical volume is mounted, to its original location and apply the changes from the last incremental backup of the original image that is recorded on the server. The changes include deletion of files.

Command: `dsmc restore image /home/proj -incremental -deletefiles`

Task Restore the /usr file system to its original location. Use the **verifyimage** option to enable detection of bad sectors on the target volume.

Command: `dsmc restore image /usr -verifyimage`

Task If bad sectors present on the target volume, use the **imagnetofile** option to restore the /usr file system to the /home/usr.img file to avoid data corruption.

Command: `dsmc restore image /usr /home/usr.img -imagnetofile`

Related information

“Verifyimage” on page 565

“Imagetofile” on page 418

Restore NAS

The **restore nas** command restores the image of a file system that belongs to a Network Attached Storage (NAS) file server. When you are using an interactive command-line session with a non-administrative ID, you are prompted for an administrator ID.

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 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 with the **RESTORE NODE** server command, if you know the fully qualified name of each file or directory and the image in which that object was backed up.

Use the **nasnodename** option to specify the node name for the NAS file server. The NAS node name identifies the NAS file server to the Tivoli Storage Manager server. You must register the NAS node name at the server. Place the **nasnodename** option in your client system-options file (**dsm.sys**). The value in the client system-options file is the default, but this value can be overridden on the command line.

You can use the **pick** option to display a list of NAS images that are 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 with the **pick** option, do not use the **monitor** option or you serialize the restores. To start multiple restore processes simultaneously when you are restoring multiple images, do not specify **monitor=yes**.

Use the **monitor** option to specify whether you want to monitor a NAS file system image restore and display processing information on your screen.

Use the **monitor process** command to display a list of current restore processes for all NAS nodes for which your administrative user ID has authority. The administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web.

Use the **cancel process** command to stop NAS restore processing.

Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: **/vol/vol0**.

Supported Clients

This command is valid for AIX, and Solaris clients only.

Syntax

```
➤—REStore NAS—┐—sourcefilespec—┐—destinationfilespec—➤  
                └—options—┘
```

Parameters

sourcefilespec

Specifies the name of the NAS file system image you want to restore. This parameter is required unless you use the pick option to display a list of NAS images from which to choose. You cannot use wildcard characters when you specify the *sourcefilespec*.

destinationfilespec

Specifies the name of an existing mounted file system on the NAS device over which you want to restore the image. This parameter is optional. The default is the original location of the file system on the NAS device.

Table 110. Restore NAS command: Related options

Option	Where to use
dateformat "Dateformat" on page 344	Client user option file (dsm.opt) or command line.
inactive "Inactive" on page 419	Command line only.
mode "Mode" on page 454	Command line only.
monitor "Monitor" on page 457	Command line only.
nasnodename "Nasnodename" on page 460	Client system options file (dsm.sys) or command line.
numberformat "Numberformat" on page 465	Client user option file (dsm.opt) or command line.
pick "Pick" on page 472	Command line only.
pitdate "Pitdate" on page 473	Command line only.
pittime "Pittime" on page 474	Command line only.
timeformat "Timeformat" on page 549	Client user option file (dsm.opt) or command line.

Examples

Task Restore the NAS file system image /vol/vol1 to the /vol/vol2 file system on the NAS file server called nas1.

Command: restore nas -nasnodename=nas1 /vol/vol1 /vol/vol2

Task Restore inactive NAS images.

Command: restore nas -nasnodename=nas2 -pick -inactive

Related information

"Nasnodename" on page 460

"Monitor" on page 457

"Cancel Process" on page 641

Restore VAPP

Use the **restore vapp** command to restore a single VMware vApp.

To use this command, you must have a license to use Tivoli Storage Manager for Virtual Environments V7.1, or later. Before you can use this command, the Tivoli Storage Manager nodes must be configured for the vCloud environment. In the Data Protection for VMware vCloud GUI, go to the Configuration window and click **Run Configuration Wizard**.

Supported Clients

This command is valid for Linux clients that are configured to perform vCloud Director vApp backups.

Syntax

```
►--Restore VAPP-- vapp_specification-----►
►--ASNODENAME=provider_vdc_node--VAPPName=new_vapp_name-----►
►--VCDHost=vCloud_server--VCDUser=vCloud_administrator--VCDPw=password-----►
►VMCHost=hostname--VMCUser=username--VMCPw=password-----►
```

Parameters

vapp_specification

Required parameter. Specifies the vApp that you want to restore.

The vApp specification must include an organization name and an organization virtual data center name and a single vApp name. The following parameters identify the vApp to restore:

ORG=*organization_name*

organization_name specifies the name of the organization that the organization virtual data center (**ORGVDC**) belongs to.

ORGVDC=*org_vdc_name*

Specifies the name of the organization virtual data center, within the organization, that contains the vApps to restore.

VAPP=*vapp_name*

Specifies a single vApp to restore. Wildcards and some special characters cannot be included in vApp name on this command. Specifically, you cannot include any of the following characters: in a vApp name: " ' : ; * ? , < > / \ |.

You must specify these keyword parameters in the following order: **ORG=**, **ORGVDC=**, **VAPP=**. If any value contains a space character, enclose the entire specification in quotation marks.

-ASNODENAME=provider_vdc_node

Required parameter. Specifies the node name of the Tivoli Storage Manager data mover node that is configured to protect vApps in a provider virtual data center. This parameter can be set on the command line or in the client options file (dsm.opt).

-VAPPName=new_vapp_name

Optional parameter. Specifies a new name for the restored vApp. You must specify a new vApp name if the vApp you are restoring exists in the inventory. The restore operation fails if the vApp exists and a new vApp name is not provided.

-VCDHost=vCloud_server

-VCDUser=vCloud_administrator

-VCDPw=password

Optional parameters for this command. These three parameters specify the following information:

- The host name or IP address of the vCloud Director server (**VCDHost**).
- The account name for a user on that server (**VCDUser**).
- The password that is associated with the user account (**VCDPw**).

These options are typically not specified on the command line. Instead, the server address is specified on the **VCDHost** option in the client options file. The administrator name and password are then associated with that host by a **set password** command, where you specify that host, the administrator name, and the administrator password. You also specify **TYPE=VCD** on the **set password** command to indicate that these credentials are for logging on to a vCloud server.

You can specify these parameters on the command line. However, passing these parameters on the command line does supply your login credentials in plain text. Adding these parameters on the command line should be done only as a temporary measure when you are troubleshooting logins.

-VMCHost=hostname

-VMCUser=username

-VMCPw=password

Optional parameters for this command. These three parameters specify the following information:

- The host name or IP address of the vCenter or ESXi server that hosts the vApps that you want to protect (**VMCHost**).
- The account name for a user on that server (**VMCUser**).
- The password that is associated with the user account (**VMCPw**).

These options are typically not specified on the command line. Instead, the server address is specified on the **VMCHost** option in the client options file. The administrator name and password are then associated with that host by a **set password** command, where you specify that host, the administrator name, and the administrator password. You also specify **TYPE=VM** on the **set password** command to indicate that these credentials are for logging on to a vCenter or ESXi server.

You can specify these parameters on the command line. However, passing these parameters on the command line does supply your login credentials in plain text. Adding these parameters on the command line should be done only as a temporary measure when you are troubleshooting logins.

Related options

Table 111. Restore VAPP command: Related options

Option	Where to use
asnodename	Client user-options file (dsm.opt) or client systems option file (dsm.sys), on the command line, or on the General tab in the Preferences editor.
inactive	Command line only.
pick	Command line only.
pitdate	Command line only.
pittime	Command line only.
vmchost	Client user-options file (dsm.opt) or client systems option file (dsm.sys), or on the command line. This option can also be included on the server in a client option set.
vmcpw	Specify this option on the command line or on a set password command where TYPE=VM is specified. This option can also be included on the server in a client option set.
vmcuser	Client user-options file (dsm.opt) or client systems option file (dsm.sys), or on the command line. You can also set this option on a set password command where TYPE=VM is specified. This option can also be included on the server in a client option set.

Example commands for restoring vApps

The following is an example command to restore a vApp that is named master_vApp2. This vApp has two virtual machines: VM1 and VM2.

```
dsmc restore vApp org=Orion,orgvdc=Development,vapp=master_vApp2
IBM Tivoli Storage Manager
Command Line Backup-Archive Client Interface
  Client Version 7, Release 1, Level 0.0
  Client date/time: 11/15/2013 00:34:06
(c) Copyright by IBM Corporation and other(s) 1990, 2013. All Rights Reserved.
```

```
Node Name: DM_pvdc1_DM1
Session established with server BORODINSERV: AIX
  Server Version 6, Release 3, Level 0.0
  Server date/time: 11/15/2013 09:27:49  Last access: 11/14/2013 14:26:10
```

```
Accessing as node: OVDC_ORG1_ORGVDC2
Restore function invoked.
Organization: Orion
Organization vDC: Orion_Silver
```

```
Restore VAPP command started.  Total number of vApps to process: 1
Total number of virtual machines to process: 2
```

```
Starting Full VAPP restore of VMware vApp 'master_vApp2'
```

```
Restoring vApp configuration.
```



```

Restore of Virtual Machine 'VM1' started

Starting Full VM restore of VMware Virtual Machine 'VM1'
  target node name='OVDC_Orion_DEVELOPMENT', data mover node name='DM_pvdc1_DM1'

... Creating vApp 'master_vApp2' by importing VM 'VM1'
Putting vApp into maintenance mode.  Restore of Virtual Machine 'VM2' started
Starting Full VM restore of VMware Virtual Machine 'VM2'
  target node name='OVDC_Orion_DEVELOPMENT', data mover node name='DM_pvdc1_DM1'

... Importing VM 'VM2'  Removing maintenance mode from vApp.
Restore processing finished.
Successful Full VAPP restore of VMware vApp 'master_vApp2'
  target node name='OVDC_Orion_DEVELOPMENT', data mover node name='DM_pvdc1_DM1'
Total number of objects restored:          1
Total number of objects failed:            0
Total number of bytes transferred:        4.50 GB
Data transfer time:                        316.78 sec
Network data transfer rate:                7,748.33 KB/sec
Aggregate data transfer rate:              5,972.31 KB/sec
Elapsed processing time:                   00:06:50
Restore VAPP command complete
Total number of vApps backed up successfully: 1
Total number of virtual machines backed up successfully: 2
Total number of vApps failed: 0
Total number of vApps processed: 1

```

Related reference:

“Asnodename” on page 319

“Pick” on page 472

“Pitdate” on page 473

“Pittime” on page 474

“Vmchost” on page 570

“Vmcpw” on page 571

“Vmcuser” on page 573

Restore VM

Use the **restore vm** command to restore a virtual machine that was previously backed up.

Restore VM for VMware virtual machines

The **Restore VM** command can be used to restore VMware virtual machines or VMware virtual machine templates.

If you have the backup-archive client installed on a separate system that is configured as a vStorage backup server, you can restore full virtual machine backups to the ESX or ESXi server that they came from, or to a different server. To restore a full virtual machine backup to a different server, use the **-host** option. The backup-archive client copies the data from the Tivoli Storage Manager server over either the LAN or SAN. The client then writes the data directly to the ESX server, by using the transport method that is specified in the client options file.

Restoring a full virtual machine backup creates a new virtual machine; the configuration information and content of the new machine is identical to what it was when the backup occurred. All virtual machine disks are restored to the specified point-in-time, as virtual disks in the newly created virtual machine.

When you restore a specific disk, by using the **:vmdk=** syntax, an existing virtual machine is updated with the specified virtual disk data. Only the specified disks are restored to the existing virtual machine; other disks in the virtual machine are not altered. The existing virtual machine that you are restoring the disk to must be powered off before you initiate the restore operation.

To create a new virtual machine, specify the **-vmname** parameter and provide a name for the new virtual machine. The **-vmname** parameter creates a new virtual machine with a configuration that is identical to what it was when the backup occurred. If you also specify the **:vmdk=** syntax, data is restored to any disks that are included in the **:vmdk=** parameters; disks that are not included are restored, but only as unformatted disks that do not contain data.


Virtual machines are restored to their original vApp, resource pool, cluster, or folder if the containers exist. During a restore operation, if the destination target (a vCenter or ESXi host) does not have the required containers, the virtual machine is restored to the top-level default location on the target ESXi host. If you use the command-line client to restore a virtual machine, and if the virtual machine cannot be restored to its original inventory location, an informational message (ANS2091I) is displayed. If you use the Java GUI to restore a virtual machine, and if the virtual machine cannot be restored to its original inventory location, the informational message is not displayed, but the virtual machine is still restored to the top-level default location.

Backup management tags that were backed up with the run **backup vm** command are restored with the virtual machine. Backup management tags are used to exclude virtual machines from backups and to specify the retention policy of backups.

Supported Clients

This command is valid on supported Linux clients that are installed on a vStorage backup server for a VMware virtual machine.

Syntax



Parameters

Any parameter that contains spaces must be enclosed in quotation marks (" ").

sourcevmspec

Specifies the name of the virtual machine (or virtual machine template) that was backed up.

VMName

This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.

Specifies the new name for the virtual machine after it is restored (if you do not want to use the name that is specified by *sourcevmspec*). You cannot use wildcards in the virtual machine name.

DATACENTER

This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.

Specifies the name of the data center to restore the virtual machine to as it is defined in the vSphere vCenter. If the data center is contained in a folder, you must specify the *-datacenter* option when you restore the virtual machine and include the folder structure of the data center in the data center name. For example, the following syntax is valid:

-datacenter=folder_name/datacenter_name

When you restore a virtual machine by using the GUI, you must restore the virtual machine to a different location. If you restore to the original location, you cannot specify the folder name of the data center. Without a folder name to help locate the original data center, the restore operation fails.

HOST

This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.

Specifies the domain name of the ESX host server to restore to as it is defined in the vSphere vCenter.

DATASTORE

This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.

Specifies the VMware datastore to restore the virtual machine to. The datastore can be on a SAN, NAS, iSCSI device, or VMware virtual volume (vVol). You can specify only one datastore when you restore a virtual machine. If you do not specify a **datastore** parameter, the virtual machine's VMDK file is restored to the datastore it was on when the backup was created.

:vmdk=all-vmdk

This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.

This option specifies that all virtual disks (*.vmdk files) are included when the virtual machine is restored. This is the default.

:vmdk=cnfg

This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.

This option specifies that the virtual machine configuration information is restored. The configuration information is always restored when a new virtual machine is created. However, by default the configuration is not restored when you update an existing virtual machine with selected virtual disks.

Ordinarily, restoring configuration information to an existing virtual machine fails because the restored configuration information conflicts with the existing virtual machine configuration information. Use this option if the existing configuration file for a virtual machine on the ESX server has been deleted, and you want to use the backed up configuration to re-create it.

:vmdk=disk label

This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.

This option is used to specify the disk label of the virtual disks to include in the restore operation. You specify this option only if you want to selectively restore data from specific disks.

Note: On the **Restore VM** command, the label names of the vmdk files that you want to include (:vmdk= parameter) in a **Restore VM** operation must be specified as the English-language label name, as it is displayed in the output of the Backup VM *vmname* -preview command. Examples of the English vmdk labels are "Hard Disk 1", "Hard Disk 2", and so on.

:-vmdk=disk label

This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.

This option is used to specify the disk label of one or more virtual disks to exclude from the restore operation.

Note: On the **Restore VM** command, the label names of the vmdk files that you want to exclude (:-vmdk= parameter) from a **Restore VM** operation must be specified as the English-language label name, as it is displayed in the output of the Backup VM *vmname* -preview command. Examples of the English vmdk labels are "Hard Disk 1", "Hard Disk 2", and so on.

Table 112. Restore VM command: Related options when restoring VMware virtual machines

Option	Where to use
datacenter	Command line or options file. This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.
datastore	Command line or options file. This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.
host	Command line or options file. This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.
inactive	Command line.
pick	Command line. This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.
pitdate	Command line. This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.
pittime	Command line. This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.

Table 112. Restore VM command: Related options when restoring VMware virtual machines (continued)

Option	Where to use
vmbackdir	Command line or client options file.
vmbackuptype	Command line or client options file.
vmdefaultdvportgroup	Command line or client options file
vmdefaultdvswitch	Command line or client options file
vmdefaultnetwork	Command line or client options file
vmstortransport	Command line or client options file. This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.

Examples

Task Restore the most recent backup version of myVM to its original name. Use the VMware management interface to delete the original virtual machine, before you restore it using this syntax.

```
dsmc restore vm myvm
```

Task Restore the most recent backup version of *myvm* to a new machine that is created with the name “Test Machine”, and with the restore target for the data center, ESX host, and datastore all specified on the command.

```
dsmc restore vm myvm -vmname="Test Machine"
    -datacenter="myDatacenter" -host="myHostName"
    -datastore="myDatastore"
```

Task Restore the most recent backup version of *myvm*. Restore to a data center named *mydatacenter*. The data center is within the vCenter; the relative path within the vCenter is *dirA/datacenters/*.

```
dsmc restore vm myvm -vmname="Test Machine"
    -datacenter="dirA/datacenters/myDatacenter"
    -host="myHostName" -datastore="myDatastore"
```

Task Restore a virtual machine template back to the same location and name.

```
dsmc restore vm vmTemplateName
```

Task Restore a virtual machine template to a new location.

```
dsmc restore vm vmTemplateName -vmname=newName
    -datastore=newDatastore -host=newHost
    -datacenter=newDatacenter
```

Task Restore only Hard Disk 2 and Hard Disk 3 to the existing virtual machine that is named vm1.

```
dsmc restore vm "vm1:vmdk=Hard Disk 2:vmdk=Hard Disk 3"
```

Task Restore all disks to the existing virtual machine named vm1, but do not restore the data from Hard Disk 4.

```
dsmc restore vm "vm1:-vmdk=Hard Disk 4"
```

Task Restore only the data from hard Disk 1 to the existing virtual machine vm1; do not update any configuration information.

Note: When you restore an existing virtual machine, the default behavior is to not update the configuration information.

```
dsmc restore vm "vm1:vmdk=Hard Disk 1:-vmdk=cnfg"
```

Task Restore all disks to the existing virtual machine named vm1.

```
dsmc restore vm "vm1:vmdk=all-vmdk"
```

This command updates all virtual disks on an existing virtual machine, named vm1. Note that this action is different from the action that is performed by `dsmc restore vm vm1`, which creates a new virtual machine named vm1 (vm1 must not exist in order for `dsmc restore vm vm1` to succeed).

Related concepts:

“Virtual machine exclude options” on page 394

“Virtual machine include options” on page 428

Related tasks:

“Preparing the environment for full backups of VMware virtual machines” on page 203

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 that are 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. The backup-archive client uses the **preservepath** option with the subtree value as the default for restoring files.

Note: When a directory is retrieved, its modification date and time is set to the date and time of the retrieval, not to the date and time the directory had when it was archived. This is because the backup-archive client retrieves the directories first, then adds the files to the directories.

Supported Clients

This command is valid for all clients.

Syntax

```
►► RETrieve [—options] [—sourcefilespec]
               [—destinationfilespec]
```

Parameters

sourcefilespec

Specifies the path and file name that you want to retrieve. Use wildcard characters to specify a group of files or all the files in a directory.

destinationfilespec

Specifies the path and file name where you want the files to be written. If you do not specify a destination, the client restores the files to the original source path.

Note: If you do not specify a destination, the backup-archive client determines whether the original file system can be reached. If the original file system cannot be reached, the client does not restore the file.

This failure can also occur if you remove the **virtualmountpoint** option from the `dsm.sys` file. In this case, you can specify a different destination, or restore the original **virtualmountpoint** option to the `dsm.sys` file, restart the client, and try the command again.

Table 113. Retrieve command: Related options

Option	Where to use
dateformat "Dateformat" on page 344	Client user options file (<code>dsm.opt</code>) or command line.
description "Description" on page 351	Command line only.
dironly "Dironly" on page 357	Command line only.
filelist "Filelist" on page 402	Command line only.
filesonly "Filesonly" on page 406	Command line only
followsymbolic "Followsymbolic" on page 407	Client user options file (<code>dsm.opt</code>) or command line.
fromdate "Fromdate" on page 409	Command line only
fromnode "Fromnode" on page 410	Command line only.
fromowner "Fromowner" on page 411	Command line only
fromtime "Fromtime" on page 412	Command line only
ifnewer "Ifnewer" on page 417	Command line only
pick "Pick" on page 472	Command line only.
preservepath "Preservepath" on page 480	Command line only.
replace "Replace" on page 489	Client user options file (<code>dsm.opt</code>) or command line.
subdir "Subdir" on page 539	Client user options file (<code>dsm.opt</code>) or command line.
tapeprompt "Tapeprompt" on page 541	Client user options file (<code>dsm.opt</code>) or command line.
timeformat "Timeformat" on page 549	Client user options file (<code>dsm.opt</code>) or command line.

Table 113. Retrieve command: Related options (continued)

Option	Where to use
todate "Todate" on page 554	Command line only.
totime "Totime" on page 554	Command line only.

Examples

Task Retrieve a single file named budget.

```
retrieve /home/devel/projecta/budget
```

Task Retrieve all files with an extension of .c from the /home/devel/projecta directory.

```
retrieve "/home/devel/projecta/*.c"
```

Task Retrieve all files in the /home directory.

```
retrieve /home/
```

Task Retrieve all files with a file extension of .c from the /home/devel/projecta directory to the /home/newdevel/projectn/projecta directory. If the /projectn or the /projectn/projecta directory does not exist, it is created.

```
retrieve "/home/devel/projecta/*.c" /home/newdevel/projectn/
```

Task Retrieve files in the /user/project directory. Use the **pick** option.

```
ret "/user/project/*" -pick
```

Task Retrieve all files that were archived from the /proj directory with the description "2012 survey results".

```
retrieve "/proj/*" -desc="2012 survey results"
```

Task Retrieve archived file /home/devel/budget with description "my budget" to the /dev/rmt1 tape drive.

```
mkfifo fifo
dd if=fifo of=/dev/rmt1&
dsmc retrieve -replace=yes -description="mybudget"
/home/devel/budget fifo
```

Task Retrieve a file from the renamed file space Jaguar_OLD to its original location. Enter both the source and destination as follows:

```
ret Jaguar_OLD/user5/Documents/myresume.doc /Users/user5/Documents/
```

Related information

"Client options reference" on page 316

Retrieve archives from file spaces that are not Unicode-enabled

If you want to retrieve archives from file spaces that were renamed by the Unicode-enabled client, you must specify the source on the server and a destination on the client.

This section applies to Mac OS X only. For example, assume that Jaguar is the name of your startup disk and you archive all of the .log files in the

/Users/user5/Documents directory. Before the archive takes place, the server renames the file space to Jaguar_OLD. The archive places the data specified in the current operation into the Unicode-enabled file space named /. The new Unicode-enabled file space now contains only the Users/user5/Documents directory and the *.log files specified in the operation.

If you want to retrieve a file from the *renamed* (old) file space to its original location, you must enter both the source and destination as follows:

```
retrieve Jaguar_OLD/Users/user5/Documents/mylog.log /Users/user5/Documents/
```

Schedule

The **schedule** command starts the client scheduler on your workstation. The client scheduler must be running before scheduled work can start.

Authorized User: The **schedule** command starts the client scheduler on your workstation. The client scheduler must be running before scheduled work can start.

Note:

1. The **schedule** command cannot be used if the `managedservices` option is set to `schedule` (does not apply to Mac OS X).
2. For Mac OSX only, to use the **schedule** command, specify `managedservices none` in the `dsm.sys` file.
3. This command is valid only on the initial command line. It is not valid in interactive mode or in a macro file.

If the `schedmode` option is set to `polling`, the client scheduler contacts the server for scheduled events at the hourly interval you specified with the `querschedperiod` option in your client user-options file (`dsm.opt`). If your administrator sets the `querschedperiod` option for all nodes, that setting overrides the client setting.

If you are using TCP/IP communications, the server can prompt your workstation when it is time to run a scheduled event. To do so, set the `schedmode` option to *prompted* in the client user-options file (`dsm.opt`) or on the **schedule** command.

You can use the `sessioninitiation` option with the **schedule** command to control whether the server or client initiates sessions through a firewall.

After you start the client scheduler, it continues to run and to start scheduled events until you press **Ctrl+C**, stop the scheduler process with the UNIX **kill** command, start the workstation again, or turn off the workstation to end it.

After you start the client scheduler, it continues to run and to start scheduled events until you press **Ctrl+C**, press the **Q** key twice, start the workstation again, or turn off the workstation to end it.

Note: You *cannot* enter this command in interactive mode.

Supported Clients

This command is valid for all clients.

Syntax

►—Schedule —options—►

Parameters

Table 114. Schedule command: Related options

Option	Where to use
maxcmdretries "Maxcmdretries" on page 450	Client system options file (dsm.sys) or command line.
password "Password" on page 468	client user options file (dsm.opt)
querschedperiod "Querschedperiod" on page 484	Client system options file (dsm.sys) or command line.
retryperiod "Retryperiod" on page 501	Client system options file (dsm.sys) or command line.
schedlogname "Schedlogname" on page 506	Client system options file (dsm.sys) or command line.
schedmode "Schedmode" on page 509	Client system options file (dsm.sys) or command line.
sessioninitiation "Sessioninitiation" on page 514	Client system options file (dsm.sys) or command line.
tcpclientport "Tcpclientport" on page 545	Client system options file (dsm.sys) or command line.

Examples

Task Start the client scheduler.

Command: dsmc sch -password=notell

Task For AIX: Start the scheduler at system bootup time by entering this command in the /etc/inittab file. Ensure that the *passwordaccess* option is set to *generate*.

Command: tsm::once:/usr/bin/dsmc sched > /dev/null 2>&1 #TSM

Task Interactively start the scheduler and keep it running in the background.

Command: nohup dsmc sched 2> /dev/null &

When you run the **schedule** command, all messages that regard scheduled work are sent to the dsm Sched.log file or to the file you specify with the schedlogname option in your client system-options file (dsm.sys). If you do not specify a directory path with the file name in the schedlogname option, the dsm Sched.log resides in the current working directory, except for Mac OS X. For Mac OS X, the dsm Sched.log resides in the /Library/Logs/tivoli/tsm/ directory.

Important: To prevent log write failures and process termination in certain cases, set the DSM_LOG environment variable to name a directory where default permissions allow the required access.

Related information

“Sessioninitiation” on page 514

Selective

The **selective** command backs up files that you specify. If you damage or mislay these files, you can replace them with backup versions from the server.

When you run a selective backup, all the files are candidates for backup unless you exclude them from backup, or they do not meet management class requirements for serialization.

During a selective backup, copies of the files are sent to the server even if they did not change since the last backup - which can 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 you back up a specific path and file, the client recursively backs up all subdirectories under that path, and any instances of the specified file that exist under any of those subdirectories.

During a selective backup, a directory path might be backed up, even if the specific file that was targeted for backup is not found. For example, the following command still backs up `dir1` and `dir2` even if the file `bogus.txt` does not exist.

```
selective /Users/user1/Documents/dir1/bogus.txt
selective "/dir1/dir2/bogus.txt"
```


If the **selective** command is retried because of a communication failure or session loss, the transfer statistics displays the number of bytes that the client attempts to transfer during *all* command attempts. Therefore, the statistics for bytes transferred might not match the file statistics, such as those for file size.

You can use the `removeoperandlimit` option to specify that the 20-operand limit is removed. If you specify the `removeoperandlimit` option with the **selective** command, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

filespec

Specifies the path and name of the file you want to back up. Use wildcard characters to include a group of files or to include all files in a directory.

To include multiple file specifications, separate each *filespec* with a space character. If multiple file specifications are included, and two or more of the specifications have common parent directories, then it is possible for the common directory objects to be backed up more than once. The conditions under which this behavior occurs are runtime-dependent, but the behavior itself has no adverse effects.

For example, if the *filespec* is `/home/amr/ice.doc /home/amr/fire.doc`, then `/home` and `/home/amr` might be backed up twice. The file objects, `ice.doc` and `fire.doc`, are backed up only once.

If you want to avoid including the shared parent directory more than once, use separate, non-overlapping **selective** commands to back up each file specification.

If you back up a file system, include a trailing slash (`/home/`).

There is a limit of 20 operands. This limit prevents excessive sessions that are caused when wildcards are expanded by the UNIX shell command processor. You can prevent shell expansion from causing you to go over the 20-operand limit by placing quotation marks around file specifications that contain wildcards (`"home/docs/*"`).

You can use the **removeoperandlimit** option to specify that the 20-operand limit is removed. If you specify the **removeoperandlimit** option, the 20-operand limit is not enforced and is restricted only by available resources or other operating system limits. For example, remove the 20 operand limit to backup 21 file specifications:

```
selective -removeoperandlimit filespec1 filespec2 ... filespec21
```

You can use the **filelist** option, instead of file specifications, to identify which files to include in this operation. However, these two methods are mutually exclusive. You cannot include file specification parameters and use the **filelist** option. If the **filelist** option is specified, any file specifications that are included are ignored.

Table 115. Selective command: Related options

Option	Where to use
changingretries "Changingretries" on page 333	Client system options file (dsm.sys) or command line.
compressalways "Compressalways" on page 338	Client user options file (dsm.opt) or command line.
compression "Compression" on page 339	Client user options file (dsm.opt) or command line.

Table 115. Selective command: Related options (continued)

Option	Where to use
dirsonly "Dirsonly" on page 357	Command line only.
filelist "Filelist" on page 402	Command line only.
filesonly "Filesonly" on page 406	Command line only.
preservelastaccessdate "Preservelastaccessdate" on page 479	Client user options file (dsm.opt) or command line.
removeoperandlimit "Removeoperandlimit" on page 488	Command line only.
snapshotcachesize "Snapshotcachesize" on page 525	Client options file (dsm.opt) or with the include.fs option.
snapshotproviderfs "Snapshotproviderfs" on page 526	System-options file (dsm.sys) within a server stanza or with the include.fs option.
snapshotroot "Snapshotroot" on page 528	Command line only.
subdir "Subdir" on page 539	Client user options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 541	Client user options file (dsm.opt) or command line.

Examples

Task Back up the proja file in the /home/devel directory.

Command: selective /home/devel/proja

Task Back up all files in the /home/devel directory whose file names begin with proj.

Command: selective "/home/devel/proj*"

Task Back up all files in the /home/devel directory whose file names begin with proj. Back up the single file that is named budget in the /user/home directory.

Command: selective "/home/devel/proj*" /user/home/budget

Task Back up the /home file system.

Command: selective /home/ -subdir=yes

Task Assuming that you initiated a snapshot of the /usr file system and mounted the snapshot as /snapshot/day1, run a selective backup of the /usr/dir1/sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name /usr.

Command: dsmc sel "/usr/dir1/sub1/*" -subdir=yes
-snapshotroot=/snapshot/day1

Associate a local snapshot with a server file space

Use the `snapshotroot` option with the **selective** command in conjunction with an independent software vendor application that provides a snapshot of a logical volume, to associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server. The `snapshotroot` option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot.

AIX only: You can perform a snapshot-based selective backup by specifying the option `snapshotproviderfs=JFS2`.

Set Access

The **set access** command gives users at other nodes access to your backup versions or archived copies.

You can also use the **set access** command to give users at other nodes access to your backup images.

You can give another user access to a specific file or image, multiple files or images, or all files in a directory. When you give access to another user, that user can restore or retrieve your objects. Specify in the command whether you are giving access to archives or backups.

For VMware virtual machines, you can give a user at another node access to the backups of a specific virtual machine.

For VMware vCloud vApps, you can give another user at another node access to backups of one or more VMware vApps. That user can then restore or retrieve a backed up vCloud vApp, including all of the virtual machines that comprise the vApp.

When a node is exported to another Tivoli Storage Manager server, the access rules can change on the importing server. If an access rule is applied to all file spaces on the exporting server, the access rule on the importing server is restricted to only those file spaces that are imported. The file spaces are restricted in the access rule on the importing server for security reasons. Additionally, the access rules do not recognize the first occurrence of a wildcard character in the file specification when you restore or retrieve. This means that if you restore or retrieve with a wildcard character in the file specification, subdirectories are ignored.

Tip: If you export a node to another Tivoli Storage Manager server, do not use a single wildcard character as the file specification in the access rule. Instead, create an access rule for each file space.


Note: You cannot give access to both archives and backups using a single command.

When an existing file space is renamed during Unicode conversion, any access rules that are defined for the file space remain applicable to the original file space. However, new access rules must be defined to apply to the new Unicode file space.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

Archive

Permits access to archived files or images.

Backup

Permits access to backup versions of files or images.

filespec

Specifies the path, file, image, or directory to which you are giving access to another node or user. Use wildcard characters to specify a group of files or images, or all files in a directory; all objects in a directory branch; or all objects in a file system. Use a single asterisk "*" for the file spec to give access to all files or images owned by you and backed up on the server. When the command set access backup "*" node is entered, no check is made with the server; it is assumed you have at least one object backed up.

If you give access to a branch of the current working directory, you only need to specify the branch. If you give access to objects that are not in a branch of the current working directory, you must specify the complete path. The file spec to which you gave access must have at least one backup version or archive copy object (file or directory) on the server.

To specify all files in a named directory, enter /home/mine/proj1/* on the command line.

To give access to all objects below a certain level, use an asterisk, directory delimiter, and an asterisk at the end of your file spec. For example, to give access to all objects below home/test, use file spec home/test/*/*.

Important: Use of the form /*/* alone will not give access to objects in the named directory; only those in directories below the named directory are accessible.

The rules are essentially the same when considering the root directory. Enter /* on one set access command and /*/* on another if you want another user to have access to all files and directories in and below the root directory. The first /* gives access to all directories and all files in the root directory. The second /*/* allows access to all directories and files below the root directory.

For example:

- Your directory structure is multilevel: /home/sub1/subsub1.
- The /home directory contains the h1.txt and h2.txt files.

- The /home/sub1 directory contains file s1.htm.
- The /home/sub1/sub2 directory contains the ss1.cpp file.

To allow access to all files in the /home/sub1/sub2 directory, enter:

```
set access backup /home/sub1/sub2/* * *
```

To allow access to only those files in the /home directory, enter:

```
set access backup /home/* * *
```

To allow access to all files in all directories in and below the /home directory, enter:

```
set access backup /home/* * *
set access backup /home/*/* * *
```

image-fs

The name of the image file system to be shared. This can be specified as an asterisk (*) to allow access to all images owned by the user granting access.

-TYPE=VM *vmname*

This parameter is required if you are using this command to provide another user with access to VMware virtual machine backups. The *vmname* option can be specified only if -TYPE=VM is specified; *vmname* is the name of the VMware virtual machine that you are permitting access to.

-TYPE=VAPP

This parameter is required if you are using this command to provide another user with access to vCloud Director vApps that were backed up by the **backup vapp** command.

To allow access to backed-up vApps, the command syntax depends on the client operating system:

```
set access backup -TYPE=VAPP "ORG=orgname,ORGVDC=orgvdcname,
VAPP=vappname" node user -ASNODENAME=providerVdcNode
```

The **ORG=** and **ORGVDC=** values specify the organization and organization virtual datacenter to which a vApp belongs.

The **VAPP=** value can specify an individual vApp name or you can specify an asterisk (*) character, or you can omit the name; either notation permits access by the node to all vApps in the specified organization and organization virtual datacenter. The specified node has access to all of the virtual machines that comprise the vApp.

The **-ASNODENAME=** parameter specifies the node that is associated with the vCloud provider that the organization virtual datacenter belongs to.

node

Specifies the client node of the user to whom you are giving access. Use wildcards to give access to more than one node with similar node names. Use an asterisk (*) to give access to all nodes.

user

This is an optional parameter that restricts access to the named user at the specified node. To allow any authorized user to access your backed up or archived data, specify **root** as the user.

Examples

Task Give the user named JONES on the node named NODE1 access to all of

the backed up vApps from the organization named XYZ and the provider virtual datacenter named VDC2. The provider virtual datacenter is mapped to a node called ABC.

```
set access backup -TYPE=VAPP "ORG=XYZ,ORGVDC=VDC2,VAPP=*"
NODE1 JONES -ASNODE=ABC
```

Task Give the user at node_2 authority to restore the budget file from the /home/user directory.

```
set access backup /home/user/budget node_2
```

Task Give node_3 the authority to retrieve all files in the /home/devel/proja directory.

```
set ac archive /home/devel/proja/ node_3
```

Task Give all nodes whose names end with bldgb the authority to restore all backup versions from directories with a file space name of project.

```
set ac b "{project}/*" "*bldgb"
```

Task Give any authorized user on node1 authority to retrieve all files in the /home/devel/projb directory.

```
set access archive /home/devel/projb/ node1 root
```

Task Give user serena at node_5 authority to restore all images of the file space mounted on directory /home/devel/proja.

```
set acc backup "home/devel/proja/*/*" node_5 serena
```

Task Give the node named **myOtherNode** the authority to restore files backed up by the VMware virtual machine named **myTestVM**.

```
set access backup -TYPE=VM myTestVM myOtherNode
```

Set Event

Using the **set event** command, you can specify the circumstances for when archived data is deleted.

You can use the **set event** command in the following ways:

- 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 to run when a particular event occurs (Notify the server that an event occurred)

Objects that are affected can be specified with a standard file specification (including wildcards), a list of files whose names are in the file that is 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 that match the filespec are affected. If you want to affect certain versions of a file, use the **-pick** option and select from the displayed list.

Interaction with down-level servers

If the **set event** command is issued when the client is connected to a server that does not support event-based policy (previous to Tivoli Storage Manager 5.2.2), the command is rejected with an error message that indicates the current server does not support event-based policy.

Supported Clients

This command is valid for all clients.

Syntax

```
► SET Event--TYPE= Hold Release Activeretention --<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.

activeretention

Signals the server that the controlling event occurred and starts to run the expiration clock.

-pick

Provides a list of objects from which the user can select to apply the event.

The following options can also be used and serve their usual purpose:

- *Dateformat*
- *Numberformat*
- *Noprompt*
- *Subdir*
- *Timeformat*

Examples

Task The following example displays the verbose and statistics output from the **set event** command `set event type=hold /home/accounting/ledgers/*05.books`, with objects rebound (as opposed to archived or some other notation).

```
Rebinding--> 274 /home/accounting/ledgers/
               jan05.books
Rebinding--> 290 /home/accounting/ledgers/
               feb05.books
```

```

Total number of objects archived:      0
Total number of objects failed:        0
Total number of objects rebound:      2
Total number of bytes transferred:    0 B
Data transfer time:                    0.00 sec
Network data transfer rate:            0.00 KB/sec
Aggregate data transfer rate:          0.00 KB/sec
Objects compressed by:                 0%
Elapsed processing time:                00:00:02

```

Task The `-pick` option used with the `set event` command `set event type=activate /user/tsm521/common/unix` shows the event type instead of the command name:

TSM Scrollable PICK Window - Retention Event : ACTIVATE

#	Archive Date/Time	File Size	File
1.	08/05/2003 08:47:46	766 B	/user/tsm521 /common/unix
2.	08/01/2003 10:38:11	766 B	/user/tsm521 /common/unix
3.	08/05/2003 08:47:46	5.79 KB	/user/tsm521 /common/unix
4.	08/01/2003 10:38:11	5.79 KB	/user/tsm521 /common/unix
5.	08/05/2003 08:47:46	10.18 KB	/user/tsm521 /common/unix

Related information

“Dateformat” on page 344

“Numberformat” on page 465

“Noprompt” on page 464

“Subdir” on page 539

“Timeformat” on page 549

Set Netappsvm

The **set netappsvm** command associates the logon credentials for a cluster management server, which are specified on the **set password** command, with a NetApp storage virtual machine, and the data storage virtual machine (SVM) name (data Vserver). You must enter this command before you can create a snapshot difference incremental backup of a clustered NetApp volume.

This command is typically entered only once. The parameters are stored and are reused the next time that you backup a clustered volume that is managed by the storage virtual machine. If you move an storage virtual machine to another cluster management server, you must reenter this command and specify the new cluster management server. If necessary, change the login credentials by using the **set password** command.

Supported clients

This command is valid for AIX and Linux backup-archive clients that complete snapshot difference backups of clustered-data ONTAP-c-mode file-server volumes.

Syntax

► SET NETAPPSVM *svm_hostname cms_hostname svm_name* ►
 └─ remove *svm_hostname* ─┘

Parameters

svm_hostname

Specifies the host name or IP address of the storage virtual machine that manages the volumes and logical interfaces (LIFs), for the volumes that you want to protect.

cms_hostname

Specifies the host name or IP address of the cluster management server. Specify the same host name that you specified for this cluster management server when you used the **set password** command to establish the login credentials.

svm_name

Specifies the name of the data SVM that manages the mounted volume. Contact the NetApp SVM administrator to obtain the data SVM name that is assigned to the virtual machine.

-remove *svm_hostname*

Disassociates the SVM from the cluster management server that it was previously associated with. Specify a SVM host-name

You can specify this parameter if you accidentally associated a storage virtual machine with a 7-mode file server. If you remove a 7-mode file server and then associate a cluster management server, set the logon credentials for the cluster management server by using the **set password** command.

Examples

Configure the credentials and access to a storage virtual machine:

```
set netappsvm svm_example.com cms_filer1.example.com svm_2
dsmc set password cms_filer1.example.com user_name password
```

Remove the associations that were created for the storage virtual machine:

```
set netappsvm -remove svm_example.com
```

Related tasks:

“Protecting clustered-data ONTAP NetApp file server volumes” on page 107

Set Password

The **set password** command changes the Tivoli Storage Manager password for your workstation, or sets the credentials that are used to access another server.

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.

Passwords can be up to 63 character in length. Password constraints vary, depending on where the passwords are stored and managed, and depending on the version of the Tivoli Storage Manager server that your client connects to.

If your Tivoli Storage Manager server is at version 6.3.3 or later, and if you use an LDAP directory server to authenticate passwords

Use any of the following characters to create a password:

a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9 . ! @ # \$ % ^ & * _ - + = ` ()
| { } [] : ; < > , ? / ~

Passwords are case-sensitive and are subject to more restrictions that can be imposed by LDAP policies.

If your Tivoli Storage Manager server is at version 6.3.3 or later, and if you do not use an LDAP directory server to authenticate passwords

Use any of the following characters to create a password:

a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9 . ! @ # \$ % ^ & * _ - + = ` ()
| { } [] : ; < > , ? / ~

Passwords are stored in the Tivoli Storage Manager server database and are not case-sensitive.

If your Tivoli Storage Manager server is earlier than version 6.3.3

Use any of the following characters to create a password:

a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9
_ - & + .

Passwords are stored in the Tivoli Storage Manager server database and are not case-sensitive.

Remember:

On the command line, enclose all parameters that contain one or more special characters in quotation marks. Without quotation marks, the special characters can be interpreted as shell escape characters, file redirection characters, or other characters that have significance to the operating system.

On AIX, HP-UX, Linux, Mac, and Solaris systems:

Enclose the command parameters in single quotation marks (').

Command line example:

```
dsmc set password -type=vmguest 'Win 2012 SQL'  
'tsml2dag\administrator' '7@#$$^&7'
```


Quotation marks are not required when you type a password with special characters in an options file.

Supported Clients

This command is valid for all clients.

Syntax

► SET Password oldpw—newpw



Parameters

oldpw

Specifies the current password for your workstation.

newpw

Specifies the new password for your workstation.

other_server other_user_id other_password

These three parameters specify the attributes that the client uses to access another server, such as a filer, an ESXi host, or a vCloud Director server.

other_server

Specifies the host name or IP address of the server that the client can access to protect files.

other_user_id

The user ID of an account on the server that the client uses to log on to the other server. The account must have the privileges that are necessary to perform the operations that are run after the user is logged on to the other server.

other_password

The password that is associated with the user ID on the other server.

TYPE

Specifies whether this password is for the backup-archive client or for another type of server.

Use `TYPE=TSM` to specify the password for your backup-archive client. The default type is `TYPE=TSM`.

Use `TYPE=FastBack`, on Linux and Windows clients, to store the Tivoli Storage Manager FastBack credentials that are required for mounting and dismounting the FastBack volumes on the Windows FastBack server or DR Hub server.

The password file on the vStorage backup server must have either the Windows administrator ID for the VMware virtual center system, or the UNIX user ID for a specific ESX server. For a Fastback proxy backup, the password file must contain the FastBack administrator ID and password. Here are some examples:

```
dsmc set password 192.0.2.24 admin admin 123 -type=fastback
```

```
dsmc set password 192.0.2.24 WORKGROUP:admin admin 123 -type=fastback
```

```
dsmc set password windserv administrator windpass4 -type=fastback
```

Important: You must define the user credentials that are required to mount and unmount FastBack volumes from a repository to the backup-archive client before you enter the backup-archive FastBack subcommand. Use the `fbserver` option to define the credentials.

Here is a brief description of the various configurations and credentials that you need:

- The backup-archive client is installed on a dedicated vStorage backup server. The client on the vStorage backup server must connect to multiple network share repositories.

Follow these steps for each of the network share repositories where the client is connected:

1. Configure the repository for remote network access from FastBack Manager. Refer to the Tivoli Storage Manager FastBack product documentation on IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS9NU9/welcome>.
This step establishes a domain name, a network share user ID, and a network share password to connect remotely to the repository.
2. On the backup-archive client workstation, manually enter the following command:

```
dsmc set password type=fastback FBServer domain:networkaccessuserid
networkaccesspassword
```

The *fbserver* option specifies the short host name of the FastBack server workstation. For a FastBack DR Hub, the *fbserver* option specifies the short name of the workstation where the DR Hub is installed.

Networkaccessuserid is either the Windows administrator ID or the FastBack administration password.

Domain is the domain name of the user ID.

Networkaccesspassword is either the Windows administrator ID or the FastBack administration password.

3. These credentials are retrieved based on the short host name that you specify with the *fbserver* option.

Use *TYPE=FILER*, on AIX, Linux, and Windows systems, to specify that this password is for snapshot difference operations on a file server.

For *TYPE=FILER*, you must specify a file server name, and the user ID and the password that is used to access the file server. For example: `dsmc set password -type=filer myfiler filerid filerpasswd`.

When you specify *TYPE=FILER*, the password is stored in the password (TSM.sth) file without validating that the password is valid. Passwords that are stored with *TYPE=FILER* can be shared between client nodes. For example, a password that is stored by *NODE_A* can be used by *NODE_B*. Only one set of credentials is stored per file server.

Use *TYPE=VCD* to set a password that is used to log on to a vCloud Director server. For example:

```
dsmc SET PASSWORD -type=VCD vcloud_host vcloud_administrator
vcloud_password
```

where:

vcloud_hostname

Specifies the host name or IP address of the vCloud Director host. This host name must match the host name syntax that is used in the ***vcdhost*** option. That is, if ***vcdhost*** uses an IP address instead of a host name, this command must provide the IP address, and not a short host name or a fully qualified host name.

vCloud_administrator

Specifies the vCloud system administrator account that is used to log on to the vCloud Director host.

password

Specifies the password that is associated with the vCloud system administrator login account.

Use TYPE=VM to set the password that is used to log on to an ESX or vCenter server.

```
dsmc SET PASSWORD -type=VM hostname administrator password
```

where:

hostname

Specifies the VMware VirtualCenter or ESX server that you want to back up, restore, or query. This host name must match the host name syntax that is used in the **vmhost** option. That is, if **vmhost** uses an IP address instead of a host name, this command must provide the IP address, and not a short host name or a fully qualified host name.

administrator

Specifies the account that is needed to log on to the vCenter or ESXi host.

password

Specifies the password that is associated with the login account that you specified for the vCenter or ESXi administrator.

Use the Preferences editor to set the vmhost, vmcuser, and vmcpw options.

You can also set the **vmhost** option in the client options file and then use the **set password** command to associate that host name with the administrator account and the administrative account password that is used to log on to that host. For example, set password TYPE=VM myvmhost.example.com administrator_name administrator_password.

Use TYPE=VMGUEST, on Linux and Windows clients, if you use the INCLUDE.VMTSMVSS option to protect a virtual machine. Use the following format for the **set password** command:

```
set password -type=vmguest guest_VM_name administrator password
```

where:

guest_VM_name

Specifies the name of the virtual machine guest that you want to protect.

administrator

Specifies the account that is needed to log on to the guest VM.

password

Specifies the password that is associated with the login account.

If you use the same credentials to log on to multiple virtual machines that are protected by the INCLUDE.VMTSMVSS option, you can set the password for the all of the virtual machines by specifying the **ALLVM** parameter. The **ALLVM** parameter causes the same credentials to be used when the client logs on to any guest that is included in an INCLUDE.VMTSMVSS option. The following commandTYPE=TSM is an example of how to use **ALLVM**. In this example, the user name "Administrator" and the password "Password" are used to log on to any virtual machine that you included on an INCLUDE.VMTSMVSS option:

```
set password -type=vmguest ALLVM Administrator Password
```


You can also set a combination of shared and individual credentials. For example, if most virtual machines in your environment use the same credentials, but a few virtual machines use different credentials, you can use multiple **set password** commands to specify the credentials. For example, assume that most virtual machines use "Administrator1" as the login name and "Password1" as the password. Assume also that one virtual machine, named VM2, uses "Administrator2" as the login name and "Password2" as the password. The following commands are used to set the credentials for this scenario:

```
set password -type=vmguest ALLVM Administrator1 Password1 (sets
credentials for most of the VMs).

set password -type=vmguest VM2 Administrator2 Password2 (sets unique
credentials for VM2).
```

Examples

The following examples use the **set password** command.

Task Change your password from osecret to nsecret.

```
set password osecret nsecret
```

Task Set up a user ID and password for the root user on the file server myFiler.example.com.

```
dsmc set password -type=filer myFiler.example.com root

Please enter password for user id "root@myFiler.example.com":
***** Re-enter the password for verification:***** ANS0302I
Successfully done.
```

Task Set up a user ID and password for the root user on the file server myFiler.example.com.

```
dsmc set password -type=filer myFiler.example.com root secret
```

Task Set up a user ID and password for the FastBack server myFastBackServer. Use the **-fbserver** option in the **archive fastback** and **backup fastback** commands for the server name.

```
dsmc set password -type=FASTBack myFastBackServer myUserId
'pa$password'
```

Important:

1. The `dsmc set password -type=fastback` command must be repeated on a dedicated client proxy workstation once for each FastBack repository where the backup-archive client is expected to connect.
2. For network share repositories, issue the `dsmc set password -type=fastback` command in this format: `dsmc set password -type=fastback myFBServer domainName:userId password`. The server name that is specified, which is `myFBServer` in this example, must match the name that you specify on the **-fbserver** option on a **backup fastback** or **archive fastback** command.
3. For a FastBack server or a Tivoli Storage Manager FastBack Hub, the user ID and password that are specified must have FastBack administrator privileges. You must issue the `dsmc set password -type=fastback` command once for each FastBack Server branch repository on the FastBack DR Hub that the backup-archive client is expected to connect to.

Task The backup-archive client is connecting to a FastBack server repository whose short host name is myFBServer. user ID is the FastBack network user ID that has read/write access to the repository share. DOMAIN is the domain to which the user ID belongs. myNetworkPass is the corresponding password for the user ID.

```
dsmc set password -type=fastback myFbServer DOMAIN:USERID
myNetworkPass
```

Task The backup-archive client is connecting to a repository on a DR Hub machine whose short host name is myFbDrHub. The user ID is the Windows administrator ID. DOMAIN is the domain to which the DR Hub machine belongs. myNetworkPass is the corresponding password for the administrator ID.

```
dsmc set password -type=fastback myFbDrHub DOMAIN:administrator
adminPasswd
```

Related reference:

“Snapdiff” on page 519

Set Vmtags

The **set vmtags** command creates backup management tags and categories that can be added to VMware virtual machines. You can manage Tivoli Storage Manager backups of virtual machines by specifying these tags with tools such as VMware vSphere PowerCLI version 5.5 R2 or later.

If you are using the IBM Data Protection extension to manage backups, you do not need to run the **set vmtags** command first. The tags and categories are created for you.

If you are writing scripts to apply these tags to virtual machines, you need only to issue the **set vmtags** command once so that backup management tags are created before they are added to the virtual machines.

When you run this command, the following categories and tags are created:

Category	Tag	Tag description
Backup Management (IBM Spectrum Protect)	Excluded	The object is excluded from backups by IBM Spectrum Protect
Management Class (IBM Spectrum Protect)	<i>Management_class_name</i>	The policy used for retention settings in IBM Spectrum Protect

Where:

Excluded

Excludes a virtual machine from scheduled backup services.

Management_class_name

The name of the retention policy that defines how long backup versions are kept on the Tivoli Storage Manager server or how many backup versions can exist on the server before they are expired.

If you do not use the *Management_class_name* tag value, the default retention policy for the VMware datacenter node is used.

Then, you can use scripts to apply these tags to virtual machines to exclude the virtual machines from scheduled backup services, or to specify the retention policy for the backups on the Tivoli Storage Manager server.

If the backup management tags already exist, running the **set vmtags** command does not create the tags again.

Requirements: Before you run the **set vmtags** command, ensure that the following requirements are met:

- VMware vCenter Server must be at Version 6.0 Update 1 or later.
- The vmchost option must be configured in the dsm.opt file on Windows clients or dsm.sys file on Linux clients. The user name and password that are associated with the vmchost value must also be set. If not already set, you can use the **dsmc set password** command to set the user name and password.

Supported clients

This command is valid only on supported Linux x86_64 clients that are installed on a vStorage backup server that protects VMware virtual machines.

Syntax

►►—SET VMTAGS—◄◄

Parameters

No parameters are required for this command.

Examples

Task Create backup management tags and categories that can be added to virtual machines:

```
dsmc set vmtags
```

Related concepts:

“Management classes and copy groups” on page 276

Related reference:

“Vmchost” on page 570

“Vmtagdatamover” on page 591

“Set Password” on page 734

Appendix. Accessibility features for the Tivoli Storage Manager product family

Accessibility features help users who have a disability, such as restricted mobility or limited vision to use information technology products successfully.

Accessibility features

The IBM Tivoli Storage Manager family of products includes the following accessibility features:

- Keyboard-only operation using standard operating-system conventions
- Interfaces that support assistive technology such as screen readers

The command-line interfaces of all products in the product family are accessible.

Tivoli Storage Manager Operations Center provides the following additional accessibility features when you use it with a Mozilla Firefox browser on a Microsoft Windows system:

- Screen magnifiers and content zooming
- High contrast mode

The Operations Center and the Tivoli Storage Manager server can be installed in console mode, which is accessible.

The Operations Center help system is enabled for accessibility. For more information, click the question mark icon on the help system menu bar.

Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

IBM and accessibility

See the IBM Human Ability and Accessibility Center (<http://www.ibm.com/able>) for information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

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

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
United States of America*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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 want 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
US*

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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

Portions of this code are derived from IBM® Corp. Sample Programs.

© Copyright IBM® Corp. _enter the year or years_. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights Except as expressly granted in this permission, no other permissions,

licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

A glossary is available with terms and definitions for the IBM Tivoli Storage Manager family of products.

See Tivoli Storage Manager glossary.

To view glossaries for other IBM products, see IBM Terminology.

Index

Special characters

* ? 219

Numerics

128-bit AES encryption support 152

256-bit AES encryption support 152

A

absolute mode 279

absolute option 316

access

permissions, archive 254

access control lists

extended permissions, backup 169

permissions, backup 169

accessibility features 743

ACL

file systems supported 153

active backup versions

displaying 143, 238, 668

restoring 238

administrative client

allowing secure sessions within a private network 542

afmskipuncachedfiles 317

AIX

configuring for snapshots 105

workload partition (WPAR)

backup 198

restore 235

AIX client

client components 3

communication methods 4

installing 11

uninstalling 13

AIX components

installable 3, 4

AIX disk space 3

AIX hardware requirements 3

AIX software requirements 3

AIX system requirements 3

application program interface (API)

setting 59

archive

a list of files 251

assign description to 351

assigning description on command line 251

associating local snapshot with server file space 251, 528

binding management class to 318

binding management classes to files 251

btrfs 189

command 612

compressing files 339

configuring client node proxy support 252

copy group attributes 277

copy mode 279

delete file spaces 174, 247

delete files after 351

delete individual archives from server file space 253, 643

archive (*continued*)

display the last modification date and last access

datecreation date 666

files only 560

grace period retention 276

hard links 255

how managed 257, 275

improving speed using share memory 289

include files for 421

information, query 666

list of files 402

more than one file specification 251

number of attempts to archive open files 333

only files; not directories 251

overriding management class during 282

primary tasks 249

process directories only (not files) 357

query user access 665

removeoperandlimit 488

retrieving using command line 256

running 249

shared data on multiple clients under a single node

name 252, 319

starting a web client session 140

subdirectories 251

summary of options 290

suppress confirmation prompt before deleting 464

symbolic links 254

using commands 250, 254

archive copy group 276

archive fastback

command 614

archive maximum file size 157

archmc option 318

archsymbk as file option 319

asnodename option 319

asnodename session settings 321

auditlogging option 321

auditlogname option 323

authentication

Tivoli Storage Manager client 135

authorization

options 304

authorized user

definition 47

tasks 47

authorized users

enabling encryption 50

authorizing

user to restore or retrieve your files 244

auto-update 325

autodeploy option 325

autofsrename option 326

automated client failover

configuration and use 87

configuring 90

determining the replication status 92

force failover 94

other components 90

overview 87

preventing 93

automated client failover (*continued*)

- requirements 88
- restoring data 229
- restrictions 88
- retrieving data 229
- testing the connection 94

automating backup services

- displaying scheduled work 266, 268
- options for 269
- process commands after backup 475
- process commands before backup 478
- starting client scheduler 63

automount option 328

B

back up

- btrfs 189
- network-attached storage (NAS) 626
- new or changed files 160
- number of attempts to back up open files 333
- parallel 579, 580, 581
- shared data on multiple clients under a single node
 - name 319
- symbolic links 214
- VM templates 578

back up hard links 216

back up open files 218

back up sparse files 217

back up volume 158

backing up 168

- in parallel sessions 205

backing up data 203

backmc option 329

backup

- copy mode 279
- displaying processing status 210
- Encrypted File Systems (EFS) 200
- grace period retention 276
- image 181
 - client domain 366
 - with incremental backup 624
- image: static, dynamic, snapshot 181
- improving speed using share memory 289
- incremental
 - associating local snapshot with server file space 660
- incremental-by-date
 - client command line 170
- multi-session, send files contiguously to the server 334
- one server session per file specification 334
- overview 147
- primary tasks 147
- process directories only (not files) 357
- query user access 665
- selective
 - associating local snapshot with server file space 728
- selective backup using client command line 170
- skip acl processing 517
- starting a web client session 140
- summary of options 290

backup chain integrity checks 594, 596

backup comparison: incremental, incremental-by-date 165

backup considerations 213

backup copy group 147, 276

- attributes 277

backup fastback command 616

backup files

- assigning management class 281

backup group command 618

backup image

- btrfs 189
- using DSM_DIR to point to plug-in library 57

backup image command 621

- supported devices 183

backup image, volume device type support 183

backup maximum file size 157

backup NAS

- using DSM_DIR to point to plug-in library 57

backup nas command 626

backup planning 147

backup set

- enabling GUI for local restore 224
- restore 221, 224
- restoring in a SAN environment 702

backup sets

- restore considerations 226, 701

backup vapp command 628

backup vm command 634

backup with client node proxy

- agent node 175
- overview 175
- target node 175

backup-archive client

- installation 10
- overview 1

backupset

- enabling GUI for local restore of 446

backupsetname option 330

basesnapshotname option 331

batch mode 606

- starting a session 137

bottom-up processing

- include-exclude list 125
- include-exclude options file 125

Bourne and Korn shell

- pointing to client user-options file 53

Bourne and Korn shell variables, setting 58

btrfs 221, 707

btrfs file system

- archiving 189
- backing up 189
- backup image 189
- protecting 190
- protecting subvolumes 191
- restore image 189
- restoring 189
- retrieving 189

C

C shell variables

- setting 59

c-mode 107

cadlistenonport option 332

cancel process command 641

cancel restore command 642

central scheduling

- summary of options 302

Certificate Authorities

- root certificates 72

Certificate Authorities 72

changingretries option 333

- CIFS
 - backing up file systems 197
- class option 333
- classic (standard) restore 242
- client
 - automatic update 2
 - client TCP/IP address other than the one for first server
 - contact 545
 - client TCP/IP port number other than the one for first server
 - contact 546
 - registering with server 113
 - setting password 114
 - size for the TCP/IP sliding window for client node 548
- client acceptor daemon
 - configuring to manage scheduler 62
 - manage scheduler, web client, or both 448
- client command options
 - overview 608
- client components
 - AIX client 3
 - Linux on Power Systems client 5
 - Linux on System z client 7
 - Linux x86_64 client 6
 - Mac OS X client 7
 - Oracle Solaris client 8
- client management service 46
- client node proxy
 - archive overview 252
 - backup 177
 - scheduling 177
 - support 252
- client options
 - display current settings 682
 - exclude
 - exclude.archive 116
 - exclude.attribute.symmlink 116
 - exclude.backup 116
 - exclude.compression 116
 - exclude.dir 116
 - exclude.file 116
 - exclude.file.backup 116
 - exclude.fs 116
 - exclude.image 116
 - guitreeviewafterbackup 413
 - order of processing (precedence) 307
 - overriding using command line 307
 - overview 608
 - using with commands 307
- client options file
 - creating and modifying 51
 - overview 50
 - required options for 51
- client options reference 316
- client scheduler
 - displaying scheduled work 266, 268
 - options for 269
 - run at startup 264
 - starting 63, 723
 - starting automatically 141
- client system options
 - exclude
 - exclude.archive 116
 - exclude.attribute.symmlink 116
 - exclude.backup 116
 - exclude.compression 116
 - exclude.dir 116
 - exclude.file 116
- client system options (*continued*)
 - exclude (*continued*)
 - exclude.file.backup 116
 - exclude.fs 116
- client system-options file
 - copying and modifying 51
 - specifying include-exclude options 115
- client user-options
 - customizing 55
- client user-options file
 - overriding using commands 307
- client-node proxy support 175
- client-server communication
 - client TCP/IP address other than the one for first server
 - contact 545
 - client TCP/IP port number other than the one for first server
 - server contact 546
 - establishing 51
 - identify your workstation to the server 462
 - maximum disk I/O buffer size client uses when reading or writing files 359
 - method 335
 - reconnection attempts after failure 337
 - reconnection interval after failure 337
 - size for the TCP/IP sliding window for client node 548
 - size of internal TCP/IP communication buffer 543
 - specifying default server 350
 - specifying name of server to contact for services 513
 - specifying number of kilobytes client buffers before sending transaction to server 555
 - TCP/IP address for dsmdad 544
 - TCP/IP address of Tivoli Storage Manager server 548
 - TCP/IP port address of Tivoli Storage Manager server 547
 - TCP/IP port address on which to establish shared memory connection 516
 - whether to send small transactions to server without buffering them first 546
- closed registration
 - permissions 114
 - using 114
- cluster environment
 - installing Tivoli Storage Manager 95
- clustered data ONTAP 107
- collecting diagnostic information 46
- collocatebyfilespec option 334
- command line
 - archiving files 254
 - assigning description to archive 251
 - display current settings for client options 682
 - displaying
 - processing status 210
 - ending a session 144
 - entering commands 608
 - general rules when entering options with commands 308
 - NAS file systems backup 195
 - overriding management class during archive 282
 - overview of parameters 608
 - performing image backup 188
 - performing large restore operations 241
 - performing point-in-time restore 233
 - restrictions for NAS file systems 192
 - retrieving archived files 256
 - return codes for operations 273
 - specifying file specification 609
 - specifying options file during session 467
 - starting a session 137

- command line (*continued*)
 - using wildcard characters 610
- command parameters
 - overview 608
- command processing, summary of options 304
- command session
 - ending 606
 - starting 606
- commands
 - archive 612
 - archive fastback 614
 - backup fastback 616
 - backup group 618
 - backup image 621
 - backup nas 626
 - backup vapp 628
 - backup vm 634
 - batch mode 606
 - cancel process 641
 - cancel restore 642
 - delete access 642
 - delete archive 643
 - delete backup 645
 - delete filespace 648
 - delete group 649
 - entering 608
 - entering on command line 608
 - expire 651
 - general rules when entering options with 308
 - help 652
 - incremental 654
 - interactive (loop) mode 607
 - loop 660
 - macro 661
 - maximum file specifications permitted 609
 - monitor process 662
 - overview of parameters 608
 - preview archive 663
 - preview backup 664
 - query access 665
 - query archive 666
 - query backup 668
 - query backupset 671, 672
 - query filespace 674
 - query group 676
 - query image 678
 - query inlexcl 679
 - query mgmtclass 680
 - query node 681
 - query options 682
 - query restore 684
 - query schedule 684
 - query session 685
 - query systeminfo 686
 - query vapp 687
 - query VM 691
 - restart restore 693
 - restore 693
 - restore backupset 698, 702
 - restore backupset considerations 226, 701
 - restore group 705
 - restore image 707
 - restore NAS 710
 - restore vapp 712
 - restore vm 715
 - retrieve 720
 - schedule 723

- commands (*continued*)
 - scheduled, enabling or disabling 269
 - selective backup 725
 - set access 728
 - set event 731
 - set netappsvm 733
 - set password 734
 - set vmtags 740
 - specifying file specification 609
 - using 603
 - using in executables 273
 - using in shell scripts 273
 - using options with 307
 - using wildcard characters 610
- commmethod option 335
- Common Internet File System (CIFS)
 - backing up file systems 197
- commrestartduration option 337
- commrestartinterval option 337
- communication methods
 - installable software 5, 6, 7, 8, 9
- Shared Memory
 - AIX client 4
 - HP-UX Itanium 2 client 5
 - Linux on System z client 7
 - Linux x86_64 client 6
 - Oracle Solaris client 9
- summary 288
- TCP/IP
 - AIX client 4
 - HP-UX Itanium 2 client 5
 - Linux on Power Systems client 5
 - Linux on System z client 7
 - Linux x86_64 client 6
 - Mac OS X client 8
 - Oracle Solaris client 9
- communications
 - establishing through firewall 66
 - establishing with Secure Sockets Layer (SSL) 68
- compressalways option 338
- compression
 - disabling processing 427
 - enabling processing 427
 - include-exclude statements 427
- compression and encryption processing
 - back up 427
 - exclude from backup 427
 - exclude options 427
- compression option 339
- compression processing
 - exclude from backup 427
 - exclude options 427
 - include files for 421
- concurrent backups 205
- configure the client for data deduplication 83
- configuring
 - client acceptor-managed scheduler 62
 - optional tasks 47
 - required tasks 47
 - the client scheduler 61
- configuring support for client node proxy backups 175
- configuring the web client 60
- console option 341
- containing quotation marks 138
- control files 568
- copy destination attribute 280
- copy frequency attribute 278

- copy group name attribute 277
- copy groups 276
 - archive 276
 - backup 276
- copy mode parameter
 - absolute 279
 - modified 279
- copy serialization attribute 279
- copy type attribute 278
- createnewbase 342
- createnewbase option 342

D

- data
 - restoring 240
- data deduplication 79
- data deduplication client configuration 83
- data deduplication files
 - exclude 85
- datacenter option 343
- datastore option 344
- date format
 - specifying 344
- dateformat option 344
- dedupcachepath option 348
- dedupcachesize option 348
- deduplication option 349
- default client user-options file
 - creating and modifying 53
 - example of 53
- default domain
 - excluding domains from backup 170, 361
- default management class 275
- default policy domain 275
- defaultserver option 350
- delete
 - file space 174, 247
 - NAS or client objects 333
- delete access command 642
- delete archive command 643
- delete backup command 645
- delete group command 649
- delete individual backups from server file space 173
- deleted file systems 218
- deletefiles option 351
- deleting
 - individual archives from server file space 253, 643
 - individual backups from server file space 645
- description option 351
- detail option 277, 353
- diagnostics
 - options 307
- diffsnapshot option 354
- diffsnapshotname option 355
- directories
 - assigning management class for 356
 - excluding 116
 - excluding from backup processing 389
 - incremental backup processing overview 160
 - processing during incremental-by-date 164
 - specifying on command line 609
- dirmc option 356
- dirsonly option 357
- disability 743
- disablenqr option 358
- disaster recovery 247

- disk recovery 247
- disk space requirements
 - client 3
- disk space, AIX 3
- disk space, HP-UX Itanium 2 5
- disk space, Linux on Power Systems 5
- disk space, Linux System z 7
- disk space, Linux x86_64 6
- disk space, Mac OS X 8
- disk space, Solaris 8
- diskbuffsize option 359
- diskcachelocation option 360
- displaying
 - archive information 666
 - online help 144
 - policy information 277
 - restartable restore sessions 684
 - scheduled events 684
 - session information 685
- domain
 - back up using the GUI 169
 - include for full vm backups 369
 - include for image backup 366
 - include for incremental backup 361
 - include for NAS image backup 366
 - incremental backup 216
 - specifying drives in the default 169
- domain option 361
- domain.image option 366
- domain.nas option 366
- domain.vcd option 368
- domain.vmfull option 369
- dontload option 374
- downloading maintenance updates 46
- DSM_CONFIG 57
 - adding to .cshrc file 59
 - pointing to client user-options file 53, 57
 - using on Solaris 57
- DSM_DIR
 - adding to .cshrc file 59
 - pointing to dsm.sys file 57
 - pointing to executable files 57
 - pointing to resource files 57
 - set for image or NAS backup or restore 57
- DSM_LOG
 - adding to .cshrc file 59
 - set to point to dsmerror.log, dsmwebcl.log, dsm Sched.log 57
- dsm.opt file
 - creating 53
 - creating and modifying 51
 - customizing 55
 - example of 53
 - required options for 51
 - specifying a drive specification using wildcards 121
- dsm.opt.smp file 53
- dsm.smp file
 - copying to dsm.opt 51
 - location 51
- dsm.sys file
 - creating 51
 - example of 51
- dsm.sys.smp file 51
- dsmerlog.pru file 385
- dsmerlog.log
 - set DSM_LOG to point to 57
- dsmerlog.log file 385

- DSMI_CONFIG environment variable
 - API, UNIX and Linux 59
- DSMI_DIR environment variable
 - API, UNIX and Linux 59
- DSMI_LOG environment variable
 - API, UNIX and Linux 59
- dsm Sched.log 504, 507
- dsm webcl.log 504, 507
- dual boot systems
 - duplicate names 159
- dynamic and shared serialization 279
- dynamicimage option 376

E

- efsdecrypt option 376
- enablearchiveretentionprotection option 377
- enablededupcache option 378
- enableinstrumentation option 380
- enablelanfree option 382
- enabling encryption for authorized users 50
- Encrypted File Systems (EFS)
 - backup file systems 200
 - restore file systems 234
- encrypting data during archive 152
- encrypting data during backup 152
- encryption
 - multiple clients under a single node name 319
 - of file data 152
 - saving encryption key password 384
- encryption considerations 50
- encryption processing
 - encryption methods available 152
 - excluding files from 389
 - include files for 421
 - query systeminfo command 686
- encryptiontype option 152, 383
- encryptkey option
 - encryptkey=generate 384
 - encryptkey=prompt
 - encryptkey=save 384
- enhanced query schedule 266
- enhanced query schedule command 684
- environment prerequisites
 - AIX client 3
 - HP-UX Itanium 2 client 4
 - Linux for zSeries client 7
 - Linux on Power Systems 5
 - Linux x86_64 client 6
 - Solaris client 8
- environment variables
 - DSM_CONFIG 57
 - DSM_DIR 57
 - DSM_LOG 57
 - LANG 57
 - setting API 59
 - setting Bourne and Korn shell 58
 - setting C shell 59
- error log 57
 - pruning 388
 - specifying path and file name 387
- error processing, summary of options 305
- errorlogmax option 385
- errorlogname option 387
- errorlogretention option 385, 388
- estimate function 169

- event logging
 - scheduler 268
- event-based policy retention protection
 - archive 285
 - backup 285
- exclude
 - EXCLUDE.VMDISK 394
- exclude data deduplication files 85
- exclude options 389
 - exclude.archive 116
 - exclude.attribute.sym link 116
 - exclude.backup 116
 - exclude.compression 116
 - exclude.dir 116
 - exclude.file 116
 - exclude.file.backup 116
 - exclude.fs 116
 - exclude.image 116
 - preview 124
 - processing 125
 - wildcard characters 121, 122
- exclude.image option 116
- EXCLUDE.VMDISK 394
- excluding files
 - system files 119
 - using wildcard characters 122
 - wildcard characters 121
- excluding files from backup services 152
- executable file
 - return codes from 273
- expire command 651
- extended permissions
 - archive 254

F

- failover
 - client 87
 - configuration and use 87
 - configuring the client 90
 - determining the replication status 92
 - disabling 93
 - other components 90
 - requirements 88
 - restore 229
 - restrictions 88
 - retrieve 229
- fbbranch option 396
- fbclient option 396
- fbpolicyname option 398
- fbreposlocation option 399
- fbserver option 400
- fbvolumename option 401
- file space
 - delete 174, 247, 648
 - determining fsID 353
 - excluding 116
 - NAS or client objects 333
 - performing an image backup 621
- file spaces 158
- file specification
 - maximum allowed on commands 609
- file systems
 - ACL support for 153
 - Btrfs 153
 - define virtual mount point for 566
 - deleted 218

- file systems (*continued*)
 - excluding from backup processing 389
 - GPFS, multinode cluster environment 153, 361
 - image backup of 181
 - QFS, restrictions 153
 - supported 153
- filelist option 402
- filename option 405
- files
 - archive a list of 251, 402
 - archive using commands 254
 - archived, overriding management class 282
 - archives, how managed 257
 - archiving 249, 612
 - archiving more than one file specification 251
 - assigning management classes 214
 - authorizing another user to restore or retrieve 244
 - back up hard-linked 216
 - back up open 218
 - binding management classes to 283
 - compressing during archive or backup 339
 - definition of changed 160
 - delete after archive 351
 - delete individual archives from server file space 253, 643
 - delete individual backups from server file space 645
 - encryption 152
 - excluding groups 121, 122
 - include-exclude
 - creating in Unicode format 421
 - including groups 121, 122
 - managing growth during compression 338
 - maximum file size for operations 157
 - performing large restore operations 241
 - processing include-exclude 125
 - query archive information 666
 - query backup information 668
 - query user access 665
 - renaming file spaces that are not Unicode to
 - Unicode-enabled 326, 725
 - restore hard-linked 216
 - restore or retrieve to another workstation 246
 - restore sparse 217
 - restore, using commands 244
 - restoring 239
 - restoring files belonging to another node 245
 - retrieve archived 256
 - retrieve using commands 256
 - retrieving files belonging to another node 245
 - sorting list of 143
- filesonly option 406
- firewall
 - establishing communications through 66, 414, 547
 - specifying TCP/IP ports for the web client 599
 - using web client through 599
 - whether server or client initiates sessions through 514
- folders
 - incremental backup processing overview 160
- followsymbolic option 407
- force incremental backup 316
- forcefailover option 409
- format
 - summary of options 304
- format and language
 - summary of options 304
- fromdate option 409
- fromnode option 410
- fromowner option 411

- fromtime option 412
- full backups, creating 204
- full incremental
 - comparing with incremental-by-date 165
 - comparing with journal-based, incremental-by-date 165
 - definition 160
 - description 160
 - when to use 165
- fuzzy backup 279

G

- getting started
 - changing your password 129
 - client scheduler 129
 - command-line session 129
 - displaying online help 129
 - ending a session 129
 - GUI session 129
 - sorting file lists 129
 - web client session 129
- GPFS file system
 - multinode cluster environment 153, 361
 - scheduling 179
 - storage pools 231
- graphical user interface
 - changing password 141
 - delete individual files or images from server file
 - space 645
 - displaying active and inactive backup versions 143, 238
 - displaying online help 144
 - displaying processing status 210
 - enabling for local backupset restore 446
 - enabling local backup set 224
 - ending a session 144
 - performing image backup 187
 - starting a session 136
 - using to back up objects 169
- group backup
 - display active and inactive objects 419
 - display all members of 517
 - overview 174
 - specify name of group 412
 - specify virtual file space name for 565
 - specifying full or differential 454
- groupname option 412
- GUI
 - ending a session 144
 - overriding management class during archive 282
 - performing point-in-time restore 233
 - starting a session 136
- guitreeviewafterbackup option 413

H

- hard links
 - archive and retrieve 255
 - back up 216
 - restore 216
- hard mounts, NFS 217
- hardware requirements, AIX 3
- hardware requirements, HP-UX Itanium 2 5
- hardware requirements, Linux on Power Systems 5
- hardware requirements, Linux System z 7
- hardware requirements, Linux x86_64 6
- hardware requirements, Mac OS X 8

- hardware requirements, Solaris 8
- help
 - displaying online 144
 - Internet resources 144
 - online forum 144
 - service and technical support 144
- help command 652
- host option 414
- HP-UX Itanium 2 client
 - communication methods 5
 - increasing default limit of data segment size 16
 - installing 14
 - uninstalling 16
- HP-UX Itanium 2 disk space 5
- HP-UX Itanium 2 hardware requirements 5
- HP-UX Itanium 2 software requirements 5
- HP-UX Itanium 2 system requirements 5
- httpport option 414

I

- IBM Knowledge Center xiv
- IBM PowerHA SystemMirror cluster
 - scheduling 178
- ieobjtype option 416
- ifnewer option 417
- image
 - restoring 221
 - using chkdsk to repair 221
 - using chkdsk tool to repair 707
 - using fsck to repair 221, 707
- image backup
 - considerations 182
 - deleting 645
 - excluding files from 389
 - include files for; assign management class to 421
 - include.dedup 421
 - incremental-by-date image backup 186
 - perform 181
 - point-in-time restore 624
 - revoke access 642
 - specifying selective or incremental 454
 - static, dynamic, snapshot 181
 - using command line 188
 - using the GUI 187
 - using with file system incremental 186
 - using with incremental-by-date 185
 - volume device type support 183
 - with incremental backup 185, 624
- image backup, considerations 182
- image to file
 - restoring 231
- imagegapsize option 417
- imagnetofile option 418
- inactive backup versions
 - displaying 143, 238, 668
 - restoring 238
- inactive option 419
- inclexcl option 420
- include option
 - management class 281
 - processing 125
 - wildcard characters 121, 122
- include VM templates in back ups 578
- include-exclude list 152
 - creating 115
 - preview 124
- include-exclude list (*continued*)
 - query order of processing 679
 - size restriction 125
- include-exclude options file
 - bottom-up processing 125
 - overview 152
 - specifying path and file name of 420
 - to manage archives 257
 - Unicode-enabled file spaces 420
- include-exclude processing
 - options for 116
 - overview 116
- include.vm option 430
- include.vmsnapshotattempts option 433
- include.vmtsmvss option 435
- incrbydate option 436
- incremental backup
 - associating local snapshot with server file space 528
 - back up new and changed files with modification date later than last backup 436
 - by date 170
 - client command line 170
 - client domain 361
 - command line 170
 - description 160
 - directories, processing overview 160
 - folders, overview 160
 - GPFS, multinode cluster environment 153, 361
 - memory-conserving algorithm 452
 - new and changed files 160
 - new and changed files with modification date later than last backup 436
 - of directories
 - processing overview 160
 - of folders
 - processing overview 160
 - optimizing memory during 152
 - overview 160
 - process a list of files 402
 - skip acl update checking 518
 - symbolic links 214
 - using client Java GUI 169
 - with image backup 185, 624
- incremental command 654
 - journal-based backup 657
- incremental option 438
- incremental-by-date
 - client command line 170
 - command line 170
 - comparing with incremental 165
 - comparing with incremental, journal-based 165
 - description 164
 - of directories
 - processing overview 164
 - overview 160
 - when to use 165
- incremental-by-date backup 164
 - using client Java GUI 169
 - using with image backup 185
- incremental, associating local snapshot with server file space 180
- input strings
 - containing blanks 138
- installation
 - backup-archive client 10
 - installation requirements
 - AIX client 3

installation requirements *(continued)*

- client 3
 - HP-UX Itanium 2 client 4
 - Linux for zSeries client 7
 - Linux on Power Systems 5
 - Linux x86_64 client 6
 - Solaris client 8

installation steps

- Mac OS X client 41
- Tivoli Storage Manager on Oracle Solaris clients 43

installing

- AIX client 11
- HP-UX Itanium 2 client 14
- Linux on Power Systems (Little Endian) client 22
- Linux on Power Systems client 17
- Linux on System z client 37
- Linux x86_64 client 28
- Mac OS X client 41
- overview 1
- Solaris client 43
- Ubuntu Linux on Power Systems (Little Endian) API 26
- Ubuntu Linux on Power Systems (Little Endian) client 26
- Ubuntu x86_64 client 33

installing the client management service 46

installing Tivoli Storage Manager

- cluster environment 95

instrlogmax option 438

instrlogname option 439

instrumentation log

- collecting performance information 380
- controlling the size 438
- specifying path and file name to store performance information 439

interactive mode 607

interactive session

- ending 660
- starting 138, 660
- using 660

J

Java GUI

- configuration restrictions 137

journal based backup

- restoring 164

journal based backups

- restoring 164

journal configuration file

- how to configure 73

journal daemon

- journal configuration file settings 73

journal database files

- errorlog 75
- journalldir 75

journal-based backup 162, 657

- comparing with incremental, incremental-by-date 165
- excluding directories 117
- excluding files 117
- include-exclude options
 - journal-based backup 117
- performing traditional full incremental, instead of 463, 657
- specifying configuration settings 73
- when to use 165

JournalSettings stanza 75

K

keyboard 743

Knowledge Center xiv

L

LAN-based image backup

- snapshot image backup 621

LAN-free data movement 382

- enabling communications for 149, 441, 443
- options 149
- prerequisites 149
- shared memory port for 442

lanfreecommmethod option 441

lanfreeshmport option 442

lanfreessl option 444

lanfreetcppport option 443

lanfreetcpsrveraddress option 444

LANG environment variable

- setting language locale 55

language locales

- supported 55

last access date

- specifying whether to update during backup or archive 160, 479

latest option 445

Linux Logical Volume Manager

- snapshot image backup of volumes 181

Linux on Power Systems (Little Endian) client

- installing 22
- uninstalling 25

Linux on Power Systems client

- client components 5
- communication methods 5
- installing 17
- uninstalling 21

Linux on Power Systems components

- installable 5

Linux on Power Systems disk space 5

Linux on Power Systems hardware requirements 5

Linux on Power Systems software requirements 5

Linux on Power Systems system requirements 5

Linux on System z client

- client components 7
- communication methods 7
- installing 37
- uninstalling 40

Linux on System z components

- installable 7

Linux System z disk space 7

Linux System z hardware requirements 7

Linux System z software requirements 7

Linux System z system requirements 7

Linux x86_64 client

- client components 6
- communication methods 6
- installing 28
- uninstalling 31, 35

Linux x86_64 components

- installable 6

Linux x86_64 disk space 6

Linux x86_64 hardware requirements 6

Linux x86_64 software requirements 6

Linux x86_64 system requirements 6

local backup set

- enabling GUI for local restore 224

- local snapshot
 - associating a local snapshot with a server file space 180
- localbackupset option 446
- log
 - See also* schedule log
 - controlling the size 438
 - DSM_LOG environment variable 387, 439, 506
 - error log, pruning 385
 - errorlogname option 387
 - errorlogretention option 387
 - instrlogmax option 438
 - intrlogname option 439
 - schedlogname option 506, 723
 - schedlogretention option 506, 723
 - specifying path and file name 387, 439, 506, 723
 - web client 504
- logical volume
 - image backup of 181
 - restoring 221, 231
- logs
 - dsmsched.log 507
 - dsmsched.pru 507
 - dsmwebcl.log 507
 - dsmwebcl.pru 507
 - truncating application logs 435
- loop command 660
- LVM
 - bring up an application after LVM starts 477
 - quiesce an application before LVM starts 483

M

- Mac OS X client
 - client components 7
 - communication methods 8
 - installation steps 41
 - installing 41
- Mac OS X components
 - installable 7
- Mac OS X disk space 8
- Mac OS X hardware requirements 8
- Mac OS X software requirements 8
- Mac OS X system requirements 8
- Macintosh client
 - environment prerequisites 7
 - installation requirements 7
- macro command 661
- maintenance 325
 - auto-update 2
- makesparsefile option 447
- manageservices option 448
- management class
 - assigning 214
- management classes
 - assigning to directories 282, 356
 - assigning to files 281
 - binding archive files to 251
 - binding to files 283
 - default 276
 - displaying 277
 - displaying information about 680
 - how Tivoli Storage Manager uses 152
 - overriding during archive processing 282
 - overriding the default 281
 - processing 281
 - questions to consider 280
 - selecting for files 280

- management classes (*continued*)
 - specifying with include option 281
 - using management class, example 281
- maxcmdretries option 450
- mbobjrefreshthresh 451
- mbpctrefreshthresh 451
- memory
 - optimizing when constrained 152
- memoryefficientbackup option 452
- messages
 - displaying on screen 564
 - stop displaying 486
- migrating backup-archive clients 1
- migration
 - web client 1
 - web client language files 1
- migration of file spaces to Unicode 159
- mode option 454
- mode parameter 278
- modes
 - batch 606
 - interactive (loop) 607
- modified mode 278
- monitor option 457
- monitor process command 662
- myreplicationserver option 458

N

- NAS
 - assigning management class to file systems 421
 - backing up file systems 192
 - deleting file spaces 174, 247, 648
 - query node command 681
 - restore file systems 236, 710
 - restore NAS command 710
 - specifying full or differential backup 454
- NAS file systems backup
 - command line 195
 - web client
 - GUI 194
- nasnodename option 460
- netapp file server 107
- Network Attached Storage (NAS)
 - backup file systems 192
- Network Attached Storage (NAS) file server
 - deleting file spaces 174, 247
- Network Data Management Protocol (NDMP) 9
- Network File System (NFS)
 - backup file systems 197
- network-attached storage (NAS)
 - display nodes for which admin ID has authority 681
- network-attached storage (NAS)
 - backup file systems 626
 - cancel backup and restore processes 641, 662
 - deleting file spaces 648
 - display file spaces on server 674
 - excluding files from backup 389
 - monitoring backup or restore operations 457
 - querying file system images belonging to 668
 - restore file systems 236, 710
 - specifying for query 557
 - specifying node name for operations 460
 - specifying whether to save table of contents for each file system backup 552
- new for backup-archive client V7.1.8 xvii

- NFS
 - backing up file systems 197
 - hard mounts 217
 - soft mounts 217
 - virtual mount points 214
- nfstimeout option 217, 461
- NLSPATH environment variable
 - displaying help browser menu in your language locale 55
 - to display help browser menu in your language locale 55
- no query restore 242
- node
 - specifying type to query 557
- node name 51
- Node name field 245
- nodename option 462
- nojournal option 463
- noprompt option 464
- nrtablepath option 464
- numberformat
 - specifying 465
- numberformat option 465

O

- online help
 - displaying 144
 - online forum 144
 - service and technical support 144
- open registration
 - permissions 115
 - using 115
- operating system requirements
 - clients 3
- optfile option 467
- options 535
 - absolute 316
 - afmskipuncachedfiles 317
 - archive, summary 290
 - archmc 318
 - archsmlinkasfile 319
 - asnodename 319
 - auditlogging 321
 - auditlogname 323
 - authorization options 304
 - autodeploy 325
 - autofsrename 326
 - automount 328
 - backmc 329
 - backup
 - excluding system state 389
 - backup, summary 290
 - backupsetname 330
 - basesnapshotname 331
 - cadlistenonport 332
 - central scheduling, summary 302
 - changingretries 333
 - class 333
 - collocatebyfilespec 334
 - command processing, summary 304
 - commmethod 335
 - commrestartduration 337
 - commrestartinterval 337
 - communication, summary 288
 - compressalways 338
 - compression 339
 - console 341
 - createnewbase 342
- options (*continued*)
 - datacenter 343
 - datastore 344
 - dateformat 344
 - dedupcachepath 348
 - dedupcachesize 348
 - deduplication 349
 - defaultserver 350
 - deletefiles 351
 - description 351
 - detail 353
 - diagnostics 307
 - diffsnapshot 354
 - diffsnapshotname 355
 - dirmc 356
 - dironly 357
 - disablenqr 358
 - diskbuffsize 359
 - diskcachelocation 360
 - domain 361
 - domain.image 366
 - domain.nas 366
 - domain.vcd 368
 - domain.vmfull 369
 - dontload 374
 - dynamicimage 376
 - efsdecrypt 376
 - enablearchiveretentionprotection 377
 - enablededupcache 378
 - enableinstrumentation 380
 - enablelanfree 382
 - encryptiontype 152, 383
 - encryptkey
 - encryptkey=generate 384
 - encryptkey=prompt 384
 - encryptkey=save 384
 - errorlogmax 385
 - errorlogname 387
 - errorlogretention 388
 - exclude
 - exclude.archive 116, 389
 - exclude.attribute.symblink 116, 389
 - exclude.backup 116, 389
 - exclude.compression 116, 389
 - exclude.dir 116, 389
 - exclude.encrypt 389
 - exclude.file 116, 389
 - exclude.file.backup 116, 389
 - exclude.fs 116, 389
 - exclude.fs.nas 389
 - exclude.image 116, 389
 - wildcard characters 121, 122
 - exclude.dedup 389
 - EXCLUDE.VMDISK 394
 - fbbranch 396
 - fbclient 396
 - fbpolicyname 398
 - fbreposlocation 399
 - fbserver 400
 - fbvolumename 401
 - filelist 402
 - filename 405
 - filesonly 406
 - followsymbolic 407
 - forcefailover 409
 - format and language, summary 304
 - format, summary 304

options (*continued*)

- fromdate 409
- fromnode 410
- fromowner 411
- fromtime 412
- general rules when entering with commands 308
- groupname 412
- guitreeviewafterbackup 413
- host 414
- httpport 414
- ieobjtype 416
- ifnewer 417
- imagegapsize 417
- imagetofile 418
- inactive 419
- inexcl 420
- include
 - wildcard characters 121, 122
- include.archive 421
- include.attribute.symlink 421
- include.backup 421
- include.compression 421
- include.encrypt 421
- include.file 421
- include.fs.nas 421
- include.image 421
- include.vm 430
- include.vmsnapshotattempts 433
- include.vmtsmvss 435
- incrbydate 436
- incremental 438
- instrlogmax 438
- instrlogname 439
- lanfreecommmethod 441
- lanfreeshmport 289, 442
- lanfreessl 444
- lanfreetcpport 443
- lanfreetcpserveraddress 444
- latest 445
- localbackupset 446
- makesparsefile 447
- managedservices 448
- maxcmdretries 450
- mbobjrefreshtresh 451
- mbpctrefreshtresh 451
- memoryefficientbackup 452
- mode 454
- monitor 457
- myreplicationserver 458
- nasnodename 460
- nfstimeout 461
- nodename 462
- nojournall 463
- noprompt 464
- nrtablepath 464
- numberformat 465
- optfile 467
- order of processing (precedence) 307
- password 468
- passwordaccess 470
- passworddir 471
- pick 472
- pitdate 473
- pittime 474
- postnschedulecmd 475
- postschedulecmd 475
- postsnapshotcmd 477

options (*continued*)

- prenschedulecmd 478
- preschedulecmd 478
- preservelastaccessdate 479
- preservepath 481
- presnapshotcmd 483
- queryschedperiod 484
- querysummary 485
- quiet 486
- quotesareliteral 487
- removeoperandlimit 488
- replace 489
- replserverguid 490
- replservername 492
- replsslport 494
- repltcpport 495
- repltcpserveraddress 497
- resourceutilization 499
- restore and retrieve, summary 299
- retryperiod 501
- revokeremoteaccess 502
- schedcmddisabled 503, 504
- schedcmduser (server defined only) 269
- schedlogmax 504
- schedlogname 506
- schedlogretention 507
- schedmode 509
- schedrestretrdisabled 510
- scrolllines 511
- scrollprompt 512
- servername 513
- sessioninitiation 514
- shmport 516
- showmembers 517
- skipacl 517
- skipaclupdatecheck 518
- snappdiff 106, 519
- snappdiffhttps 524
- snapshotcachesize 525
- snapshotproviderfs 526
- snapshotproviderimage 527
- snapshotroot 528
- specifying in commands 307
- srvoptsetencryptiondisabled 530
- srvprepostscheddisabled 531
- srvprepostsnapdisabled 532
- ssl 533
- sslacceptcertfromserv 534
- sslrequired 536
- stagingdirectory 538
- subdir 539
- system state
 - exclude from backup processing 389
- tapeprompt 541
- tcpadminport 542
- tcpbuffsize 543
- tcpcadaddress 544
- tcpclientaddress 545
- tcpclientport 546
- tcpnodelay 546
- tcpport 547
- tcpserveraddress 548
- tcpwindowsize 548
- timeformat 550
- toc 552
- todate 554
- totime 555

- options (*continued*)
 - transaction processing, summary 306
 - txnbytelimit 555
 - type 557
 - update time 557
 - useexistingbase 558
 - usereplicationfailover 559
 - v2archive 560
 - vappmc 561
 - vcdhost 562
 - vcdpw 562
 - vcduser 563
 - verbose 564
 - verifyimage 565
 - virtual machine exclude options 394
 - virtual machine include options 428
 - virtualfsname 565
 - virtualmountpoint 566
 - virtualnodename 567
 - vmbackdir 568
 - vmbackupmailboxhistory 569
 - vmbackuptype 570
 - vmchost 570
 - vmcpw 571
 - vmcuser 573
 - vmdatastorethreshold 574
 - vmdefaultdvportgroup 576
 - vmdefaultdvswitch 576
 - vmdefaultnetwork 577
 - vmenabletemplatebackups 578
 - vmlimitperdatastore 579
 - vmlimitperhost 580
 - vmmaxparallel 581
 - vmmaxvirtualdisks 583
 - vmmc 584
 - vmnoprdmdisks 584
 - vmnovrdmdisks 585
 - vmpreferdagpassive 586
 - vmprocessvmwithprdm 588
 - vmprocesswithindependent 587
 - mskipctlcompression 589
 - mskipmaxvirtualdisks 590
 - mskipmaxvmdks 591
 - vmtagdatamover 591
 - vmtimeout 598
 - vmverifyifaction 594
 - vmverifyiflatest 596
 - vmvstortransport 597
 - web client, summary 306
 - webports 599
 - wildcardsareliteral 600
- Oracle Solaris client
 - client components 8
 - communication methods 9
- Oracle Solaris components
 - installable 8

P

- parallel backups 205, 579, 580, 581
- parameters
 - yes and no, alternatives 316
- partial incremental
 - definition 160
 - incremental-by-date
 - running 170

- password
 - changing 141, 734
 - number of characters 141
 - setting 468
 - setting for client 114
 - specifying directory location to store encrypted password file 471
 - specifying whether to generate automatically or set as user prompt 470
 - using 137
 - valid characters 141
- password location 133
- password option 468
- password store 133
- passwordaccess option 470
- passwordddir option 471
- performance
 - improving speed of backups, restores, archives, retrieves 289
 - transaction options 306
 - transaction processing 555
- performing traditional full incremental backup 657
- permissions
 - access, saving standard and extended 254
- pick option 472
- PIHDW plugin 374
- pitdate 473
- pittime option 474
- plug-in library
 - for image or NAS backup or restore 57
- point-in-time restore
 - image backup 624
- policies, storage management 275
- policy domains
 - default policy domain 275
 - standard policy domain 275
- policy sets
 - active policy set 275
- portable media
 - restoring backup sets 224
- postnschedulecmd option 475
- postsnapshotcmd option 477
- Preferences editor
 - excluding domains from back up 169
- preschedulecmd option 478
- preschedulecmd option 478
- preservelastaccessdate option 479
- preservepath option 481
- Presnapshotcmd option 483
- preview
 - include-exclude list 124
- preview archive command 663
- preview backup command 664
- processing aliases 426
- processing options
 - authorization 304
 - backup and archive 290
 - central scheduling 302
 - communication 288
 - diagnostics 307
 - error processing 305
 - format 304
 - format and language 304
 - overview 287
 - restore and retrieve 299
 - specifying in commands 307
 - transaction processing 306

- processing options (*continued*)
 - using 129, 131, 287
 - web client 306
- processing symbolic links 426
- processing symbolic links and aliases 393
- processing time
 - estimating 169
- protecting Btrfs file systems 190
- protecting Btrfs subvolumes 191
- proxied session restrictions 176, 252
- publications xiv

Q

- QFS file system
 - restrictions 153
- query
 - amount of information that displays on screen 511
 - backups, establish point-in-time 473, 474
 - based on date and time of backup, archive 409, 412
 - description for 351
 - display active and inactive objects 419
 - files for another node 410
 - group
 - command 676
 - display members of 517
 - include-exclude list 679
 - NAS or client objects 333
 - nodes to which client has proxy authority 252
 - nodes to which client has proxy node authority 175
 - process directories only (not files) 357
 - scrolling preferences after displaying information on screen 512
 - system information 686
- query access command 665
- query archive command 666
- query backup command 668
- query backupset command 671, 672
- query filespace command 674
- query group command 676
- query image command 678
- query inclexcl command 679
- query mgmtclass command 277, 680
- query node command 681
- query options command 682
- query restore command 684
- query schedule
 - enhanced 266
- query schedule command 684
- query schedule command, enhanced 684
- query session command 685
- query systeminfo command 686
 - encryption processing 686
- query vapp command 687
- query VM command 691
- queryschedperiod option 484
- querysummary option 485
- quiesce applications 435
- quiet option 486
- quotesareliteral option 487

R

- raw logical volume
 - image backup of 181
 - restoring 221

- rebinding files to a different management class 284
- registering
 - client with server 113
 - using closed registration 114
 - using open registration 114, 115
- removeoperandlimit option 488
- replace option 489
- replserverguid option 490
- replservername option 492
- replsslport option 494
- repltcpport option 495
- repltcpserveraddress option 497
- resourceutilization option 499
- restart restore command 693
 - restart interrupted restore 244
- restartable restore 242
- restartable restore sessions, display 684
- restore 707
 - active version 238
 - authorizing another user 244
 - backup set
 - supported tape devices 698, 702
 - backup sets
 - overview 224
 - backups, establish point-in-time 473, 474
 - based on date and time of backup 409, 412
 - btrfs 189
 - classic (also known as standard) 242
 - create list of backup versions to 472
 - data using command line 240
 - disk 247
 - display active and inactive objects 419
 - during failover 229
 - enable SELinux 248
 - Encrypted File Systems (EFS) 234
 - files and directories 239
 - files belonging to another node 245
 - files for another node 410
 - files for another user 411
 - files to another workstation 246
 - from file spaces that are not Unicode-enabled 698
 - from portable media
 - overview 224
 - group
 - command 705
- GUI, displaying active and inactive versions 143
- image 221
 - considerations 707
 - enable detection of bad sectors on target volume 565
 - to a file 418
 - using chkdsk tool to repair 221
 - using DSM_DIR to point to plug-in library 57
 - using fsck tool to repair 221
- image to file 231
- image, suppress confirmation prompt 464
- improving speed using share memory 289
- inactive version 238
- list of files 402
- local backup set using the GUI 224
- logical volume 221, 231
- most recent backup version 445
- NAS
 - using DSM_DIR to point to plug-in library 57
- NAS file systems 236
 - command line 238
 - web client 236
- no query 242

- restore (*continued*)
 - non-root users on RHEL 5 248
 - overview 221
 - performing large operations 241
 - primary tasks 221
 - process directories only (not files) 357
 - processing status window 239
 - raw logical volume 221
 - replace existing file with latest backup 417
 - restartable 242
 - sorting file list 143
 - standard (also known as classic) 242
 - starting a web client session 140
 - summary of options 299
 - symbolic links
 - UNIX and Linux restrictions 693
 - to different workstation 567
 - using commands 244
 - using fsck tool to repair 707
 - whether to prompt before overwriting existing files 489
- restore backupset command 698, 702
- restore backupset command considerations 226, 701
- restore command 693
 - performing large operations 241
- restore group command 705
- restore hard links 216
- restore image
 - btrfs 189
- restore image command 707
- restore maximum file size 157
- restore NAS command 710
- restore sparse files 217
- restore vapp command 712
- restore vm command 715
- restoring point-in-time 233
 - using command line 233
 - using GUI 233
- restrictions
 - asnodename option 319
 - asnodename session settings 321
 - specifying full path with client acceptor daemon 467
 - within a proxied session 176, 252
- retain extra versions attribute 278
- retain only versions attribute 279
- retain versions attribute 280
- retention grace period
 - archive 276, 284
 - backup 276, 284
- retrieve
 - archive copies 256
 - archived files using commands 256
 - authorizing another user 244
 - based on date and time of archive 409, 412
 - btrfs 189
 - description for 351
 - during failover 229
 - files belonging to another node 245
 - files for another node 410
 - files to another workstation 246
 - hard links 255
 - improving speed using share memory 289
 - list of files 402
 - primary tasks 249
 - process directories only (not files) 357
 - replace existing file with latest archive if existing file is newer 417
 - running 256

- retrieve (*continued*)
 - sorting file list 143
 - starting a web client session 140
 - summary of options 299
 - symbolic links 254
 - to different workstation 567
 - whether to prompt before overwriting existing files 489
- retrieve command 720
- retrieve maximum file size 157
- retryperiod option 501
- return codes for operations 273
- revokeremoteaccess option 502
- root user
 - acquiring root user access 47
- root user tasks
 - creating default client user-options file 53
 - setting up 51
- running a snapshot difference backup
 - with HTTPS 168
- running a snapshot differential backup
 - with HTTPS 168

S

- SAN
 - restoring backup sets using 702
- schedcmddisabled option 503, 504
- schedcmduser option (server defined only) 269
- schedlogmax option 504
- schedlogname option 506
- schedlogretention option 507
- schedmode option 509
- schedrestretrdisabled option 510
- schedule command 723
- schedule log
 - controlling the size 504
 - specifying number of days to keep entries and whether to save pruned entries 507
 - specifying path and file name to store schedule log information 506
- scheduled (automated) backups
 - closing files before back up 218
 - displaying scheduled work 266, 268
 - options for 269
 - process commands after backup 475
 - process commands before backup 478
 - restart applications after back up 218
 - starting 63
- scheduled commands
 - enabling-disabling 269
- scheduled events, displaying 684
- scheduled services
 - defining schedules for UID other than zero 269
 - disabling scheduled commands 503, 504
 - restrictions for NAS file systems 192
- scheduler
 - configuring 61
 - displaying scheduled work 266, 268
 - event logging 268
 - managed by client acceptor daemon 448
 - number of hours between contacts to server for scheduled work 484
 - number of minutes between attempts to process scheduled commands 501
 - options for 269
 - polling mode or prompted mode 509
 - resolving memory retention after scheduled backups 448

- scheduler (*continued*)
 - starting 63
 - whether server or client initiates sessions through firewall 514
 - whether to disable execution of restore or retrieve operations 510
- scheduler comparison
 - CAD versus traditional scheduler 61
- scheduling
 - client node proxy 177
 - client node proxy backup 175
 - GPFS file system 179
 - IBM PowerHA SystemMirror cluster 178
- scrolllines option 511
- scrollprompt option 512
- Secure Sockets Layer (SSL)
 - establishing communications with 68
- selective backup 528, 725
 - associating local snapshot with server file space 180
 - client command line 170
 - command line 170
 - overview 160, 168, 170
 - symbolic links 214
 - using the client Java GUI 169
- selective command 725
- self-contained application protection 435
- serialization
 - copy serialization
 - dynamic 279
 - shared static 279
 - static 279
- server
 - communicating with 51
 - establishing communications through firewall 66
 - establishing communications with 51
 - establishing communications with Secure Sockets Layer (SSL) 68
 - identify to begin a stanza containing options for 513
 - specifying name of server to contact for services 513
 - TCP/IP address of Tivoli Storage Manager server 548
 - TCP/IP port address for 547
- server options
 - Sslfipsmode 536
- servername option 513
- service and technical support 144
- session information, displaying 685
- sessioninitiation option 514
- set access command 728
 - restore-retrieve authorization 244
- set DSM_DIR to point to 57
- set event command 731
- set netappsvm 107
- set password command 734
- set vmtags command 740
- setting
 - user privileges 140
- setting environment variables
 - API, UNIX and Linux
 - DSMI_CONFIG 59
 - DSMI_DIR 59
 - DSMI_LOG 59
- setting language locale 55
- setting up
 - required root user tasks 51
- shared dynamic serialization 279, 333
- shared memory communication method
 - options 289
- shared static serialization 279, 333
- shell scripts
 - return codes from 273
 - using commands in 273
- shmport option 516
- showmembers option 517
- skipacl option 517
- skipaclupdatecheck option 518
- snapdiff option 106, 519
- snapdiffhttps option 524
- snapshot difference 106, 519
 - with HTTPS 167
- snapshot differential backup
 - with HTTPS 167
- snapshot differential backup with HTTPS connection 524
- snapshot-differential-incremental backup 519
- snapshotcachesize option 525
- snapshotproviderfs option 526
- snapshotproviderimage option 527
- snapshotroot option 528
- snapshots
 - configuring 105
- soft mounts, NFS 217
- software requirements, AIX 3
- software requirements, HP-UX Itanium 2 5
- software requirements, Linux on Power Systems 5
- software requirements, Linux System z 7
- software requirements, Linux x86/x86_64 6
- software requirements, Mac OS X 8
- software requirements, Solaris 8
- Software updates 46
- Solaris client
 - installing 43
 - installing in Solaris zones 43
- Solaris disk space 8
- Solaris hardware requirements 8
- Solaris software requirements 8
- Solaris system requirements 8
- Solaris zones 168
- sparse files
 - back up 217
 - restore 217
 - specifying how to restore or retrieve 447
- special file systems 153, 214
- specifying whether to update last access date 479
- srvoptsetencryptiondisabled option 530
- srvprepostscheddisabled option 531
- srvprepostsnapdisabled option 532
- SSL (Secure Socket Layer)
 - establishing communications with 68, 71
- ssl option 533
- sslacceptcertfromserv option 534
- Sslfipsmode option 536
- sslrequired option 536
- stagingdirectory option 538
- standard (classic) restore 242
- standard management class
 - copy destination 280
 - copy frequency 278
 - copy group name 277
 - copy mode
 - absolute 279
 - modified 279
 - copy serialization 279
 - copy type 278
 - deduplicate data attribute 280
 - default values 277

- standard management class *(continued)*
 - retain extra versions 278
 - retain only version 279
 - retain versions 280
 - versions data deleted
 - active versions 278
 - inactive versions 278
 - versions data exists 278
- standard policy domain 275
- start the client scheduler at startup 264
- starting
 - automatically 141
 - overview 1
- starting a session
 - batch mode 137
 - interactive mode 138
- static serialization 279
- storage
 - displaying restartable restore sessions 684
- Storage Agent
 - for LAN-free data movement 149
 - using for LAN-free data movement 382
- storage area network
 - for LAN-free data movement 149
 - restoring backup sets using 382, 702
 - using for LAN-free data movement 382
- storage management policies 275
 - assigning management classes to files 214
 - copy groups 276
 - default management class 275
 - display on backup-archive client or web client GUI 214
 - include-exclude list 276
 - management classes 276
 - policy domains
 - default 275
 - standard 275
 - policy sets
 - active policy set 275
- storage pools
 - GPFS 231
- subdir option 539
- subdirectories
 - archive 251
 - include in backup 170
- support
 - gathering system information for 341, 405, 686
- supported language locales 55
- swing-enabled browser
 - necessary to run web client 140
- symbolic links
 - archiving and retrieving 254
 - back up 214
 - backing up 426
 - exclude from backup 426
 - exclude options 426
 - processing 426
 - restoring 407
 - UNIX and Linux restrictions 693
- symbolic links and aliases
 - backing up 393, 426
 - exclude from backup 393, 426
 - exclude options 393, 426
 - processing 393, 426
- syntax diagram
 - reading xiv
 - repeating values xiv
 - required choices xiv

- system files
 - excluding 119
- system information
 - gathering 341, 405
- system requirements, AIX 3
- system requirements, HP-UX Itanium 2 5
- system requirements, Linux on Power Systems 5
- system requirements, Linux System z 7
- system requirements, Linux x86_64 6
- system requirements, Mac OS X 8
- system requirements, Solaris 8
- system state
 - display active and inactive objects 419

T

- tagging support in VMware 591
- tapeprompt option 541
- tasks
 - assigning management classes to directories 282
 - closed registration 113
 - display management classes 277
 - GUI, override management class 282
 - open registration 113
 - password, change 141
 - root user 47
 - sessions, ending 141
- TCP/IP communication method
 - options 288
- tcpadminport option 542
- tcpbuffsize option 543
- tcpcadaddress option 544
- tcpclientaddress option 545
- tcpclientport option 546
- tcpnodelay option 546
- tcpserveraddress option 548
- tcpwindowssize option 548
- time format
 - specifying 550
- timeformat option 550
- Tivoli Storage Manager
 - client components
 - AIX client 3
 - Linux on Power Systems client 5
 - Linux on System z client 7
 - Linux x86_64 client 6
 - Mac OS X client 7
 - Oracle Solaris client 8
 - communication methods
 - AIX client 4
 - HP-UX Itanium 2 client 5
 - Linux on Power Systems client 5
 - Linux on System z client 7
 - Linux x86_64 client 6
 - Mac OS X client 8
 - Oracle Solaris client 9
 - installation requirements 7
 - online forum 145
 - password 137
 - upgrading from earlier versions of the product 1
- Tivoli Storage Manager client
 - authentication 135
- Tivoli Storage Manager FastBack configuration 94
- Tivoli Storage Manager FastBack data backup 209
- Tivoli Storage Manager FastBack data restore 209
- Tivoli Storage Manager FastBack installation requirements 9

- Tivoli Storage Manager on Mac OS X client
 - uninstalling 42
- Tivoli Storage Manager on Oracle Solaris clients
 - installation steps 43
 - uninstalling 45
- Tivoli Storage Manager password
 - using 136
- toc option 552
- todate option 554
- totime option 555
- traditional full incremental backup 162
- transaction processing
 - summary of options 306
 - txnbytelimit option 555
- TSM.sth file
 - HP-UX restriction 384
- tsmjbbd.ini
 - configuring 73
- txnbytelimit option 555
- type option 557

U

- Ubuntu Linux on Power Systems (Little Endian) API
 - installing 26
 - uninstalling 27
- Ubuntu Linux on Power Systems (Little Endian) client
 - installing 26
 - uninstalling 27
- Ubuntu x86_64 client
 - installing 33
- Unicode
 - migrating file spaces to 159
 - renaming file spaces that are not Unicode to
 - Unicode-enabled 326, 725
 - restore from file spaces that are not Unicode-enabled 698
- uninstalling
 - AIX client 13
 - HP-UX Itanium 2 client 16
 - Linux on Power Systems (Little Endian) client 25
 - Linux on Power Systems client 21
 - Linux on System z client 40
 - Linux x86_64 client 31, 35
 - Tivoli Storage Manager on Mac OS X client 42
 - Tivoli Storage Manager on Oracle Solaris clients 45
 - Ubuntu Linux on Power Systems (Little Endian) API 27
 - Ubuntu Linux on Power Systems (Little Endian) client 27
- UNIX
 - file systems, ACL support 153
 - saving standard access permissions 254
- UNIX and Linux
 - cluster environment 95
 - installing Tivoli Storage Manager 95
 - restrictions
 - restoring symbolic links 693
- updatectime option 557
- updating the client automatically 2
- upgrading backup-archive clients 1
- upgrading the backup-archive client from earlier versions of
 - the product 1
- useexistingbase option 558
- user privileges
 - setting 140
- usereplicationfailover option 559
- using multiple sessions 210

V

- v2archive option 560
- vappmc option 561
- vcdhost option 562
- vcdpw option 562
- vcduser option 563
- verbose option 564
- verifyimage option 565
- versions data
 - deleted attribute 278
 - deleted parameter 278
 - exists attribute 278
 - exists parameter 278
- virtual machine
 - exclude options 394
 - include options 428
- virtual mount point, setting 169
- virtualfsname option 565
- virtualmountpoint option 566
- virtualnodename option 567
 - restore or retrieve to another workstation 246
- VM 201
- vmbackdir option 568
- vmbackupmailboxhistory 569
- vmbackuptype option 570, 584
- vmchost option 570
- vmcpw option 571
- vmctlmc option
 - options
 - vmctlmc 572
- vmcuser option 573
- vmdatastorethreshold
 - option 574
- vmdefaultdvportgroup option 576
- vmdefaultdvswitch option 576
- vmdefaultnetwork option 577
- vmenabletemplatebackups option 578
- vmlimitperdatastore option 579
- vmlimitperhost option 580
- vmmaxparallel option 581
- vmmaxvirtualdisks 583
- vmnoprdmdisks 584
- vmnovrdmdisks 585
- vmpreferdagpassive option 586
- vmprocessvmwithprdm 588
- vmprocesswithindependent 587
- vmskipctlcompression option 589
- vmskipmaxvirtualdisks 590
- vmskipmaxvmdks 591
- vmtagdatamover
 - option 591
- vmtimeout option 598
- vmverifyifaction 594
- vmverifyiflatest 596
- vmvstortransport option 597
- VMware virtual machine backups 204
 - types 201
- volume label
 - duplicate names 158
- volume name 158
- volume naming precautions
 - dual boot systems 159
 - UNIX mount point 158
- vStorage backup server
 - off-host backup 203

W

web client

- configuration overview 59
- configuring 60
- enable to run in a swing-enabled browser 140
- establishing communications through firewall 66, 414
- GUI 194
- NAS
 - file systems backup 194
- restrict administrator from accessing client running web client 502
- restrictions for NAS file systems 192
- specifying TCP/IP port address for 414
- starting 140
- summary of options 306
- supported browsers 140
- unsupported functions 147
- using through a firewall 599

web client configuration overview 59

webports option 599

wildcard characters

- guidelines 610
- include or exclude files 120
- include or exclude groups of files 121
- specifying a drive specification in dsm.opt 121
- to include or exclude groups of files 122
- using with commands 219
- using with file specifications 219

wildcardsareliteral option 600

Z

Zettabyte file systems (ZFS) 199

- backup file systems 199



Product Number: 5608-E01
5608-E02
5608-E03

Printed in USA