

IBM Tivoli Storage Manager for Windows
Backup-Archive Clients
Version 7.1.6

Installation and User's Guide



IBM Tivoli Storage Manager for Windows
Backup-Archive Clients
Version 7.1.6

Installation and User's Guide



Note:

Before you use this information and the product it supports, read the information in “Notices” on page 793.

Seventh edition (June 2016)

This edition applies to version 7, release 1, modification 6 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, 2016.

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

Contents

Tables	xi	Configuring the language for displaying the Java GUI	33
About this publication	xiii	Web client configuration overview	33
Who should read this publication	xiii	Configuring the web client on Windows systems	34
Publications	xiii	Configuring the scheduler	36
Conventions used in this publication	xiv	Comparison between client acceptor-managed services and traditional scheduler services	36
Reading syntax diagrams	xiv	Configuring the client to use the client acceptor service to manage the scheduler	37
What's new for Version 7.1.6	xvii	Starting the client scheduler	38
Chapter 1. Installing the Tivoli Storage Manager backup-archive clients	1	Scheduling events using the GUI	39
Upgrading the backup-archive client	1	Configuring Tivoli Storage Manager client/server communication across a firewall	39
Upgrade path for clients and servers	1	Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer	42
Considerations for migrating between processor architectures x32 and x64	1	Certificate Authorities root certificates	44
Additional upgrade information	2	Configure your system for journal-based backup	45
Automatic backup-archive client deployment	3	Configuring the journal engine service	45
Client environment requirements	5	JournalSettings stanza (Windows)	47
Windows client environment requirements	5	JournalExcludeList stanza	49
Windows client installable components	6	JournaledFileSystemSettings stanza	49
System requirements for Windows clients	6	Overriding stanzas	53
Windows client communication methods	6	Client-side data deduplication	53
Backup-archive client features that are available on Windows platforms	6	Configuring the client for data deduplication	57
Windows supported file systems	7	Excluding files from data deduplication	59
NDMP support requirements (Extended Edition only)	7	Automated client failover configuration and use	60
Installation requirements for backing up and archiving Tivoli Storage Manager FastBack client data	8	Automated client failover overview	60
Client configuration wizard for Tivoli Storage Manager FastBack	9	Requirements for automated client failover	61
Installing the backup-archive client from the DVDs	9	Restrictions for automated client failover	62
Windows backup-archive client installation overview	9	Failover capabilities of Tivoli Storage Manager components	63
Windows client installation might require a reboot	10	Configuring the client for automated failover	63
Installation procedures	10	Determining the status of replicated client data	65
Installing the Windows client for the first time	11	Preventing automated client failover	66
Upgrading the Windows client	15	Forcing the client to fail over	67
Reinstalling the Windows client	19	Configuring the client to back up and archive Tivoli Storage Manager FastBack data	67
Silent installation	19	Configuring the backup-archive client to protect FastBack client data	69
Modifying, repairing, or uninstalling the Windows client	23	Configuring the backup-archive client in a cluster server environment	70
Troubleshooting problems during installation	25	Protecting data in MSCS clusters (Windows Server clients)	71
Software updates	26	Configuring cluster protection (Windows Server clients)	71
Installing the client management service to collect diagnostic information	26	Configure the web client in a cluster environment	72
Chapter 2. Configure the Tivoli Storage Manager client	27	Configure the web client to process cluster disk resources	72
Client options file overview	27	Frequently asked questions	79
Creating and modifying the client options file	29	Configuring online-image backup support	82
Create a shared directory options file	31	Configuring Open File Support	83
Creating multiple client options files	31	Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups	83
Environment variables	32	Protecting clustered-data ONTAP NetApp file server volumes	85

SnapMirror support for NetApp snapshot-assisted progressive incremental backup (snapdiff)	88
Register your workstation with a server	91
Closed registration	91
Open registration	91
Creating an include-exclude list	92
Include-exclude options	93
Exclude file spaces and directories.	93
Include-exclude statements for networked file systems.	94
Exclude files and directories from a journal-based backup	95
Control processing with exclude statements	95
System files to exclude	96
Exclude files with UNC names	97
Include and exclude files that contain wildcard characters.	97
Include and exclude groups of files with wildcard characters.	98
Examples using wildcards with include and exclude patterns	100
Determine compression and encryption processing	100
Preview include-exclude list files	101
Include and exclude option processing	102
Processing rules when using UNC names	104
Explicit use of UNC names for remote drives	104
Conversion of DOS pathnames for fixed and remote drives	104
Character-class matching examples	104

Chapter 3. Getting started 105

Backup-archive client operations and security rights	105
Backup Operators group operations	107
Considerations before you start using a Backup Operators group account	107
Permissions required to restore files that use adaptive subfile backup	108
Permissions required to back up, archive, restore or retrieve files on cluster resources	108
Tivoli Storage Manager client authentication	108
User account control	109
Enabling client access to network shares when UAC is enabled	109
Starting a Java GUI session	110
Tivoli Storage Manager password.	111
Setup wizard	111
Starting a command-line session	111
Using batch mode	112
Issuing a series of commands by using interactive mode	112
Displaying Euro characters in a command-line prompt	113
Use options on the DSMC command	113
Specifying input strings that contain blank spaces or quotation marks	114
Starting a web client session	115
User privileges	115
Start the client scheduler automatically.	116

Changing your password	116
Sorting file lists using the backup-archive client GUI	117
Displaying online help	119
Ending a session	119
Online forums	119

Chapter 4. Backing up your data . . . 121

Planning your backups (Windows)	121
Which files are backed up	122
Open file support for backup operations	123
Backing up data using the GUI	125
Specifying drives in your domain	126
Backing up data using the command line	126
Deleting backup data.	128
When to back up and when to archive files	129
Pre-backup considerations (Windows)	130
LAN-free data movement	130
LAN-free prerequisites	130
LAN-free data movement options	131
Unicode file spaces (Windows)	131
Incremental backups on memory-constrained systems	132
Incremental backups on systems with a large number of files	132
Performing a backup with limited bandwidth	133
Control processing with an include-exclude list	134
Data encryption during backup or archive operations	134
Maximum file size for operations.	135
How Tivoli Storage Manager handles long user and group names	135
Incremental, selective, or incremental-by-date backups (Windows)	136
Full and partial incremental backup	136
Journal-based backup.	138
Incremental-by-date backup	140
Comparing incremental-by-date, journal-based, and NetApp snapshot difference to full incremental and partial incremental backups	141
Snapshot differential backup with an HTTPS connection	142
Running a snapshot differential backup with an HTTPS connection	143
Selective backup	144
Backing up files from one or more file spaces for a group backup (Windows)	144
Backing up data with client-node proxy support (Windows)	145
Enabling multiple node operations from the GUI	146
Setting up encryption	147
Scheduling backups with client-node proxy support	147
Associate a local snapshot with a server file space (Windows)	148
Backing up Windows system state	149
Backing up Automated System Recovery files	150
Preparation for Automated System Recovery	151
Creating a client options file for Automated System Recovery	151

Backing up the boot drive and system drive for Automated System Recovery	152
Image backup	153
Performing prerequisite tasks before creating an image backup	154
Utilizing image backups to perform file system incremental backups	155
Method 1: Using image backups with file system incremental backups	155
Method 2: Using image backups with incremental-by-date image backups	156
Comparing methods 1 and 2	157
Performing an image backup using the GUI	157
Performing an image backup using the command line	158
Back up NAS file systems using Network Data Management Protocol	159
Backing up NAS file systems with the web client GUI using NDMP protocol	160
Back up NAS file systems using the command line	161
Methods for backing up and recovering data on NAS file servers accessed by CIFS protocol	163
Support for CDP Persistent Storage Manager	164
Backing up VMware virtual machines	164
Preparing the environment for full backups of VMware virtual machines	167
Creating full backups for VMware virtual machines	170
Preparing the environment for file-level backups of VMware virtual machines	171
Creating file-level backups for VMware virtual machines	173
Parallel backups of virtual machines	174
Back up virtual machines on a Hyper-V system	178
Hyper-V backup support limitations	180
Back up and archive Tivoli Storage Manager FastBack data	181
Backing up Net Appliance CIFS share definitions	182
Display backup processing status	182
Backup (Windows): Additional considerations	185
Open files	185
Ambiguous file space names in file specifications	186
Management classes	186
Deleted file systems	187
Removable media backup	187
Fixed drives	188
NTFS and ReFS file spaces	188
Universal Naming Convention names	188
Examples: UNC names in domain lists	188
Examples: UNC name backup	189
Microsoft Dfs file protection methods	190

Chapter 5. Restoring your data 193

Duplicate file names	193
Universal Naming Convention names restore	194
Active or inactive backup restore	194
Restoring files and directories	195
Restoring data by using the GUI	195

Examples for restoring data using the command line	195
Examples: Restoring large amounts of data	197
Standard query restore, no-query restore, and restartable restore	198
Restoring Windows system state	199
Restoring Automated System Recovery files	201
Restoring the operating system when the computer is working	201
Recovering a computer when the Windows OS is not working	202
Creating a bootable WinPE CD	202
Restoring the Windows operating system with Automated System Recovery	202
Microsoft Dfs tree and file restore	203
Restoring an image	203
Restoring an image using the GUI	204
Restoring an image using the command line	205
Restore data from a backup set	206
Restore backup sets: considerations and restrictions	208
Backup set restore	209
Restoring backup sets using the GUI	209
Backup set restores using the client command-line interface	211
Restore Net Appliance CIFS shares	211
Restoring data from a VMware backup	212
Restoring full VM backups	213
Scenarios for running full VM instant access and full VM instant restore from the backup-archive client command line	214
Full VM instant restore cleanup and repair scenarios	217
Recovering from non-standard error conditions	219
Scenario: Restoring file-level VM backups	219
Restoring full VM backups that were created with VMware Consolidated Backup	222
Restore Windows individual Active Directory objects	224
Reanimate tombstone objects or restoring from a system state backup	224
Restoring Active Directory objects using the GUI and command line	225
Restrictions and limitations when restoring Active Directory objects	226
Preserve attributes in tombstone objects	227
Modifying the client acceptor and agent services to use the web client	227
Restoring or retrieving data during a failover	228
Authorizing another user to restore or retrieve your files	229
Restoring or retrieving files from another client node	231
Restoring or retrieving your files to another workstation	232
Deleting file spaces	232
Restoring data to a point in time	233
Restore NAS file systems	235
Restoring NAS file systems using the web client	235

Restoring NAS files and directories using the web client	236
Options and commands to restore NAS file systems from the command line	237

Chapter 6. Archive and retrieve your data (Windows) 239

Archive files.	239
Snapshot backup or archive with open file support	240
Archiving data with the GUI	240
Archive data examples by using the command line.	241
Associate a local snapshot with a server file space (Windows)	242
Archiving data with client node proxy	242
Deleting archive data.	244
Retrieve archives	244
Retrieving archives with the GUI.	245
Retrieve archive copies by using the command line.	245

Chapter 7. Tivoli Storage Manager scheduler overview 247

Examples: Blank spaces in file names in schedule definitions	248
Preferential start times for certain nodes	248
Scheduler processing options	249
Evaluate schedule return codes in schedule scripts.	250
Return codes from preschedulecmd and postschedulecmd scripts.	251
Client-acceptor scheduler services versus the traditional scheduler services	252
Setting the client scheduler process to run as a background task and start automatically at startup	252
Examples: Display information about scheduled work	254
Display information about completed work	256
Examples: event logs	257
Specify scheduling options	260
Enable or disable scheduled commands	260
Change processing options used by the scheduler service.	261
Manage multiple schedule requirements on one system	261

Chapter 8. Client return codes 265

Chapter 9. Storage management policies 267

Policy domains and policy sets	267
Management classes and copy groups	268
Display information about management classes and copy groups	269
Copy group name attribute.	269
Copy type attribute	269
Copy frequency attribute	270
Versions data exists attribute	270

Versions data deleted attribute	270
Retain extra versions attribute.	270
Retain only version attribute	271
Copy serialization attribute.	271
Copy mode parameter	272
Copy destination attribute	272
Retain versions attribute.	272
Deduplicate data attribute	273
Select a management class for files	273
Assign a management class to files	273
Override the management class for archived files	274
Select a management class for directories	275
Bind management classes to files.	275
Rebind backup versions of files	276
Retention grace period	276
Event-based policy retention protection.	277
Archive files on a data retention server.	278

Chapter 10. Tivoli Storage Manager Client Service Configuration Utility . . . 279

Install the backup-archive scheduler service	279
Using the Client Service Configuration Utility (Windows)	279
Examples: Automating backups	279
Examples: Configuring the CAD to manage an existing scheduler service	281
Creating a new scheduler and associating a CAD to manage the scheduler.	282
dsmcutil command	283
Dsmcutil commands: Required options and examples	283
Dsmcutil valid options	290

Chapter 11. Processing options 295

Processing options overview	295
Communication options.	296
TCP/IP options	296
Named Pipes option	297
Shared memory options	297
Backup and archive processing options.	297
Restore and retrieve processing options.	306
Scheduling options	309
Format and language options	310
Command processing options	311
Authorization options	311
Error processing options.	312
Transaction processing options	312
Web client options.	313
Diagnostics options	313
Using options with commands	314
Entering options with a command	314
Initial command-line-only options	319
Client options that can be set by the Tivoli Storage Manager server	320
Client options reference	321
Absolute	322
Adlocation	323
Archmc.	323
Asnodename	324

	Session settings and schedules for a proxy		Virtual machine exclude options	404
	operation	326	Fbbranch	406
	Asrmode	327	Fbclientname	407
	Auditlogging	328	Fbpolicyname	409
	Auditlogname	330	Fbreposlocation	410
	Autodeploy	332	Fbserver	411
	Autofsrename	333	Fbvolumename	412
	Backmc	335	Filelist	414
	Backupsetname	336	Filename	417
	Basesnapshotname	337	Filesonly	418
	Cadlistenonport	338	Forcefailover	419
	Casesensitiveaware	339	Fromdate	420
	Changingretries	340	Fromnode	420
	Class	341	Fromtime	421
	Clientview	342	Groupname	422
	Clusterdisksonly	342	Guitreeviewafterbackup	422
	Clusternode	345	Host	423
	Collocatebyfilespec	346	Httpport	423
	Commmethod	347	Hsmreparsetag	424
	Commrestartduration	348	Ieobjtype	425
	Commrestartinterval	349	Ifnewer	426
	Compressalways	350	Imagegapsize	427
	Compression	351	Imagetofile	428
	Console	352	Inactive	428
	Createnewbase	354	Incllexcl	429
	Datacenter	356	Considerations for Unicode-enabled clients	430
	Datastore	356	Include options	430
	Dateformat	356	Compression, encryption, and adaptive	
	Dedupcachepath	359	subfile backup processing	435
	Dedupcachesize	359	Processing NAS file systems	436
	Deduplication	360	Virtual machine include options	437
	Deletefiles	361	Incrbydate	445
	Description	362	Incremental	446
	Detail	363	Incrthreshold	446
	Diffsnapshot	364	Instrogrammax	448
	Diffsnapshotname	366	Instrogramname	449
	Dirmc	367	Journalpipe	450
	Dirsonly	368	Lanfreecommmethod	450
	Disablenqr	368	Lanfreeshmport	452
	Diskbuffsize	369	Lanfreetcppport	453
	Diskcachelocation	370	Lanfreessl	453
	Domain	371	Lanfreetcpserveraddress	454
	Domain.image	374	Language	455
	Domain.nas	375	Latest	456
	Domain.vcd	376	Localbackupset	456
	Domain.vmfile	378	Managedservices	457
	Domain.vmfull	381	Maxcmdretries	459
	Enable8dot3namesupport	386	Mbobjrefreshthresh	460
	Enablearchiveretentionprotection	388	Mbpctrefreshthresh	461
	Enablededupcache	389	Memoryefficientbackup	462
	Enableinstrumentation	390	Mode	463
	Enablelanfree	392	Monitor	468
	Encryptiontype	393	Myprimaryserver	468
	Encryptkey	394	Myreplicationserver	469
	Errorlogmax	396	Namedpipename	471
	Errorlogname	397	Nasnodename	471
	Errorlogretention	398	Nodename	472
	Exclude options	400	Nojournal	474
	Controlling compression processing	403	Noprompt	474
	Adaptive subfile backup processing	403	Nrtabpath	475
	Processing NAS file systems	404	Numberformat	476

Optfile	478	Subfilecachesize	551
Password	478	Systemstatebackupmethod	552
Passwordaccess.	480	Tapeprompt	553
Pick	481	Tcpadminport	554
Pitdate	482	Tcpbuffsize.	555
Pittime	483	Tpcadaddress	556
Postschedulecmd/Postnschedulecmd	484	Tcpclientaddress	557
Postsnapshotcmd	486	Tcpclientport	557
Preschedulecmd/Prenschedulecmd	487	Tcpnodelay	558
Preserveaccessdate.	489	Tcpport	559
Preservepath	490	Tcpserveraddress	560
Presnapshotcmd.	492	Tcpwindow size	560
Queryschedperiod	494	Timeformat	562
Querysummary	495	Toc	563
Quiet	496	Todate.	564
Quotesareliteral	497	Totime.	565
Replace	499	Txnbytelimit	566
Replserverguid.	500	Type	567
Replservername.	501	Usedirectory	568
Replsslport.	503	Useexistingbase	569
Repltcpport.	504	Usereplicationfailover.	570
Repltcpserveraddress	505	V2archive	570
Resetarchiveattribute	507	Vappmc.	571
Resourceutilization.	508	Vcdhost	572
Regulating backup and archive sessions	509	Vcdpw	573
Regulating restore sessions	510	Vcduser	573
Multiple client session considerations	510	Verbose	574
Retryperiod.	511	Verifyimage.	575
Revokeremoteaccess	512	Virtualfsname	575
Runasservice	512	Virtualnodename	576
Schedcmddisabled	513	Vmautostartvm	577
Schedcmdexception	514	Vmbackdir	578
Schedlogmax.	515	Vmbacknodelete.	579
Schedlogname	517	Vmbackupmailboxhistory.	580
Schedlogretention	518	Vmbackuptype	581
Schedmode	519	Vmbackvcbtransport	582
Schedrestretrdisabled	521	Vmchost	583
Scrolllines.	521	Vmcpw	584
Scrollprompt	522	Vmctlmc	585
Sessioninitiation	523	Vmcuser	586
Shmport	525	Vmdatastorethreshold	587
Showmembers.	526	Vmdefaultdvportgroup	588
Skipmissingsyswfiles	526	Vmdefaultdvswitch	589
Skipntpermissions	528	Vmdefaultnetwork	590
Skipntsecuritycrc	528	Vmdiskprovision	591
Snapdiff.	529	Vmenabletemplatebackups	591
Snapdiffhttps	534	Vmexpireprotect	593
Snapshotproviderfs	535	Vmiscsiadapter.	594
Snapshotproviderimage	536	Vmiscsiserveraddress	595
Snapshotroot	537	Vmlimitperdatastore.	596
Srvoptsetencryptiondisabled.	539	Vmlimitperhost.	597
Srvprepostscheddisabled	540	Vmlist.	598
Srvprepostsnapdisabled.	541	Vmmaxparallel	598
Ssl	542	Vmmaxvirtualdisks	599
Sslfipsmode.	543	Vmmc	601
Sslrequired.	544	Vmmountage	601
Ssldisablelegacytls.	546	Vmnoprmdisks	602
Stagingdirectory	547	Vmnovrmdisks	603
Subdir.	548	Vmpreferdagpassive	604
Subfilebackup	549	Vmprocessvmwithindependent	604
Subfilecachepath	550	Vmprocessvmwithprdm.	606

Vmrestoretype	607
Vmskipctlcompression	609
Vmskipmaxvirtualdisks	610
Vmskipmaxvmdks.	611
Vmstoragetype	611
Vmtagdatamover.	612
Vmtempdatastore	615
Vmverifyifaction	615
Vmverifyiflatest	617
Vmvstortransport	618
Vssaltstagingdir	619
Vssusesystemprovider	620
Vmtimeout	621
Webports	622
Wildcardcardsareliteral.	623

Chapter 12. Using commands 625

Start and end a client command session	628
Process commands in batch mode	628
Process commands in interactive mode. . . .	629
Enter client command names, options, and parameters	629
Command name	630
Options	630
Options in interactive mode	630
Parameters	630
File specification syntax	631
Wildcard characters	632
Client commands reference.	633
Archive	633
Open file support	635
Archive FastBack	636
Backup FastBack	639
Backup Group	642
Backup Image	644
Offline and online image backup	646
Utilizing image backup to perform file system incremental backup	647
Backup NAS	648
Backup Systemstate	650
Backup VAPP	652
Backup VM	658
Cancel Process	670
Cancel Restore	670
Delete Access	671
Delete Archive	672
Delete Backup	674
Delete Filespace	677
Delete Group	678
Expire	680
Help	681
Incremental	683
Open file support	686
Journal-based backup.	687
Backing up NTFS or ReFS volume mount points	688
Backing up data on NTFS or ReFS mounted volumes	688
Back up Microsoft Dfs root	689
Incremental-by-Date	689
Adaptive subfile backups	690

Associate a local snapshot with a server file space	690
Loop	690
Macro	692
Monitor Process	692
Preview Archive	693
Preview Backup	694
Query Access	695
Query Adobjects	696
Query Archive	697
Query Backup	700
Query NAS file system images	703
Query Backupset	703
Query Backupset without the backupsetname parameter	705
Query Filespace	706
Query NAS file spaces	708
Query Group	709
Query Image	710
Query Inclexcl	712
Query Mgmtclass	713
Query Node	714
Query Options	715
Query Restore	716
Query Schedule	717
Query Session	717
Query Systeminfo	718
Query Systemstate	720
Query VAPP	721
Query VM	725
Restart Restore	729
Restore	730
Restoring NTFS or ReFS volume mount points	735
Restoring data on NTFS mounted volumes	735
Restore Microsoft Dfs junctions	736
Restore active files.	736
Universal Naming Convention restores.	736
Restore from file spaces that are not Unicode-enabled	737
Restore named streams	737
Restore sparse files	737
Restore Adobjects	738
Restore Backupset	739
Restore backup sets: considerations and restrictions	742
Restore backup sets in a SAN environment	743
Restore Backupset without the backupsetname parameter	744
Restore Group	746
Restore Image	748
Restore NAS	751
Restore Systemstate	753
Restore VAPP	755
Restore VM	758
Retrieve	766
Retrieve archives from file spaces that are not Unicode-enabled	769
Retrieve named streams	769
Retrieve sparse files	770
Schedule	770
Selective	771

Open file support	774
Adaptive subfile backups	774
Associate a local snapshot with a server file space	775
Set Access	775
Set Event.	778
Set Netappsvm	780
Set Password	781
Set Vmtags	787

**Appendix. Accessibility features for
the Tivoli Storage Manager product
family. 791**

Notices 793

Glossary 797

Index 799

Tables

1. Windows client communication methods	6
2. Supported features on Windows platforms	7
3. Stoppable services	16
4. File path and name limits	28
5. Client acceptor-managed services versus traditional scheduler services.	36
6. Data deduplication settings: Client and server	56
7. Options for excluding file spaces and directories	94
8. Options for controlling processing using include and exclude statements	95
9. Wildcard and other special characters	98
10. Specifying a drive specification using wildcards	99
11. Using wildcard characters with include and exclude patterns	100
12. Options for controlling compression and encryption processing.	100
13. UNC name patterns and DOS patterns	104
14. Required user security rights for Tivoli Storage Manager backup and restore services	105
15. Working with your files using the backup-archive client GUI	118
16. Planning your backups	121
17. Command line backup examples	127
18. Maximum file size	135
19. Comparing incremental image backup methods	157
20. NAS options and commands	161
21. Backup and restore capabilities for VMware virtual machines on Windows platforms	165
22. Comparison of full backups and file-level backups for VMware virtual machines on Windows systems	167
23. Client command line informational messages	182
24. UNC examples	189
25. Command-line restore examples	196
26. Backup set GUI restore restrictions	207
27. Backup set command-line restore restrictions	207
28. Components for the restore command when you restore files to the same computer	221
29. Components for the restore command when you restore files to a different computer.	221
30. NAS options and commands	237
31. Command-line archive examples	241
32. Command line examples of retrieving archives	246
33. Sample classic query schedule output	255
34. Sample enhanced query schedule output	256
35. Client return codes and their meanings	265
36. Default attribute values in the standard management class	269
37. TCP/IP options	296
38. Named Pipes communication option	297
39. Shared memory communication options	297
40. Backup and archive processing options	297
41. Restore and retrieve processing options	306
42. Scheduling options.	309
43. Format and language options	310
44. Command processing options	311
45. Authorization options.	311
46. Error processing options	312
47. Transaction processing options.	312
48. Web client options	313
49. Diagnostics options	313
50. Client command options.	315
51. Options that are valid on the initial command line only	320
52. Options that can be set by the Tivoli Storage Manager server	321
53. Setting the value of the asnodename option to distribute backups.. . . .	324
54. Clusternode and clusterdisksonly combinations.	344
55. Interaction of domain definitions from several sources.	374
56. System services components and corresponding keywords.	402
57. Other optional parameters	434
58. Incremental command: Related options	531
59. Effects of server and client SSL settings on success or failure of login attempts	545
60. Commands	625
61. Wildcard characters	633
62. Archive command: Related options	634
63. Archive FastBack command: Related options	637
64. Backup FastBack command: Related options	640
65. Backup Group command: Related options	643
66. Backup Image command: Related options	645
67. Backup NAS command: Related options	650
68. Backup VAPP command: Related options	654
69. Delete Archive command: Related options	673
70. Delete Backup command: Related options	676
71. Delete Filespace command: Related options	678
72. Delete Group command: Related options	679
73. Expire command: Related options.	681
74. Incremental command: Related options	684
75. Query Adobjects command: Related options	696
76. Query Archive command: Related options	698
77. Query Backup command: Related options	701
78. Query Backupset command: Related options	704
79. Query Backupset command: Related options	706
80. Query Filespace command: Related options	707
81. Query Group command: Related options	709
82. Query Image command: Related options	711
83. Query Mgmtclass command: Related options	714
84. Query Node command: Related options	715
85. Query Options command: Related options	716
86. Query Systeminfo command: Related options	720
87. Query Systemstate command: Related options	720
88. Query VAPP command: Related options	723

89. Query VM command: Related options for VMware virtual machine queries.	726	97. Restore Systemstate command: Related options	754
90. Query VM command: Related options for Hyper-V virtual machine queries.	728	98. Restore VAPP command: Related options	757
91. Restore command: Related options	733	99. Restore VM command: Related options when restoring VMware virtual machines	762
92. Restore Adobjects command: Related options	738	100. Restore VM command: Related options when restoring Hyper-V virtual machines	765
93. Restore Backupset command: Related options	741	101. Retrieve command: Related options	767
94. Restore Group command: Related options	747	102. Schedule command: Related options	771
95. Restore Image command: Related options	749	103. Selective command: Related options	773
96. Restore NAS command: Related options	752		

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.
- 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.

Related concepts:

“Planning your backups (Windows)” on page 121

“What's new for Version 7.1.6” 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.

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

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.
DSMI_DIR	A series of uppercase letters indicates return codes and other values.
dsmQuerySessInfo	Boldface type indicates a command that you type on a command line, the name of a function call, the name of a structure, a field within a structure, or a parameter.
<i>timeformat</i>	Boldface italic type indicates a Tivoli Storage Manager 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*—┐
└─'(1)—————►◄

Notes:

1 Specify *repeat* up to 5 times.

Required choices

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

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


►►—cmd_name—┐
└─A
└─B
└─C—————►◄

Optional choices

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

►►—cmd_name—┐
└─A—————►◄

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



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 7.1.6

IBM Tivoli Storage Manager Version 7.1.6 introduces new features and updates.

For a list of new features and updates in this release and previous Version 7 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.

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 7.1.6” on page xvii

“Planning your backups (Windows)” on page 121

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 7.1.6 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.

Considerations for migrating between processor architectures x32 and x64

When migrating between the Tivoli Storage Manager processor architectures (x32 and x64), there are some items that you need to consider for a successful migration from one architecture to another.

Consider the following items when migrating from one processor architecture to another:

- The restore of DFS links backed up from 32-bit Windows environments to 64-bit Windows environments, is not supported.
- The restore of DFS links backed up from 64-bit Windows environments to 32-bit Windows environments, is not supported.
- System state data is incompatible between different processor architectures. Therefore, when migrating client data between different processor architectures, avoid backing up the system state data from one architecture into the same file spaces containing the system state data from another architecture. To avoid commingling the data, you can do either of the following before the migration:

- Ask your Tivoli Storage Manager administrator to rename the existing system state file spaces. Once the migration is complete and you have new backups of the system state data, your administrator can delete the renamed file spaces. The file spaces to rename include one or more of the following:
 - SYSTEM STATE
 - SYSTEM SERVICES
 - ASR

Ask your Tivoli Storage Manager administrator to delete the existing file spaces for your node.

- Renamed file spaces remain on the server and managed as stabilized file spaces. These file spaces contain all the original data, which can still be restored until the file spaces are deleted.
- When an existing file space is renamed during Unicode conversion, any access rules defined for the file space remain applicable to the original file space. New access rules must be defined to apply to the new Unicode file space, if necessary.
- When backing up files to a file space that is not Unicode-enabled, the Unicode-enabled client skips the files and directories with names containing characters from a code page other than the current locale.
- If files and directories with names containing characters from a code page other than the current locale were previously backed up with a client that was not Unicode-enabled, Tivoli Storage Manager expires them in the file space that is not Unicode-enabled. However, the Unicode-enabled client can back up or archive these files to a Unicode-enabled file space.
- When migrating from a client that is not Unicode-enabled to a Unicode client, file names with double-byte characters mixed with single-byte characters might be restored in mixed single-byte character set (SBCS) format. This is only a display issue.
- When a client performs a selective backup of an object and the original file space is renamed, the new Unicode-enabled file space contains only that object. The client backs up all other directories and file spaces during the next full incremental backup.
- The Unicode-enabled client has a file and directory name size limitation of 504 bytes. Depending on the Unicode characters (2-byte, 3-byte, or 4-byte) used, this translates to a maximum of 126 to 252 characters.

Important:

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

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:

- The size of the buffer to record change notifications for a particular journal file system (**DirNotifyBufferSize**) has changed. The default value is 16 KB.
- 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 workstations that already have the backup-archive client installed.

When you schedule automatic Backup-Archive client deployments, the updated client packages (which include the client components and the API library) are installed on the workstations that receive them. A dependency check is performed by the client installation program to ensure that the API library does not conflict with the client package that is already installed.

Tivoli Storage Manager for ERP applications do not use the same installation technology that is used by the client installation program. Because of that, the client installation dependency check is not able to detect whether the API library that is being used by the Tivoli Storage Manager for ERP applications is compatible with the API library that will be installed by automatic client deployments. If a client package is automatically deployed to and installed on a workstation, the API library that is installed might not be compatible with the API library that was installed by the Tivoli Storage Manager for ERP application. The newly deployed API library can cause the Tivoli Storage Manager for ERP applications to fail.

Do not schedule automatic client deployments to workstations that have a Tivoli Storage Manager for ERP application installed on them.

The Tivoli Storage Manager server can be configured to automatically upgrade backup-archive clients on Windows computers. The existing clients must be at version 5.4, or later.

You can use the `autodeploy` option to conditionally enable automatic client deployment, if the deployment does not require a restart of the client workstation.

By default, the `autodeploy` option is enabled, and the client workstation is restarted if required.

The `autodeploy` option can be set in the following places:

- In the Administration Center. However, starting with Tivoli Storage Manager Version 7.1, the Administration Center component is no longer included in Tivoli Storage Manager distributions. If you have an Administration Center from a previous server release, you can continue to use it to set the `autodeploy` option on schedules that deploy client updates. If you do not already have an Administration Center installed, you can use Tivoli Storage Manager server commands to automatically deploy client upgrades. For information about how to manually set up automatic deployment of backup-archive clients, see <http://www.ibm.com/support/docview.wss?uid=swg21673299>.
- On a schedule definition. Schedules are defined on the server. Schedule definitions that deploy client software updates have an `action=deploy` statement. On those schedules, you can include the `autodeploy` option as part of the command that you include on the `-postnschedulecmd` statement.
- On the client node, in an options file that is associated with the client scheduler or client acceptor daemon. The deployment manager detects options files that are associated with the scheduler or client acceptor daemon. If there are multiple scheduler or client acceptor daemon processes running on the same computer at the same time, and the processes use different options files, the deployment manager uses the `autodeploy` value that is set in one of the options files.

- On the client in the client options file (dsm.opt). The autodeploy option that is set in the client options file overrides any other autodeploy setting.

To use automatic deployment but not restart the system, specify the autodeploy noreboot option.

Important: For automatic client deployment to be operational, the following conditions apply:

- The client computer must have a minimum of 2 GB free disk space.
- The passwordaccess option must be set to **generate**.

To store the password in the local Windows registry, a user must log on from the local workstation once, to provide the password. If the client automatic deployment process cannot find the node password in the registry, the following events occur:

- The deployment process does not start.
- A warning message is logged in the setup.log file.

The following messages are examples of the messages that can be logged during a deployment.

```
Sun 10/11/2009 10:16:04.35 The deployclient.bat script is started.
Sun 10/11/2009 10:16:04.37 deployclient.bat is removing the deployment
manager and temporary files possibly left from the previous installation.
Sun 10/11/2009 10:17:10.53 WARNING: Still waiting for the result of query
system information.
Ensure that "PASSWORDACCESS GENERATE" is set in the client options file
and that a connection to the server can be made without being prompted
for a password.
```

If you see these errors and PASSWORDACCESS is set to **GENERATE**, the errors are probably caused by a network problem (the client cannot be reached) or there might be a problem with the digital rights management subkey. The digital rights management key is in the registry in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup\DRM.

- The **dsmc query systeminfo** command is left running.
- The deployment process cannot start, so no messages are sent to the server.
- The client is deployed from the server as a scheduled task; the client scheduler must be running.

The scheduler must be started as a Windows service and not from the command line. To minimize the chance of a restart, the deployment manager shuts down the scheduler service before the new client is installed, and restarts the scheduler service after the installation. If the scheduler is not run as a Windows service, a restart is required when the client is deployed.

- The automatic client deployment feature stops and restarts scheduler and CAD processes on the client, but it does not stop or restart any client operations, such as a backup or a restore. It is possible that in-process client operations could be affected during an automatic deployment. To prevent an automatic deployment from interfering with client operations, schedule automatic client deployments to occur when it is not likely that the client is performing a backup or restore, or archive or retrieve operation.
- The client is deployed as a postschedule operating system command; scheduled operating system commands must be enabled on the client.

You must configure your server to perform backup-archive client automatic deployments. See the Tivoli Storage Manager server documentation for instructions.

- The command-line Windows registry utility (reg.exe) is required. This tool is generally installed as part of the operating system installation on supported Windows platforms.
- The Windows cluster services environment is not supported.
- To turn off automatic client deployment, add **autodeploy no** to the client options file.

The autodeploy option is only available for Tivoli Storage Manager V6.2 and later clients.

When you perform automatic client deployments from the Tivoli Storage Manager Server, the scheduler installs the updated client by running a command that is specified on a postschedulecmd option that is associated with the schedule. By default, the log and trace files for a deployment operation are written to the client's disk in C:\Program Files\Tivoli\TSM\IBM_ANR_WIN\Vxxxx\log; where xxxx represents the version for the newly deployed client.

If you did not use the default installation directory when you installed the client, when an automatic client deployment is performed, the log and trace data for the deployment operation are still copied to the IBM_ANR_WIN\Vxxxx\log folder, and this folder is created one directory level up (../) from where the client executable files (dsm.exe, dsmd.exe, dsmdad.exe, and so on) were installed. For example, if you originally installed the client into E:\IBM\Tivoli\baclient, the log and trace files resulting from the automatic deployment are created in E:\IBM\Tivoli\IBM_ANR_WIN\Vxxxx\log folder.

A lock file is also created on each client to prevent a client from being updated by different deployment managers at the same time. Before a deployment begins, the deployment manager looks on the client for the lock file. If the file exists and is less than 24 hours old, the deployment is canceled. If the lock file is more than 24 hours old or does not exist, the deployment is started. The lock file is deleted when the deployment completes. If the deployment manager fails to delete the lock file, you can manually delete it from the client's disk. The lock file location is *install_dir\..\IBM_ANR_WIN\mutext.txt*.

Related reference:

"Autodeploy" on page 332

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.

- "Windows client environment requirements"
- "NDMP support requirements (Extended Edition only)" on page 7

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

Windows client environment requirements

This section contains client environment information, backup-archive client components, and hardware and software requirements for the supported Windows platforms.

Windows client installable components

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

The installable components for the Windows backup-archive 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 API (32-bit and 64-bit)
- VMware vStorage API runtime files. These files are required if you are using a backup-archive client node to act as a data mover node that protects VMware virtual machines. Windows data mover nodes must have the 64-bit Windows client installed on them.

System requirements for Windows clients

The backup-archive client on Windows requires a minimum amount of disk space for installation and a supported operating system.

For software and hardware requirements for all supported versions of Windows clients, including the most recent fix packs, see technote 1197133.

Windows client communication methods

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

You can use the following communication methods with the Windows backup-archive client:

Table 1. Windows 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 all supported Windows)	IBM AIX®, HP-UX IA®, Linux, Oracle Solaris, Windows
Named Pipes	Named Pipes (Standard with all supported Windows platforms)	Windows
Shared Memory	TCP/IP (Standard with all supported Windows platforms)	Windows

Backup-archive client features that are available on Windows platforms

This topic lists which features are supported or not supported on the various Windows platforms.

Table 2 on page 7 shows the supported and unsupported features on the various Windows platforms.

Table 2. Supported features on Windows platforms

Features	32-bit: Windows 7	64-bit: Windows 7	32-bit: Windows 8 Windows Server 2008 Windows Server Core 2008	64-bit: Windows 8 Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server Core 2008 Windows Server Core 2012
Journal-based backup	yes	yes	yes	yes
Online image backup	yes	yes	yes	yes
Offline image backup	yes	yes	yes	yes
System state support with Volume Shadowcopy Services (VSS)	yes	yes	yes	yes
LAN-free operations	no	no	yes	yes
Automated System Recovery (ASR)	yes	BIOS: yes UEFI: yes	yes	BIOS: yes UEFI: yes
Open File Support (OFS)	yes	yes	yes	yes

Windows supported file systems

The Tivoli Storage Manager Windows backup-archive client is supported on specific file systems.

The Windows backup-archive client supports the following types of file systems:

- GPFS™

Restriction: The Windows client can back up or restore files in a Windows-only GPFS cluster. If the GPFS cluster contains mixed nodes (some AIX, some Linux, and some Windows), you must use the AIX or Linux clients to protect data in the cluster.

- File Allocation Table (FAT and FAT32)
- Microsoft New Technology File System (NTFS)
- Microsoft Resilient File System (ReFS). ReFS was introduced on Windows Server 2012 systems.

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

Related concepts:

“Configuring the client to back up and archive Tivoli Storage Manager FastBack data” on page 67

Client configuration wizard for Tivoli Storage Manager FastBack

The backup-archive client provides a wizard to configure the backup-archive client for Tivoli Storage Manager FastBack.

The wizard is available in a remote application (the web client) and in a local application (the Java GUI). The wizard helps you set the options to send FastBack client data to the Tivoli Storage Manager server on a scheduled basis.

You can use the client configuration wizard for Tivoli Storage Manager FastBack on 32-bit Windows systems that are supported by both Tivoli Storage Manager and FastBack. The configuration wizard must be installed on the same workstation where the FastBack server resides, or where a FastBack Disaster Recovery Hub is deployed. You can use the client configuration wizard for Tivoli Storage Manager FastBack with Tivoli Storage Manager FastBack Version 5.5.1 or later versions.

Related concepts:

“Configuring the backup-archive client to protect FastBack client data” on page 69

Installing the backup-archive client from the DVDs

You can install the backup-archive client and its components from Tivoli Storage Manager DVDs.

About this task

To install the backup-archive client from a DVD, you change to the directory that contains the installation image for your client operating system and you install the backup-archive client and its components.

The installation image for Windows backup-archive clients is in the `setup` directory.

Procedure

1. Install the backup-archive clients by starting the installation program from the appropriate installation image directory.
2. Optional: Copy the installation images from the DVD to a server workstation so that client workstations can mount the server volume and start the installation from the server copies.

Copy all of the files that are in the DVD installation directory for the client operating systems and all of the subdirectories and files that are subordinate to it.

For example, Windows users can use the `xcopy` command to copy the `setup` directory and its subdirectories from the Windows DVD to a remote disk. In this example, the H drive represents the DVD drive letter and the X drive represents the remote disk: `xcopy h:\setup* x:\ /s`. Substitute the drive letters that are appropriate for your environment.

Related tasks:

“Windows backup-archive client installation overview”

Windows backup-archive client installation overview

You can install the Tivoli Storage Manager Windows backup-archive client from the product DVD, or other installation media.

Before you begin

Before you begin a Windows client installation, ensure that the system that you want to install the client on meets the client requirements. Then, determine the type of installation that you need to perform, and follow the steps in the appropriate procedure.

For the hardware and software requirements for the Windows client, see technote 1197133.

Related concepts:

“Automatic backup-archive client deployment” on page 3

Related tasks:

“Creating and modifying the client options file” on page 29

“Starting a web client session” on page 115

Windows client installation might require a reboot

As part of the Windows client installation process, one or more Microsoft C++ redistributable packages are installed, if they are not already installed on the Windows workstation. These packages can also be automatically updated by the Windows Update service. If the packages are updated, the update can cause the system to reboot when you start the Windows client installation program.

The reboot that is triggered if the C++ redistributable packages are updated can occur, even under any of the following conditions:

- An automatic client deployment pushes a client upgrade to a node, and the client or the scheduler sets the AUTODEPLOY=NOREBOOT option.
- A manual installation or upgrade of the client is started.
- A client silent installation is started, even if the options to suppress reboot prompts, and the client reboot itself, are set.

Additionally, because the Microsoft Visual Studio C++ redistributable package is a shared Windows component, other applications that have dependencies on the package might be stopped or restarted by Windows as part of the installation or upgrade of the C++ redistributable package. Schedule client installations and upgrades during a maintenance window when other applications will not be adversely affected if they are stopped or restarted when the C++ redistributable package is installed. Monitor other applications after the client is installed to see whether there are any applications that were stopped and not restarted.

Installation procedures

The procedure that you follow to install the Tivoli Storage Manager Windows backup-archive client depends on the type of installation that you want to perform.

Procedures are provided for each of the following installation types:

Installation type	Installation description
Installing the Windows client for the first time	Describes how to install the Windows backup-archive client for the first time. This procedure presumes that the Windows computer that you are installing the client on has never had a previous version of the client installed on it before.

Installation type	Installation description
Upgrading the Windows client	Describes how to upgrade an earlier version of the Windows backup-archive client to this latest version.
Reinstalling the Windows client	Describes how to reinstall the Windows backup-archive client, if you uninstalled it.
Silent installation	Describes how to install the Windows backup-archive client silently, without user interaction during the installation procedure.
Repairing, modifying, or uninstalling the Windows client	Describes how to add or remove features from an installed backup-archive client (modify), replace damaged files or missing registry keys (repair), or uninstall the Windows backup-archive client.

Installing the Windows client for the first time

Complete this procedure to install the Windows backup-archive client for the first time.

Before you begin

If you have an earlier version of the Windows backup-archive client that is already installed on a node and you want to upgrade it to Version 7.1.6, see “Upgrading the Windows client” on page 15.

Important: You must know the host name or IP address of the Tivoli Storage Manager server, the port number that the server listens on for client communications, and the communications method to use when the client communicates with the server. Obtain this information from your Tivoli Storage Manager server administrator before you start this procedure.

Procedure

1. Install the product by using the installation DVD or by using the compressed installation file that you download from IBM.

Option	Description
Install from a DVD	<ol style="list-style-type: none"> 1. Insert the DVD into the DVD drive. 2. If the installation program does not automatically start, browse the DVD to locate the spinstall.exe file. Double-click the spinstall.exe file.

Option	Description
Install by using a compressed installation file	<ol style="list-style-type: none"> 1. The Tivoli Storage Manager V7.1 client release notes contain a link to a document that describes how to download the Windows backup-archive client distribution package from the IBM Passport Advantage® website. Follow the directions in that document to obtain the required files. 2. Copy the compressed installation file that you downloaded from Passport Advantage to a local disk or to a network-accessible share. 3. Double-click the compressed installation file. 4. By default, the uncompressed files are stored in the current disk drive, in the <i>download directory</i>\TSMClient directory. If the installation program detects files from another client installation attempt in this directory, you are prompted to determine whether you want to overwrite the old files. If you receive this prompt, enter A to overwrite the existing files; this selection ensures that only the files from the current installation are used. 5. In the <i>download directory</i>\TSMClient directory, double-click the spinstall.exe file to start the client installation program.

2. Select a language to use for this installation and click **OK**.
3. If the installation wizard indicates that one or more Microsoft C++ redistributable files must be installed, click **Install**. These files are needed to run the Windows client.
4. On the Tivoli Storage Manager client welcome screen, click **Next** to begin installing the client software.
5. Accept the default installation directory by clicking **Next**, or specify a different installation directory. The default installation directory is C:\Program Files\Tivoli\TSM.
6. Select the installation type: **Typical** or **Custom**.

Option	Description
Typical	<p>A typical installation installs the following components:</p> <ul style="list-style-type: none"> • The backup-archive client GUI files (needed to use the Java GUI) • The backup-archive client web files (needed to use the web client) • The client API files (as needed by your client and operating system)

Option	Description
Custom	<p>A custom installation installs the same files as a typical installation. However, you can choose to install the following optional components:</p> <ul style="list-style-type: none"> • The API SDK files (only needed if you are developing applications that work with the backup-archive client) • The Administrative Client command-line files (required to remotely run administrator functions on the Tivoli Storage Manager server) • VMware vStorage API runtime files (needed if you are using this client to back up VMware virtual machines)

7. Click **Next**, then click **Install**.
8. When the installer completes the installation, click **Finish**.
9. Verify the installation. Click **Start > All Programs > Tivoli Storage Manager**. The client components that you installed are shown in the list of Tivoli Storage Manager startable programs. The administrative command-line client, backup-archive command-line client, and the backup-archive GUI are the only components that are displayed in this list. The administrative command-line client is only shown if you perform a custom installation and you include the administrative command-line client. If you installed other components, such as the API, SDK, or VMware vStorage API runtime files, they are not shown in this list.
10. Click **Backup-Archive GUI** to start the client GUI. The TSM Client Options File Configuration Wizard starts. Click **Next** to start the wizard.
11. On the Options File Task screen, select **Create a new options file** and click **Next**.
12. On the TSM Node Name screen, specify a node name. A node name uniquely identifies your node to the Tivoli Storage Manager server. The default node name is the short host name of the Windows computer that you are installing the client on. Accept the default node name or specify a new node name. Click **Next**.
13. On the TSM Client/Server Communications screen, specify the communications method to use when the client communicates with the server and click **Next**. This information must be provided to you by your Tivoli Storage Manager server administrator. If you are not sure what to select, accept the default setting (TCP/IP). If the default setting does not work when the client attempts to connect to the server, contact the server administrator to determine which communications method to specify.
14. On the TCP/IP Options screen, specify the server address and port information that your Tivoli Storage Manager administrator provided to you. In the **Server Address** field, specify the IP address or host name of the Tivoli Storage Manager server. In the **Port Number** field, specify the port number that the server listens on for client communications. Click **Next**.
15. The Recommended Include/Exclude List screen contains a list of system files and directories that are typically included, or excluded from client operations. The excluded files are typically not required to restore your system. You can select or clear all default selections. Alternatively, you can use the Shift and

Ctrl keys to selectively include objects. To facilitate the installation process, click **Select All**; you can add or remove files from this list later, if you want to. Click **Next**.

16. The Common File Exclusion Selection screen provides a default list of file extensions that you can exclude from client operations. The file extensions that are provided in this list are typically large files, like graphics or multimedia extension. These files consume server disk space but they might not be required to restore critical data. Click **Select All** to exclude all of the default file extensions. Alternatively, you can use the Shift and Ctrl keys to selectively choose which extensions to exclude from client operations. Click **Clear All** to clear any extensions that you selected. You can modify these extensions later if you want to. Click **Next**.
17. The Domains for Backup screen specifies the default file systems and objects to include in client operations for incremental and image backups.
 - a. To configure the default file systems for incremental backups, in the **Backup Type** field, select **Incremental**. By default, **Back up all local file systems** is selected. If you do not want to back up all local file systems as the default action during incremental backups, clear this option and individually select the file systems to include. You can override the default selection when you initiate an incremental backup operation.
 - b. To configure the default file systems for image backups, in the **Backup Type** field, select **Image**. By default, **Back up all local file systems** is selected. If you do not want to back up all local file systems as the default action during image backups, clear this option and individually select the file systems to include. You can override the default selection when you initiate an image backup operation.
 - c. Click **Next**.
18. On the Confirm and apply your configuration screen, click **Apply**. The Login into a TSM server screen displays and prompts you to enter a user ID and password to log on to the Tivoli Storage Manager server. The user ID defaults to the node name that you specified in step 12 on page 13.
19. You can accept the default user ID or specify a different user ID. Specify the password that you will use when you log on to the server. Click **Login**. What happens next depends on whether the Tivoli Storage Manager server is configured for open or closed registration.

Option	Description
Server is configured for open registration	<p>The Register New Node screen prompts you for contact information and it prompts you again for the password.</p> <p>Adding text to the Contact Information field is optional, but suggested; specify your name.</p> <p>Re-enter your password, twice, in the two Password fields. If the password that you enter and confirm in these Password fields does not match what you previously specified in the Login into a TSM server screen, the password that you specify and confirm here becomes the password that is required to log on to the server.</p> <p>Click Register to register this node on the server.</p> <p>Click Finish. The graphical user interface opens and is ready for use. You can also start any of the other installed client components from the Start menu.</p>
Server uses closed registration	<p>Click Finish. Provide the information that you specified in the client configuration wizard to your Tivoli Storage Manager server administrator. Provide the administrator with the following information:</p> <ul style="list-style-type: none"> • The node name that you specified. • The user ID and password that you entered. • Your contact information, such as your name, email address, and phone number, so the administrator can contact you after your node and user information is registered on the server. <p>After the administrator registers your node, you can start any of the installed client components from the Start menu.</p>

Related concepts:

“Troubleshooting problems during installation” on page 25

Upgrading the Windows client

You can upgrade an earlier version of the Tivoli Storage Manager Windows backup-archive client to Version 7.1.6. Your previous configuration settings are preserved, where it is possible to do so. However, enhancements that are in the latest version of the client can deprecate or prohibit the use of options that were available in earlier versions of the client.

Before you begin

Wait for any in-progress backup-archive client tasks (backup, restore, archive, retrieve) to complete before you upgrade a client node.

About this task

To upgrade to the Version 7.1.6 Windows client, install the Version 7.1.6 Windows client; you do not need to uninstall previously installed client software first. The Version 7.1.6 client installation program preserves your current client options and settings (in `dsm.opt`), and it does not overwrite or delete the `dsmerror.log`, `dsm sched.log`, and `dsmwebcl.log` files, if you install the new client into the same directory that was used by the previous installation.

The Logical Volume Snapshot Agent (LVSA) component was deprecated in Tivoli Storage Manager Version 6.4. If you previously had LVSA configured as your snapshot provider, install the Version 7.1.6 client, and then configure it to use the Microsoft Volume Shadow Copy Service (VSS) as the snapshot provider in the new installation. If LVSA was installed, your client reboots after the upgrade installation completes, to allow for the removal of LVSA entries from the registry.

The installation program stops any client services that are running before it upgrades the client software. If you prefer, you can manually stop the services by using the control panel or command line. Table 3 shows the stoppable services, and the names to look for in the **Control Panel > Administrative Tools > Services** list, so you can stop them with the Control Panel. The table also provides the commands to stop them from a command prompt or a script.

Note: The service names that are shown in the table are the default names that are set by the installation program. You can change some of these service names when you configure the services by using one of the configuration wizards on the **Utilities > Setup Wizard** menus. If you change the service name, record the name that you specify and use that name to stop the services.

Table 3. Stoppable services

Control panel display name	Command-line procedure
TSM Journal Service	<code>net stop "tsm journal service"</code>
TSM Client Acceptor	<code>net stop "tsm client acceptor"</code>
TSM Client Scheduler	<code>net stop "tsm client scheduler"</code>
Remote Client Agent	<code>net stop "tsm remote client agent"</code>

Perform the following steps to upgrade an earlier version of the Windows backup-archive client to Version 7.1.6:

Procedure

1. Install the product by using the installation DVD or the compressed installation file that you downloaded from Passport Advantage.

Option	Description
DVD	<ol style="list-style-type: none">1. Insert the DVD into the DVD drive.2. If the installation program does not automatically start, browse the DVD to locate the <code>spinstall.exe</code> file. Double-click the <code>spinstall.exe</code> file.

Option	Description
Compressed installation file	<ol style="list-style-type: none"> 1. The Tivoli Storage Manager V7.1 client release notes contain a link to a document that describes how to download the Windows backup-archive client distribution package from the IBM Passport Advantage website. Follow the directions in that document to obtain the required files. 2. Copy the compressed installation file that you downloaded from Passport Advantage to a local disk or to a network-accessible share. 3. Double-click the compressed installation file. 4. By default, the uncompressed files are stored in the current disk drive, in the <i>download directory</i>\TSMClient directory. If the installation program detects files from another client installation attempt in this directory, you are prompted to determine whether you want to overwrite the old files. If you receive this prompt, enter A to overwrite the existing files; this selection ensures that only the files from the current installation are used. 5. In the <i>download directory</i>\TSMClient directory, double-click the spinstall.exe file to start the client installation program.

2. Select a language to use for this installation and click **OK**.
3. If you are prompted to install one, or more, Microsoft C++ redistributable files, the prompt indicates that your node does not have the C++ files that are required by the Windows backup-archive client. Click **Install** to install the files and continue with the client installation; or, click **Cancel** to end the installation process.
4. The backup-archive client installation program starts. On the Welcome screen, click **Next** to begin installing the new client software.
5. Accept or change the default installation directory.
6. Select the installation type: **Typical** or **Custom**.

Option	Description
Typical	<p>A typical installation installs the following components:</p> <ul style="list-style-type: none"> • The backup-archive client GUI files (required to use the Java GUI) • The backup-archive client web files (required to use the web client) • The client API files (as needed by your client and operating system)

Option	Description
Custom	<p>A custom installation installs the same files as a typical installation. However, you can choose to install the following optional components:</p> <ul style="list-style-type: none"> • The API SDK files. These files are only needed if you are developing applications that work with the backup-archive client. • The administrative client command line files. These files are needed if you want to run administrator functions on the Tivoli Storage Manager server. • VMware vStorage API runtime files. These files are required if you are using this client as a data mover to back up VMware virtual machines. Windows data mover nodes must run the 64-bit Windows client.

7. Click **Next**, then click **Install**.
8. When the installer completes the installation, click **Finish**.
9. Verify the installation. Click **Start** > **All Programs** > **Tivoli Storage Manager**. The client components that you installed are shown in the list of Tivoli Storage Manager startable programs. This list includes only the administrative command-line client, backup-archive command-line client, or the backup-archive GUI. The other installable components (the API, SDK, and VMware vStorage API runtime files) do not display in this list.
10. Click the **Backup-Archive GUI** entry in the startable programs list.
 - a. When prompted, type your user ID and password and click **Login**.
 - b. After the GUI starts, click **Help** > **About TSM**. Verify that the version shown is Version 7.1.6.

What to do next

Your previous configuration settings are preserved in the `dsm.opt` file. If you previously used LVSA as the snapshot provider, warning messages are displayed when the command-line client is started. The messages provide instructions to edit the `dsm.opt` file and remove the LVSA options. Removing the unused options is not required, but removing options that have no affect or are not used, can facilitate troubleshooting. If you are using the GUI, the messages are not displayed, but they are logged in the `dsmerror.log` file, which is in the client installation directory, in the `baclient` directory. Messages are issued when any of the following options are included in `dsm.opt`. Some of these options are valid for VSS, and if they are, the messages are displayed and logged only if they contain parameters that are specific to LVSA.

- **snapshotcachelocation**
- **snapshotfsidleretries**
- **snapshotproviderimage**
- **snapshotproviderfs**
- **snapshotcachesize**

You can set VSS options on the **Snapshot** tab in the Preferences Editor. They can also be set by running the online image support and open file support configuration wizards. To use the wizards, start the GUI and click **Utilities** > **Setup**

Wizard. Select the wizards that you want to run, click **Next**, and follow the prompts to make your selections.

Related concepts:

“Troubleshooting problems during installation” on page 25

Reinstalling the Windows client

If you uninstall the Version 7.1.6 Windows client, you can reinstall it if you need to.

About this task

If you reinstall the Windows client into the same directory that it was installed in before, the previous configuration information is detected by the installation program. Because the previous configuration information is detected, the installation process is the same as an upgrade installation; follow the steps in “Upgrading the Windows client” on page 15 to reinstall the Windows client.

If you do not want to preserve the old configuration information, you can remove it. For information about thoroughly removing client settings and files, see the article at [https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli Storage Manager/page/How to completely remove the Backup-Archive client from Microsoft Windows](https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli%20Storage%20Manager/page/How%20to%20completely%20remove%20the%20Backup-Archive%20client%20from%20Microsoft%20Windows).

If you do completely remove all configuration settings and later decide to reinstall the Windows client, follow the steps in “Installing the Windows client for the first time” on page 11. That procedure is the appropriate installation procedure to follow if you reinstall the software into a different directory, or if you reinstall the software on a system that contains no previous configuration information.

Silent installation

The backup-archive client installation program supports silent, unattended installations.

Note: The Microsoft Visual C++ 2010 and 2012 redistributable packages are required to use the backup-archive client. The graphical installation program installs these packages for you. If you are silently installing the client by using MSISEXEC, you must separately install the Microsoft Visual C++ 2010 and 2012 redistributable packages. The packages can be installed before or after the silent installation of the client is completed, but they must be installed before you use the backup-archive client.

Use the following executable files to install the C++ 2010 and 2012 redistributable packages. In the paths that are shown, the *dir* text string represents the drive and directory where you saved the files when you extracted them from the installation package.

Windows 32-bit clients

```
| dir\ISSetupPrerequisites\{270b0954-35ca-4324-bbc6-ba5db9072dad}  
| (contains MS 2010 x86 C++ Runtime - vc_redist_x86.exe)  
  
| dir\ISSetupPrerequisites\{BF2F04CD-3D1F-444e-8960-D08EBD285C3F}  
| (contains MS 2012 x86 C++ Runtime - vc_redist_x86.exe)
```

Windows 64-bit clients

```
dir\ISSetupPrerequisites\{270b0954-35ca-4324-bbc6-ba5db9072dad}  
(contains MS 2010 x86 C++ Runtime - vc_redist_x86.exe)
```

```

|      dir\ISSetupPrerequisites\{BF2F04CD-3D1F-444e-8960-D08EBD285C3F}
|      (contains MS 2012 x86 C++ Runtime - vcredist_x86.exe)
|
|      dir\ISSetupPrerequisites\{7f66a156-bc3b-479d-9703-65db354235cc}
|      (contains MS 2010 x64 C++ Runtime - vcredist_x64.exe)
|
|      dir\ISSetupPrerequisites\{3A3AF437-A9CD-472f-9BC9-8EEDD7505A02}
|      (contains MS 2012 x64 C++ Runtime - vcredist_x64.exe)

```

To install a predefined (custom) dsm.opt file, use the following instructions before you begin the silent installation.

- Place the customized copy of the dsm.opt file in the ... \CONFIG directory located within the installation image, for example:

- (32-bit): C:\tsm_images\TSMClient\Program Files\Tivoli\TSM\config
- (64-bit): C:\tsm_images\TSMClient\Program Files 64\Tivoli\TSM\config

The file must be named *dsm.opt*.

Tip: You cannot add the customized dsm.opt file if you are installing the client from an ISO image. The ISO image must first be copied to a writable location on your computer.

- The installation program copies the predefined dsm.opt file to the .. \BACLIENT directory when BOTH of the following conditions are met:
 - dsm.opt does NOT exist in the .. \BACLIENT directory. The installation program does not copy over an existing dsm.opt file.
 - dsm.opt exists in the .. \CONFIG directory of the installation image, as described earlier.

To silently install the C++ redistributables or the backup-archive client, you must turn off User Account Control (UAC).

To turn off UAC, use either the Windows Control Panel or the MSCONFIG utility.

To turn off UAC by using the Control Panel, go to the Control Panel and find **User Account Control settings**, then set the notification level to **Never Notify**.

To turn off UAC by using the MSCONFIG utility, open a command prompt window and enter **msconfig**. Click **Tools**, select the **Change UAC Settings** tool, and click **Launch**. Set the notification level to **Never Notify** and click **OK**.

After you install the C++ redistributables and the Windows client, remember to turn on UAC.

The C++ redistributables require elevated privileges to install them. Open a command prompt window as follows:

1. Click **Start Menu > All Programs > Accessories > Command Prompt**.
2. Right-click the **Command Prompt** icon to view the properties.
3. Click **Run as administrator**.
4. Click **Continue** in the permission window.
5. Start the product installation by using the command prompt window. On a 32-bit Windows system, install the 32-bit C++ redistributable packages. On a 64-bit Windows system, install the 64-bit C++ redistributable packages.

Silently installing C++ redistributables for the 32-bit client

Run the following command twice. Run it first from the directory where

the C++ 2010 vc_redist_x86.exe file is stored. Then, run it again from the directory where the C++ 2012 vc_redist_x86.exe file is stored.

```
vc_redist_x86.exe /install /quiet /norestart /log logfilename
```

For more information about the vc_redist_x86.exe command, run the following command:

```
vc_redist_x86.exe /?
```

Silently installing C++ redistributables for the 64-bit client

Run the following command twice. Run it first from the directory where the C++ 2010 vc_redist_x86.exe file is stored. Then, run it again from the directory where the C++ 2012 vc_redist_x86.exe file is stored.

```
vc_redist_x86.exe /install /quiet /norestart /log logfilename
```

For more information about the vc_redist_x86.exe command, run the following command:

```
vc_redist_x86.exe /?
```

Run the following command twice. Run it first from the directory where the C++ 2010 vc_redist_x64.exe file is stored. Then, run it again from the directory where the C++ 2012 vc_redist_x64.exe file is stored.

```
vc_redist_x64.exe /install /quiet /norestart /log logfilename
```

For more information about the vc_redist_x86.exe command, run the following command:

```
vc_redist_x64.exe /?
```

Install the Windows backup-archive client. UAC must still be turned off. If it is not turned off, turn off UAC now. Open a command prompt that has elevated privileges.

1. Click **Start Menu > All Programs > Accessories > Command Prompt**.
2. Right-click the **Command Prompt** icon to view the properties.
3. Click **Run as administrator**.
4. Click **Continue** in the permission window.
5. Start the Windows backup-archive client silent installation by using the command prompt window. On a 32-bit Windows system, install the 32-bit clients and API. On a 64-bit Windows system, install the 64-bit clients and API.

Silent client installation on Windows 32-bit systems

When you place a customized version of the **msiexec** command (which calls the Microsoft Software Installer) in a script or batch file, you can perform installations on multiple Windows systems. The following example is a sample command to install the backup-archive command-line client, client GUI, web client, API, and Administrative command-line client. You might need to customize this example to run correctly on your system. Although the command is physically spread across multiple lines in the following example, enter it on a single command line.

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

Silent client installation on Windows 64-bit systems

This example is essentially the same as the 32-bit example. However, it names the 64-bit API feature on the **ADDLOCAL=** parameter.

```
msiexec /i "Z:\tsm_images\TSM_BA_Client\IBM Tivoli Storage Manager
Client.msi" RebootYesNo="No" REBOOT="Suppress" ALLUSERS=1
INSTALLDIR="c:\program files\tivoli\tsm"
ADDLOCAL="BackupArchiveGUI,BackupArchiveWeb,Api64Runtime,
AdministrativeCmd" TRANSFORMS=1033.mst /qn /l*v "c:\log.txt"
```

The descriptions of the silent installation parameters are as follows:

msiexec

Starts the Microsoft Software Installer (MSI) program.

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

"Z:\tsm_images\TSM_BA_Client\IBM Tivoli Storage Manager Client.msi"

Specifies the complete path to the source package. The Z drive is shown in this example. Specify the drive letter for the disk drive, in your configuration, that contains the installation image.

RebootYesNo="No" REBOOT="Suppress"

Under certain conditions, a system reboot might be necessary for the installation to complete successfully. This option causes the installation program to not reboot the system if circumstances would otherwise cause the reboot to occur. While this option is convenient, use it with caution because suppressing the reboot might cause the program to behave in an unpredictable manner. The most common reason that a reboot is required is if the installation was an upgrade to an existing backup-archive client, and the installation was performed while the client programs were running. Therefore, shut down all backup-archive client programs and services before you begin the installation.

ALLUSERS=1

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

INSTALLDIR="c:\program files\tivoli\tsm"

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

For 32-bit client installations:

ADDLOCAL="BackupArchiveGUI,BackupArchiveWeb,ApiRuntime"

For 64-bit client installations:

ADDLOCAL="BackupArchiveGUI,BackupArchiveWeb,Api64Runtime"

Specifies the features to install. Specify all the components on a single line within quotation marks, separated by commas, with no spaces before or after the commas. The installable client features are shown in the following table:

Windows 32-bit client features	Windows 64-bit client features	Feature description
BackupArchiveGUI	BackupArchiveGUI	Graphical user interface
BackupArchiveWeb	BackupArchiveWeb	Backup-archive web client
ApiRuntime	Api32Runtime Api64Runtime	API Runtimes
ApiSdk	ApiSdk	API SDK

Windows 32-bit client features	Windows 64-bit client features	Feature description
AdministrativeCmd	AdministrativeCmd	Administrative Command Line
VMwareBackupTools	VMwareBackupTools	VMware vStorage API runtime files

TRANSFORMS=1033.mst

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

Transform	Language
1028.mst	CHT Traditional Chinese
1029.mst	CSY Czech
1031.mst	DEU German
1033.mst	ENG English
1034.mst	ESP Spanish
1036.mst	FRA French
1038.mst	HUN Hungarian
1040.mst	ITA Italian
1041.mst	JPN Japanese
1042.mst	KOR Korean
1045.mst	PLK Polish
1046.mst	PTB Portuguese
1049.mst	RUS Russian
2052.mst	CHS Simplified Chinese

/qn Specifies to install the product silently.

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

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

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

Related concepts:

"Troubleshooting problems during installation" on page 25

Modifying, repairing, or uninstalling the Windows client

You can modify, repair, or uninstall an existing Windows client.

About this task

Use the Windows control panel to modify, repair, or uninstall the Windows client.

Procedure

1. Click **Start > Control Panel > Uninstall a program**.
2. Select **IBM Tivoli Storage Manager Client** in the list of installed programs.
3. Select the function that you want to perform: **Repair**, **Change**, or **Uninstall**.

Option	Description
Repair	<p>Wait for any in-progress backup-archive client tasks to completed before you repair the Windows client.</p> <p>This option repairs an existing Windows client installation. If you select Repair, the files installed by the installation program are examined to determine whether they have somehow become corrupted. If a file is determined to be corrupted, the repair option attempts to replace it from the DVD or saved installation image. The repair option also repairs missing program short cuts and icons, missing files, and registry keys.</p>
Change	<p>Wait for any in-progress backup-archive client tasks to completed before you modify the Windows client.</p> <p>This option modifies an existing installation. If you select Change, the next screen that is displayed shows Modify as the option for changing installed programs. If you already installed the client and you need to add or remove components, click Change, and select Modify. Choose the icon next to the feature that you want to install or remove and select the appropriate action from the drop-down list. For example, if you selected a typical installation when you installed the client, the administrative client command line interface files are not installed. If you decide that a node needs this interface, select the icon next to Administrative Client Command Line Files and click the This feature will be installed on local hard drive option.</p> <p>Note: This option achieves the same effect as upgrading the client. The difference is that you bypass the initial steps and the installation process begins with the last installation type that you selected. If you want to change the installation type, you can click Back, and select the new installation type; then complete the information as you are prompted for it. Use the information that is provided in “Upgrading the Windows client” on page 15 (start at step 6 on page 17) if you have questions about a prompt.</p>

Option	Description
Uninstall	<p>Wait for any in-progress backup-archive client tasks to completed before you uninstall the Windows client.</p> <p>This option uninstalls the Windows client program. It does not remove any client services. It also does not remove log files, or other items that were created when you configured or used the client. Most of these artifacts remain in the installation directory (Program Files > Tivoli > TSMdirectory), but they can exist anywhere on the disk, depending on what you chose for the installation directory and other options. This option also does not remove files that were copied to the local disk if you extracted the installation files from a compressed distribution file.</p> <p>Leaving these artifacts on disk is not a problem if you want to reinstall the client in the future. However, if you want to more thoroughly remove the client and related files and settings, see the article at https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli Storage Manager/page/How to completely remove the Backup-Archive client from Microsoft Windows.</p> <p>The installation program stops any client services that are running before it uninstalls the software. If you want to stop the services yourself, type the following commands at a command prompt window:</p> <pre>net stop "tsm journal service" net stop "tsm client acceptor" net stop "tsm client scheduler" net stop "tsm remote client agent"</pre> <p>You can also use the Control Panel to stop these services. Their display names match the name used on the command line.</p> <p>Note: The service names shown here are the default names that are set by the installation program. You can change some of these service names when you configure the services using one of the configuration wizards on the Utilities > Setup Wizard menus. If you change the service name, record the name that you specify and use that name to stop the services.</p> <p>If you want to remove any of these services without uninstalling the client, perform the following steps:</p> <ol style="list-style-type: none"> 1. Click Start > All Programs > Tivoli Storage Manager > Backup-Archive GUI. 2. Click Utilities > Setup Wizard. 3. Select and run the wizard for each service that you want to remove. The setup wizard options can also remove the configuration information for online image support and open file support.

Troubleshooting problems during installation

If you are upgrading from a previous version of the backup-archive client and there are client services running (for example, Client Acceptor or Scheduler), you might see an error during the installation.

If there are other Tivoli Storage Manager client services running on any account (for example, Client Acceptor or Scheduler), you might see a request to reboot the system during installation. You must stop all instances of the Tivoli Storage Manager client on all accounts before starting the installation.

You might see the following error during installation:

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

When this error occurs, you must stop the installation. After stopping the installation process, the previous version is no longer installed. Stop the client services and retry the installation process.

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.

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` file used by the previous client installation is available in the default installation directory or the directory or file pointed to by the `DSM_CONFIG` and `DSM_DIR` environment variables, the client accesses this file for configuration information.

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

- “Creating and modifying the client options file” on page 29
- “Register your workstation with a server” on page 91

The following configuration tasks are optional:

- “Create a shared directory options file” on page 31
- “Creating multiple client options files” on page 31
- “Environment variables” on page 32
- “Configuring the language for displaying the Java GUI” on page 33
- “Configuring the web client on Windows systems” on page 34
- “Configuring the scheduler” on page 36
- “Configuring the journal engine service” on page 45
- “Configuring online-image backup support” on page 82
- “Configuring Open File Support” on page 83
- “Creating an include-exclude list” on page 92
- Configuring parallel backups of VMware virtual machines. See “Parallel backups of virtual machines” on page 174

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 Windows systems, the default client options file is named `dsm.opt`.

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 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.

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="c: d:"
domain="ALL-LOCAL -c: -systemstate"
```

- 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 Windows, a file name cannot exceed 255 bytes. Directory names, including the directory delimiter, are also limited to 255 bytes. The maximum combined length for a file name and path name is 5192 bytes. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.

File path and file name limits are shown in Table 4.

- For archive or retrieve operations, the maximum length that you can specify for a path and file name, combined, is 1024 bytes.

Table 4. File path and name limits

MBCS encoding	Path name length limits	File name length limits
1	5192 bytes	255 bytes
2	4092 bytes	127 bytes
3	2728 bytes	85 bytes

In the table, MBCS encoding has these meanings:

Basic Latin

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

Latin extensions

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

Chinese, Japanese, Korean, Vietnamese

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

If you update the client options file while a session is active, you must restart the session to pick up the changes. If you use the client GUI setup wizard to make changes, the changes are effective immediately. If you are not using the client acceptor to manage the scheduler, you must also restart the scheduler.

Related reference:

“Optfile” on page 478

“Query Options” on page 715

Creating and modifying the client options file

The client options file is an editable text file that contains configuration information for the backup-archive client.

About this task

The first time that you start the Windows backup-archive client GUI, the installation program searches for an existing client options file, called `dsm.opt`. If this file is not detected, a client options file configuration wizard starts and prompts you to specify initial client configuration settings. When the wizard completes, it saves the information that you specified in the `dsm.opt` file. By default, the `dsm.opt` file is saved to `C:\Program Files\Tivoli\TSM\baclient`.

The options file must contain the following information to communicate with the server:

- The host name or IP address of the Tivoli Storage Manager server.
- The port number that the server listens on for client communications. A default port number is configured by the client options file configuration wizard. You do not need to override this default port number unless your server is configured to listen on a different port.
- Your client node name. The node name is a name that uniquely identifies your client node. The node name defaults to the short host name of the computer that the client is installed on.

Additional client options can be specified, as needed.

Note: Client options can also be set on the server in a *client option set*. Client options that are defined on the server in a client option set override client options that are set in the client options file.

A sample options file is copied to your disk when you install the backup-archive client. The file is called `dsm.smp`. By default, the `dsm.smp` file is copied to `C:\Program Files\Tivoli\TSM\config\`. You can view the contents of this file to see examples of different options and how they are specified. The file also contains comments that explain syntax conventions for include lists, exclude lists, and wildcard use. You can also use this file as a template for your client options file by editing it and saving it as `dsm.opt` in the `C:\Program Files\Tivoli\TSM\baclient` directory.

After the initial client options file is created, you can modify the client options by adding or changing the options as needed. You can modify the `dsm.opt` file in any of the following ways:

- By running the client options file configuration setup wizard
- By using the client preferences editor
- By editing the dsm.opt file with a text editor program, such as Notepad

Perform the following steps to modify the client options:

Procedure

1. Select a method to modify the file.

Option	Description
Setup wizard	<ol style="list-style-type: none"> 1. Click Start > All Programs > Tivoli Storage Manager > Backup-Archive GUI. 2. Select Utilities > Setup Wizard > Help me configure the TSM Client Options File. On-screen text and online help is available to provide guidance as you navigate through the wizard panels. This client options file configuration wizard offers limited choices and configures only the most basic options.
Preferences editor	<ol style="list-style-type: none"> 1. Click Start > All Programs > Tivoli Storage Manager > Backup-Archive GUI. 2. Select Edit > Client Preferences. Select the tabs in the preferences editor to set client options. Specify the options in the dialog boxes, drop down lists, and other controls. Online help is provided. Click the question mark (?) icon to display the help topics for the online help for the tab that you are editing. You can set more options in the preferences editor than you can set in the setup wizard.
Edit the dsm.opt file	<ol style="list-style-type: none"> 1. Edit the dsm.opt file by using a plain text editor. Each of the options is described in detail in the documentation in "Client options reference" on page 321. This method is the most versatile way to set client options because not all options can be set in the client options file configuration wizard or in the preferences editor. 2. To comment out a setting, insert an asterisk (*) as the first character on the line that you want to comment out. Remove the asterisk to make the commented option active.

2. Save the changes.
 - a. Changes made in the client options file configuration wizard and in the preferences editor are saved and recognized by the client when the wizard completes, or when you exit the preferences editor.
 - b. If you edit the client options file with a text editor while the client is running, you must save the file and restart the client so the changes are detected.

Related concepts:

“Client options reference” on page 321

“Communication options” on page 296

Chapter 11, “Processing options,” on page 295

“Register your workstation with a server” on page 91

Related reference:

“Passwordaccess” on page 480

Create a shared directory options file

A Tivoli Storage Manager server administrator can generate client options files in a shared directory.

Windows clients can access the shared directory, and use the files there to create their own client options file.

Creating a shared directory options file is an optional root user or authorized user task.

Creating multiple client options files

You can create multiple client options files if you must work with multiple servers, or find that you need multiple sets of parameters to do back up or archive tasks.

About this task

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

Procedure

Use one of the following methods to specify or use a different client options file:

- Replace the `dsm.opt` file with the appropriate options file before you start the backup-archive client.

For example, issue the following commands to copy the `a.opt` file to `dsm.opt` and then start the backup-archive client GUI:

```
copy a.opt dsm.opt
dsm
```

- Start the backup-archive client from the command line and use the **optfile** option to specify the options file that you want to use.

For example:

```
dsm -optfile=b.opt
```

- Define the `DSM_CONFIG` environment variable to specify the options file to use before you start a backup-archive client session.

For example:

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

What to do next

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

- Define the DSM_DIR environment variable to point to the directory where all other executable files reside:
SET DSM_DIR=C:\Program Files\Tivoli\TSM\baclient
- Define the DSM_LOG environment variable to point to the directory where dsmerror.log resides:
SET DSM_LOG=C:\Program Files\Tivoli\TSM\baclient

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

Environment variables

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

About this task

You must set the environment variables if you need to run in either of the following environments:

- 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.

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

You need to set four environment variables:

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

DSM_CONFIG

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

DSM_DIR

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

DSM_LOG

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

Ensure that the environment variables meet the following guidelines:

- Include the directory where the executable files (for example, dsm.exe) reside in the current PATH environment variable. If you accepted the default installation directory using the C: drive, you can set this from a command prompt by typing:
SET PATH=C:\Program Files\Tivoli\TSM\baclient
- Specify the fully-qualified path name of your client options file (dsm.opt) using the DSM_CONFIG environment variable:
SET DSM_CONFIG=C:\Program Files\Tivoli\TSM\baclient\dsm.opt

- Define the DSM_DIR environment variable to point to the directory where the client message file dsc*.txt is located:

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

Related reference:

“Optfile” on page 478

Configuring the language for displaying the Java GUI

You can select the language to use for displaying the backup-archive client Java GUI.

About this task

The language that is displayed by the backup-archive client Java GUI is defined by the Windows display locale and not the Windows system locale. For example, if the Windows system and input locale is French, but the display locale is Russian, the language that is displayed by the Java GUI is Russian by default, if the language option is not used.

If you want to display the Java GUI in US English or another language, you can override the default display language by specifying the language option.

Tip: The language option does not affect the web client. The web client displays in the language that is associated with the locale of the browser. If the browser is running in a locale that client does not support, the web client is displayed in US English.

Procedure

Use one of the following methods to configure the language for displaying the Java GUI:

- Add the language *language* option to the client options file (dsm.opt). For example, to set the display language to US English, add the following statement:
language enu
- Complete the following steps in the backup-archive client Java GUI:
 1. In the main window of the Java GUI, click **Edit > Client Preferences**.
 2. Click the **Regional Settings** tab.
 3. Click the **Language** drop-down list and select a language.
 4. Click **OK**.

Related reference:

“Language” on page 455

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.

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

On Windows client nodes, a web client setup wizard is provided in the backup-archive client GUI. You can use the setup wizard to configure the web

client. The options that you select in the wizard are copied to the client-user options file (dsm.opt). You can also add the options directly to the dsm.opt file by editing the file and adding the web client options to it.

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 313

Related tasks:

“Configuring the web client on Windows systems”

Configuring the web client on Windows systems

On Windows systems, you can configure and start the web client by using a wizard that is available in the backup-archive client GUI, or by using both Tivoli Storage Manager and Windows commands.

Procedure

Choose one of the following methods to configure the Windows web client:

Setup method	Procedure
Setup wizard	<ol style="list-style-type: none">1. Start the backup-archive client GUI.2. Click Utilities > Setup Wizard.3. Select the Help me configure the TSM web client check box.4. Click NEXT and follow the wizard instructions to configure the web client options.

Setup method	Procedure
Command prompt	<ol style="list-style-type: none"> 1. Set the following options in the <code>dsm.opt</code> file: <code>managedservices webclient schedule and passwordaccess generate</code>. 2. Install the client acceptor service by entering the following command: <code>dsmcutil install cad /name:"TSM CAD" /node:nodename /password:password /autostart:yes</code> where: <i>TSM CAD</i> a name for the service. The default name is TSM Client Acceptor. <i>nodename</i> is the name of the client node. <i>password</i> is the Tivoli Storage Manager password. <i>/autostart:yes</i> indicates that the CAD service is started when the operating system starts. Start the service by using the Windows net start command. 3. Install the Tivoli Storage Manager remote-client-agent service by entering the following command: <code>dsmcutil install remoteagent /name:"TSM AGENT" /node:nodename /password:password /partnername:"TSM CAD"</code> where: <ul style="list-style-type: none"> • <i>TSM AGENT</i> is a name for the remote-client-agent service. The default service name is TSM Remote Client Agent. • <i>nodename</i> is the name of the client node. • <i>password</i> is the Tivoli Storage Manager password. • <i>TSM CAD</i> is the service-partner name. This name must match the service name that you specified when you installed the client acceptor service. The default name is TSM Client Acceptor. Do not start the TSM Remote Client Agent service from the Control Panel > Administrative Tools > Services view, or by using the net start command. The client acceptor service starts the remote client agent when it is needed.

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 309

"Web client options" on page 313

Related tasks:

"Starting a web client session" on page 115

Related reference:

"Httpport" on page 423

"Passwordaccess" on page 480

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 client options file (`dsm.opt`).

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 5. 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 service is started as a Windows service	Started with command dsmc 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 dsmc sched is started.

Table 5. Client acceptor-managed services versus traditional scheduler services (continued)

Client acceptor-managed services	Tivoli Storage Manager traditional scheduler services
Cannot be used with SESSIONINITiation=SERVEROnly 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 scheduler on the Windows client:

1. In the backup-archive client GUI, click **Utilities > Setup Wizard > Help me configure the TSM Client Scheduler** and click **Next**.
2. Read the information in the TSM Scheduler Wizard page and click **Next**.
3. In the Scheduler Task page, select **Install a new or additional scheduler** and click **Next**.
4. In the TSM Scheduler Name and Location page, specify a name for the client acceptor service that you want to manage the scheduler. Then, select **Use the Client Acceptor daemon (CAD) to manage the scheduler** and click **Next**.

5. If the client acceptor is already installed for use by the web client, select that name of the client acceptor from the drop-down list in the Web Service Name page. Otherwise, type the name that you want to assign to this client acceptor. The default name is **TSM Client Acceptor**. Click **Next**.
6. Follow the instructions on the remaining screens to complete the configuration. Use the following information to help you complete the wizard pages:
 - If the **sessioninitiation** option is set to **serveronly** in the client options file (dsm.opt), the client configuration wizard and scheduler service might be unable to initiate authentication with the Tivoli Storage Manager server. To avoid this problem, ensure that the **Contact the TSM Server to validate password** check box on the TSM Authentication page is cleared.
 - For the client acceptor-managed scheduler, select **Manually when I explicitly start the service** in the Service login options page.
7. Start the client acceptor service from the Services Control Panel, but do not start the scheduler service. The scheduler service is started and stopped automatically by the client acceptor service as needed.

Tip:

- You can also use the **managedservices** option in the client options file (dsm.opt) 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.opt file.
- If you do not use the client acceptor to manage the scheduler, select **Automatically when Windows boots** in the Service login options window. This setting starts the service automatically when Windows starts so that your schedules are run automatically. Alternatively, you can use the Services Control Panel or the **net start** command to start the Scheduler service.
- You can also use the Scheduler Service Configuration utility (dsmcutil.exe) to configure the scheduler. The Scheduler Service Configuration utility must be run from an account that belongs to the Administrator/Domain Administrator group. You can start multiple client scheduler services on your system.

Related concepts:

"Web client configuration overview" on page 33

"Enable or disable scheduled commands" on page 260

"Scheduling options" on page 309

Related tasks:

"Setting the client scheduler process to run as a background task and start automatically at startup" on page 252

Related reference:

"Cadlistenonport" on page 338

"Managedservices" on page 457

"Sessioninitiation" on page 523

Starting the client scheduler

To start the client scheduler, use the Services Control Panel or the **net start** command.

About this task

To avoid problems, do not run the client scheduler on the command line. The command line does not run the scheduler as a background service.

When you start the client scheduler, it runs continuously until you close the window, shut down your system, or log out of your system. If you are running the Scheduler Service, the scheduler runs until the system is shutdown or you explicitly stop it using the services control panel.

Related concepts:

Chapter 11, "Processing options," on page 295

Scheduling events using the GUI

This task guides you through the steps to schedule events using the GUI.

Procedure

1. From the backup-archive client GUI main window, click **Utilities > Setup Wizard**. The Client Configuration Assistant appears.
2. Select the **Help me configure the TSM Client Acceptor Daemon and TSM Client Scheduler** and click the **OK** button. The Scheduler Wizard panel appears.
3. Select the task that you want to perform. You can install a new client scheduler, update the settings for a scheduler, or remove a scheduler.
4. Complete each panel and click the right arrow to continue. To go back to a previous panel, click the left arrow.

What to do next

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

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 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 installing the scheduler using the setup wizard or **dsmcutil**, and the Tivoli Storage Manager server is behind a firewall, the node password will not get stored on the client workstation. As a result, the scheduler service might be unable to authenticate to the server when the server contacts the client to run a schedule. In this case, you can

run the scheduler from the command line (`dsmc schedule`), wait until a scheduled operation starts, and enter the password for your node when prompted. After you enter the password for your node, restart the scheduler service. You can also use the following **dsmcutil** command to write the password into the registry:

```
dsmcutil updatepw /node:nnn /password:ppp /validate:no
```

If **sessioninitiation** option is set to *serveronly* in your client options file (`dsm.opt`), the client setup wizard and scheduler service is unable to initiate authentication with the Tivoli Storage Manager server. To avoid this problem, when configuring the client scheduler using the setup wizard, ensure that the **Contact the TSM Server to validate password** checkbox on the TSM Authentication page is unchecked.

A similar problem can occur if an encryption key is required for backup operations. In this case, you can run the scheduler from the command line (`dsmc schedule`), wait until a scheduled backup starts, and enter the encryption key when prompted. After the password and encryption key are updated, you must restart the scheduler.

4. When configuring the 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.
5. 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 Tivoli Storage Manager 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 36

Related reference:

“Sessioninitiation” on page 523

“Tcpadminport” on page 554

“Tcpport” on page 559

“Webports” on page 622

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

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

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

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

1. Obtain the Tivoli Storage Manager server self-signed certificate (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 GSKit command-line utility (gsk8capicmd for 32-bit clients or gsk8capicmd_64 for 64-bit clients) 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.

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 GSKit command-line utility (gsk8capicmd for 32-bit clients or gsk8capicmd_64 for 64-bit clients) 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.

The **gsk8capicmd** and **gsk8capicmd_64** commands are provided by Global Security Kit (GSKit). The backup-archive client automatically installs GSKit in C:\Program Files\Common Files\Tivoli\api64\gsk8.

Before you set up the server certificate on the client, follow these steps:

1. Open a command window and change the directory to the backup-archive client directory, for example: `cd "C:\Program Files\Tivoli\TSM\baclient"`
2. Append the GSKit binary path and library path to the PATH environment variable, for example:

```
set PATH=C:\Program Files\Common Files\Tivoli\TSM\api64\gsk8\bin\;  
C:\Program Files\Common Files\Tivoli\TSM\api64\gsk8\lib64;%PATH%
```

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

```
gsk8capicmd_64 -keydb -create -populate  
-db dsmcert.kdb -pw password -stash
```

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

If you use a self-signed certificate

Each Tivoli Storage Manager server generates its own certificate. The certificate has a fixed file name of either `cert.arm` or `cert256.arm`. The certificate file is stored on the server workstation in the server instance directory, for example, `C:\Program Files\tivoli\tsm\server1\cert256.arm`. If the certificate file does not exist and you specify the **SSLTCPPORT** or **SSLTCPADMINPORT** 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:

```
gsk8capicmd_64 -cert -add -db dsmcert.kdb -stashed  
-label "TSM server <servername> self-signed key"  
-file <path_to_cert256.arm> -format ascii
```

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 44.

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:

```
gsk8capicmd_64 -cert -add -db dsmcert.kdb -stashed  
-label "XYZ Certificate Authority" -file <path to CA root certificate>  
-format ascii
```

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. Different certificates must have different labels. The label names are not important, but use meaningful names. Also, more than one CA root certificate can be added to the client key database.
3. If you do not run the preceding commands from the backup-archive client directory, you must copy dsmcert.kdb and dsmcert.sth into that directory.
4. For performance reasons, use SSL only for sessions where it is needed. Consider adding more processor resources on the Tivoli Storage Manager server system to manage the increased requirements.
5. 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.

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 **tcpport** option. It is important to understand that the server is normally set up for SSL connections on a different port. In other words, two ports are opened on the server:

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

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

If the value of **tcpport** is incorrect, the client cannot connect to the server. 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 542

"Sslfipsmode" on page 543

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.

Configuring the journal engine service

Journal-based backup can be used for all Windows clients. If you install the journal engine service and it is running, then by default the **incremental** command automatically performs a journal-based backup on selected file systems that are being monitored by the journal engine service.

About this task

Journal-based backup is enabled by installing and configuring the Tivoli Storage Manager journal service. You can install the journal service with the GUI Setup wizard or with the **dsmsutil** command. Basic journal service configuration can be done with the GUI Setup wizard, more advanced configuration can be done by editing the journal service configuration file, `tsmjbbd.ini`.

Tip: The default location for journal service configuration file is `C:\Program Files\Tivoli\TSM\baclient\tsmjbbd.ini`. If this is the first time you are

configuring the journal engine service and a copy of `tsmjbbd.ini` does not already exist, copy the sample file `C:\Program Files\Tivoli\TSM\config\tsmjbbd.ini.smp` to `C:\Program Files\Tivoli\TSM\baclient\tsmjbbd.ini`.

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

Procedure

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

Results

Journal service configuration settings are stored in the journal configuration file `tsmjbbd.ini`. This file can be installed and configured with the GUI setup wizard or be manually edited.

Follow these steps to set up multiple journal services:

1. Create and set up a separate journal configuration file (`tsmjbbd.ini`) for each journal service to be installed. Each configuration file must specify a different `JournalPipe` value, and must also specify different drives to journal, so that the two services do not interfere with each other. Multiple journal services journaling the same drive causes problems. The different services attempts to write to the same journal database unless this is specifically overridden by specifying different journal directories in the different configuration files.
2. Install the multiple journal services using the **dsmcutil.exe** tool. Use distinct names for each service, and specify the `/JBBCONFIGFILE` option to identify the `tsmjbbd.ini` to be used for that particular journal instance. For example:

```
dsmcutil install journal /name:"TSM Journal Service 1"
/JBBCONFIGFILE:c:\journalconfig\tsmjbbd1.ini
dsmcutil install journal /name:"TSM Journal Service 2"
/JBBCONFIGFILE:d:\journalconfig\tsmjbbd2.ini
```

Note: In Uniform Naming Convention (UNC) format, the **jbbconfigfile** path must contain a drive letter. In the following UNC format example, the path contains the drive letter `D$`: `\\computer7\D$\journalconfig\tsmjbbd1.ini`

3. Different backup clients (based on the distinct `dsm.opt` file used) can now connect to the desired journal service by specifying the appropriate `JournalPipe` option in the appropriate `dsm.opt`, which corresponds to the `JournalPipe` journal service setting.

Note:

1. Each journal service instance is associated to only one backup-archive client node name. Changing the association requires a restart of the journal service to recognize the new association.
2. You cannot use network and removable file systems.

Configuration settings that you apply when the journal service is started and any changes you make while the journal service is running are applied without having to restart the service. This also applies to the journal exclude list. However, some settings for journaled file systems do not take effect until the file system is brought offline and then back online.

File systems can be taken online (added) or offline (removed) without stopping and restarting the journal service. You can bring a file system offline by removing it from the list of journaled file systems in the journal configuration file `tsmjbbd.ini`, or by shutting down the journal service. You can bring a file system back online by adding it to the list of journaled file systems in the journal configuration file `tsmjbbd.ini` or by starting (restarting) the journal service.

Attention: If you take a file system offline without setting the **PreserveDbOnExit** value of 1, the journaled file system journal database is deleted.

PreserveDbOnExit=1 specifies that the journaled file system journal database is not deleted when the journal file system goes offline. The database is also valid when the journal file system comes back online.

The following is the syntax for stanza and stanza settings:

Syntax for stanzas:

[StanzaName]

Syntax for stanza settings:

stanzaSetting=value

Note:

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

Related concepts:

"Journal-based backup" on page 138

JournalSettings stanza (Windows)

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

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

Syntax for `JournalSettings` stanza:

[JournalSettings]

Syntax for stanza settings:

JournalSettings=*value*

You can specify the following `JournalSettings` values:

JournalPipe=*pipename*

Specifies the pipe name of the journal service session manager to which backup clients initially connect, when establishing a journal-based backup session. This setting is used in conjunction with the backup client option of the same name. The default pipe name is `\\.\pipe\jnlSessionMgr1`. For example, in `dsm.opt`:

```
JournalPipe \\.\pipe\jnlSessionMgr1
```

Under tsmjbbd.ini [JournalSettings] stanza:

```
JournalPipe=\\.\pipe\jnlSessionMgr1
```

Note: The same pipe name must be specified by the client using the JournalPipe option.

NlsRepos

Specifies the National Language Support repository the journal service uses for generating messages. Since the journal service is non-interactive, this only applies to messages written to the journal error log. The default value is dscameng.txt. For example:

```
NlsRepos=dscenu.txt
```

ErrorLog

Specifies the log file where detailed error messages generated by the journal service are written. Note that less detailed error and informational messages are written to the Windows application event log as well. The default value is jbberror.log. For example:

```
ErrorLog=jbberror.log
```

In Uniform Naming Convention (UNC) format, the path must contain a drive letter. In the following UNC format example, the path contains the drive letter D\$: \\computer7\D\$\temp\jbberror.log.

JournalDir

Specifies the directory where journal database files are stored and written. The default directory is the journal service installation directory. You can specify different journal locations for each file system being journaled. This is useful when running in a clustered environment because the location of the journal must be accessible by each workstation in the cluster running the journal service. Typically the journal for local resources being journaled resides in the same location and the journal for shared cluster resources (which can move from workstation to workstation) is located on the shared resource to ensure that it is accessible to both workstations.

By default, this setting applies to all journaled file systems but can be overridden by an override stanza for each journal file system. If the default value is a fully qualified path (for example c:\tsmjournal), all journal database files are written to the specified directory. If the default value does not specify a drive letter, (for example \tsmjournal) the journal database files for each journal file system are written to the specified directory on each journal file system.

In Uniform Naming Convention (UNC) format, the path must contain a drive letter. In the following UNC format example, the path contains the drive letter D\$: \\computer7\D\$\temp\tsmjournal.

The following is an example configuration stanza:

```
[JournalSettings]
;
; Store all resources in one location unless overridden
; by an override stanza
;
JournalDir=c:\tsmjournal
;
;
[JournaledFileSystemSettings.D:\]
```

```

;
; Journal for d: only is in location specified below
;
JournalDir=d:\tsmjournal

```

Note: Changes to this setting do not take effect until the journaled file systems are brought online.

JournalExcludeList stanza

This list of exclude statements filters changes from being recorded in the journal database. Changes to objects which match statements in this stanza are ignored and are not recorded in the journal database.

Note:

1. Excluding files from the journal has no bearing on those files being excluded by the backup client, other than preventing the files from being sent to the backup client to be processed during journal-based backup. A file that is not excluded from the journal should still be excluded by the backup-archive client, if there is a matching exclude statement in the client options file.
2. The journal service only provides a subset of the INCLUDE/EXCLUDE function provided by the backup-archive client. The journal service does not support INCLUDE statements and it does not support the *exclude.dir* option.

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

The following are examples of equivalent journal exclude statements:

```
dsm.opt: tsmjbbd.ini
```

```

EXCLUDE c:\testdir\...\* c:\testdir\*
EXCLUDE.DIR c:\testdir\test* c:\testdir\test*\*

```

The following pattern matching meta characters are supported:

- % Matches exactly one character.
- * Matches zero or more characters.
- %EnvVar% Expands environment variable.

The following is an exclude statement syntax example:

```

[JournalExcludeList]
%SystemRoot%\System32\Config\*
%SystemDrive%\Adsm.Sys\*
%TEMP%\*
%TMP%\*
c:\excluedir\*
c:\dir1\excludefile
*.*\*.tmp

```

Note: The `c:\excluedir*` statement matches the entire tree including subdirectories and files.

JournalFileSystemSettings stanza

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

The following is the syntax for the JournalFileSystemSettings stanza:

Syntax for *JournaledExceptionSettings* stanza:
[JournaledExceptionSettings]

Syntax for stanza settings:
JournaledExceptionSetting=value

You can specify the following *JournaledExceptionSettings* values:

DirNotifyBufferSize

Specifies the size of the buffer to record change notifications for a particular journal file system. You might need to increase this value for journaled file systems that generate a very large volume of change activity. The buffer size is limited by memory. The default value is 16 KB.

JournaledExceptions

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 service to run. Journaled file systems can be added or removed online without having to restart the service. For example:

JournaledExceptions=c: d:

JournalDbSize

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

JournalDbSize=0x10000000

NotifyBufferSize

Specifies the size of the memory buffer receiving file system change notifications for a particular journal file system. You might need to increase this value for journaled file systems that generate a very large volume of change activity. The buffer size is limited by memory. The default value is 32 KB. For example:

NotifyBufferSize=0x00008000

NotifyFilter

Specifies what file system change actions generate notifications to the journal service. ***NotifyFilter*** applies to file changes and directory modifications. Directory name changes, such as deletions and creations, are always tracked regardless of the filter value. Multiple actions can be monitored by combining (adding) values together. The default value is 0x11F (File and Dir Name, Attrib, Size, Last Write, and security Changes). You can also use the Tivoli Storage Manager Journal Engine Wizard to specify that any or all of these actions are monitored. Supported values are:

Value type	Decimal	Hex
File Name	1	0x001
Dir Name	2	0x002
Attribute	4	0x004
File size*	8	0x008
Last Write Time*	16	0x010
Last Access Time	32	0x020
Create Time	64	0x040
Security (ACL)	256	0x100

The asterisk (*) indicates that notification might be deferred until disk write cache is flushed. Name changes are object creations, deletions, or renames.

Example:

```
NotifyFilter=0x107
```

PreserveDbOnExit setting

This setting allows a journal to remain valid when a journaled file system goes offline and comes back online. This is useful for preserving the journal during system reboots, cluster failovers, and resource movement.

File systems go offline when the journal service stops or when the file system is removed from the configuration file. File systems come back online when the journal service is started or when the file system is added to the configuration file.

This setting allows a journal-based backup to continue processing when the service is restarted (or the file system comes back online) without performing a full incremental backup.

Note: Any change activity which occurs while the journal service is not running (or the file system is offline) is not recorded in the journal.

In a clustered environment, shared resources can move to different workstations in the cluster. The journal service running on each workstation in the cluster must include these shared resources in the list of journaled file systems. The journal service running on the workstation which currently owns the resource actively journals the shared resource while other journal services on workstations in the cluster which do not own the resource must defer journaling until the resource becomes available (or is moved to that workstation). The configuration settings *deferFSMonStart*, *deferRetryInterval*, and *logFSErrors* allows deferment for a file system until the file system is available and accessible.

A value of 1 specifies that the journaled file system journal database is not deleted when the journal file system goes offline. The database 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.

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

```
[JournaledFileSystemSettings.D:\]  
;  
; Do not delete the journal when D:\ goes offline  
;  
PreserveDbOnExit=1
```

deferFSMonStart setting

This setting defers an attempt to begin monitoring a file system in the following cases:

- When the specified journaled file system is not valid or available

- The journal directory for the specified journaled file system cannot be accessed or created

Resources are checked at the interval you specify using the *deferRetryInterval* setting.

The *deferFSMonStart* setting is most commonly used in a cluster environment where shared resources might move to different workstations in the cluster.

A value of 1 indicates that the setting is on. A value of 0 indicates that the setting is off. The default value is off (set to 0) .

deferRetryInterval setting

This setting specifies the value in seconds that a deferred file systems with the *deferRetryInterval* setting enabled are checked for availability and brought online. The default value is 1 second.

logFSErrors setting

This setting specifies whether errors encountered while accessing a journaled file system or journal directory are logged in the *jbberror.log* and the event log.

Use the *logFSErrors* setting with the *deferFSMonStart* setting to prevent excessive *File System unavailable* messages from being logged when bringing a journaled file system online is deferred. The first error which causes the file system to be deferred is logged. Subsequent errors are not logged. A value of 1 indicates that the setting is on. A value of 0 indicates that the setting is off.

An example to defer journaling until the file system journal directories are valid is:

```
[JournalSettings]
;
; Place journal files in directory on each journaled file system
;
journalDir=\tsmjournal

[JournaledeFileSystemSettings]
;
;journal c:, d:, and f:
;
JournaledeFileSystems=c: d: d:\mountpoint f:
;
; Override stanza to defer starting journaling for f:\
; until it is a valid file system

[JournalFileSystemSettings.f:\]
;
; Keep database valid if file system goes offline
;
PreserveDBOnExit=1
;
; Defer journaling until file system and journal directory
; are valid
;
deferFSMonStart=1
;
; Attempt to start journaling every 120 seconds when deferred
;
deferRetryInterval=120
;
;Do not log excessive resource unavailable messages
;
logFsErrors=0
```


Related concepts:

“Overriding stanzas”

Overriding stanzas

Any setting in the *JournalFileSystemSettings* stanza, except for the buffer sizes, can be overridden for a particular journaled file system by creating an override stanza.

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

Syntax for JournalFileSystemSettings stanza:

[JournalFileSystemSettings.fs]

Syntax for stanza settings:

JournalFileSystemSetting=override value

Example:

```
[JournalFileSystemSettings.C:\]  
NotifyBuffer=0x0020000  
NotifyFilter=0x107
```

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
- Subfile backup operations
- 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 6.

Table 6. 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” on page 57

Related reference:

“Deduplication” on page 360

“Exclude options” on page 400

“Dedupcachepath” on page 359

“Dedupcachesize” on page 359

“Enablededupcache” on page 389

“Ieobjtype” on page 425

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 53 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	<ul style="list-style-type: none">• Add the deduplication yes option to the <code>dsm.opt</code> file.
Preferences editor	<ol style="list-style-type: none">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.opt` 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.opt 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 : SPACEMGPOOL
Copy Group
Copy Group Name.....: STANDARD
Copy Type.....: Backup
Copy Frequency.....: 0 day(s)
Versions Data Exists...: 2 version(s)
Versions Data Deleted..: 1 version(s)
Retain Extra Versions...: 30 day(s)
Retain Only Version....: 60 day(s)
Copy Serialization.....: Shared Static
Copy Mode.....: Modified
Copy Destination.....: AVIFILEPOOL
Lan Free Destination...: NO
Deduplicate Data.....: YES

Copy Group Name.....: STANDARD
Copy Type.....: Archive
Copy Frequency.....: Cmd
Retain Version.....: 365 day(s)
Copy Serialization.....: Shared Static
Copy Mode.....: Absolute
Retain Initiation.....: Create
Retain Minimum.....: 65534 day(s)
Copy Destination.....: FILEPOOL
```

Lan Free Destination...: NO
Deduplicate Data.....: YES

ANS1900I Return code is 0.

Related concepts:

“Client-side data deduplication” on page 53

Related reference:

“Deduplication” on page 360

“Enablededupcache” on page 389

“Exclude options” on page 400

 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, images, system state objects, and ASR 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**
 - **System state**
 - **ASR**
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.

For ASR and system state, this field is filled out automatically. When you specify the image object type, the drive letter must be followed by `:**`. For example, to exclude drive E:, enter the following pattern:

`E:**`

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.opt` file:

1. Add the deduplication yes option
2. Exclude client-side data deduplication for image backup of drive. For example, to exclude drive E:, add the following statement: `EXCLUDE.DEDUP E:**`
`IEOBJTYPE=Image` to `dsm.opt`.

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 53

Related reference:

“Deduplication” on page 360

“Enablededupcache” on page 389

“Exclude options” on page 400

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 228

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 the secondary 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 the secondary server.

Each time the client logs on to the Tivoli Storage Manager server, the client attempts to contact the primary server. If the primary 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 about the secondary server that is sent to the client and saved to the client options file (dsm.opt):

```
*** 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
MYPRIMARYServer SERVER1
*** 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 475.

- 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 Tivoli Storage Manager 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 65

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 **dsmc 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.opt` 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
```

```
MYREPLICATIONServer TARGET
MYPRIMARYSERVERNAME SERVER1
```

3. Save and close the client options file.
4. 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 60

Related tasks:

“Restoring or retrieving data during a failover” on page 228

Related reference:

“Forcefailover” on page 419

“Myprimaryserver” on page 468

“Myreplicationserver” on page 469

“Nrtablepath” on page 475

“Replserverguid” on page 500

“Replservername” on page 501

“Replsslport” on page 503

“Repltcpport” on page 504

“Repltcpserveraddress” on page 505

“Usereplicationfailover” on page 570

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 228

Related reference:

“Nrtablepath” on page 475

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 65

Related reference:

“Usereplicationfailover” on page 570

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 options file (dsm.opt).
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 419

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. You can also use the Client Configuration wizard for Tivoli Storage Manager FastBack.

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 8

“Client configuration wizard for Tivoli Storage Manager FastBack” on page 9

“Configuring the backup-archive client to protect FastBack client data” on page 69

Related reference:

“Set Password” on page 781

Configuring the backup-archive client to protect FastBack client data

You can configure the backup-archive client to protect FastBack client data by using the client configuration wizard.

Before you can use the Tivoli Storage Manager Client Configuration wizard for FastBack, you must complete the following tasks:

- Ensure that either the Tivoli Storage Manager FastBack server, or the Tivoli Storage Manager FastBack Disaster Recovery Hub, is installed and configured for short-term data retention.
- Also ensure that at least one snapshot has been taken.
- Ensure that the backup-archive client is properly configured with a Tivoli Storage Manager server. Also make sure that the client acceptor service (dsmcad.exe) is running. You can use the Tivoli Storage Manager Client Configuration wizard in the backup-archive client GUI, after you install the backup-archive client.
- Complete a one-time post-installation setup for these purposes:
 - To specify a Tivoli Storage Manager FastBack user name and password to be used by the wizard, to query and mount volumes from the Tivoli Storage Manager FastBack repository
 - To run Tivoli Storage Manager scheduler scripts
- Set up the FastBack credentials file. The user ID that you specify must have Tivoli Storage Manager FastBack administrative authority.
 1. Configure the user ID and password. Run the following command on the workstation where the backup-archive client and FastBack server or Disaster Recovery Hub are installed:

```
cd <TSM_FastBack_install_location>\FastBack\shell
```

where <TSM_FastBack_install_location> is the directory location where the Tivoli Storage Manager FastBack client is installed.
 2. If it does not exist, create a folder called **FastbackTSMScripts** under the system drive of the workstation, using the following command:

```
mkdir <machine_system_drive>:\FastbackTSMScripts
```
 3. Run the **fastbackshell** command:

```
FastBackShell -c encrypt -u userName -d domain -p password -f <machine_system_drive>:\FastbackTSMScripts\credential.txt
```

The following options are used in the preceding command example:
 - -u specifies the Tivoli Storage Manager FastBack administrator user name.
 - -p specifies the Tivoli Storage Manager FastBack administrator password.
 - -d specifies the Tivoli Storage Manager FastBack domain for the user name.
 - -f specifies the output file where the encrypted credentials are to be written.

Important: The credentials file must be generated with the name "credential.txt". The credentials file must also be located in the FastbackTSMScripts directory of the system drive of the workstation, for the wizard to function properly.

You can use the client configuration wizard in the backup-archive client Java GUI or the backup-archive web client.

Follow these steps to use the client configuration wizard in the Java GUI:

1. Ensure that the backup-archive client is properly configured with a Tivoli Storage Manager server.
2. The configuration wizard starts automatically to create the configuration file.
3. Follow the instructions on the panel to complete the wizard.
4. From the backup-archive client GUI main window, select **Utilities > Setup Wizard**.
5. From the welcome page, select **Help me configure the client to protect FastBack client data** and click **Next**.
6. Use the wizard to complete the configuration process.

Follow these steps to start the client configuration wizard in the web client:

1. Ensure that the web client is properly configured with a Tivoli Storage Manager server, and that the Tivoli Storage Manager client acceptor service is running.
To configure the web client, follow these steps:
 - a. From the backup-archive client GUI main window in the Java GUI, click **Utilities > Setup Wizard**.
 - b. From the welcome page, select **Help me configure the TSM Web Client** and click **Next**. Follow the instructions on the panel to complete the wizard.
2. Start the web client. In your web browser, specify the client node name and port number where the client acceptor service is running.
For example: `http://<machine_name_or_ip_address>:1585`
3. From the backup-archive client GUI main window, click **Utilities > Setup Wizard**.
4. From the welcome page, select **Help me configure the client to protect FastBack client data**, and click **Next**.
5. Use the wizard to complete the configuration process.

Related concepts:

“Client configuration wizard for Tivoli Storage Manager FastBack” on page 9

Configuring the backup-archive client in a cluster server environment

You can install the backup-archive client software locally on each node of a Microsoft Cluster Server (MSCS) or Veritas Cluster Server (VCS) environment cluster.

You can use the backup-archive client in a VCS environment on the supported Windows Server platforms.

You can also install and configure the Scheduler Service for each cluster node to manage all local disks and each cluster group containing physical disk resources.

For example, MSCS cluster **mcs-cluster** contains two nodes: **node-1** and **node-2**, and two cluster groups containing physical disk resources: **group-a** and **group-b**. In this case, an instance of the Tivoli Storage Manager Backup-Archive Scheduler Service should be installed for **node-1**, **node-2**, **group-a**, and **group-b**. This ensures that proper resources are available to the Backup-Archive client when disks move (or fail) between cluster nodes.

The `clusternode` option ensures that the client manages backup data logically, regardless of which cluster node backs up a cluster disk resource. Use this option for client nodes that process cluster disk resources, and not local resources.

Note: You must set the `clusternode:` option to `yes` for all Tivoli Storage Manager-managed cluster operations. Inconsistent use of the `clusternode` option for a given Tivoli Storage Manager cluster node name can cause the client to invalidate the cluster node name encrypted password, and prompt the user to reenter the password during the next backup-archive client program invocation.

Use the `optfile` option to properly call the correct (cluster) `dsm.opt` for all client programs to ensure proper functionality for cluster related operations.

How you install and configure the backup-archive client in a cluster environment depends on the cluster server technology used (MSCS or VCS) and the operating system being used by the nodes in the cluster.

Related reference:

“`Optfile`” on page 478

Protecting data in MSCS clusters (Windows Server clients)

A client configuration wizard is used on nodes in an MSCS cluster environment to automate and simplify the configuration of the backup-archive client to protect cluster disk groups. The wizard can only be used on nodes that run supported Windows Server clients as their operating system.

Configuring cluster protection (Windows Server clients)

Use the Tivoli Storage Manager cluster wizard to configure the backup-archive client to protect cluster resources. The wizard gathers the information that is needed so the backup-archive client can protect cluster resources, and log on to the server.

Before you begin

Before you run the cluster configuration wizard, perform the following steps:

- Install the backup-archive client on each node in the cluster. All backup-archive clients must be the same version of the software and all clients must be installed in the same directory on each node.
- Register the nodes that you are going to run the cluster configuration wizard on. On the Tivoli Storage Manager server, use the administrative command-line client and register the node by using the **register node** command.
- Make sure that the cluster groups that will be configured are owned by the system that will run the cluster wizard. This ensures that the backup-archive client files (options file, error log, schedule log) can be created/updated on the cluster drives.

About this task

You run the wizard on only one node in the cluster; the wizard creates the necessary services on all nodes in the cluster.

The wizard can configure only one cluster group at a time. If you have multiple cluster groups to protect, run the wizard as many times as needed to configure the client to back up each group.

Procedure

1. Run `dsm.exe` to start the Java GUI.
2. In the GUI, click **Utilities > Setup Wizard > Help me protect my cluster**.

3. Choose **Configure a new or additional cluster group**, the first time you run the wizard on a node. On subsequent wizard sessions, you can choose to update a previously configured cluster group or to remove a saved configuration.
4. Select the name of the cluster group that you want to protect.
5. Select the disks in the cluster group that you want to protect. You cannot use the wizard to back up the quorum drive.
6. Specify the disk location where you want the wizard to store the client options file (dsm.opt) that the wizard creates. The client options file must be on one of the drives in the cluster group that you selected in step 4. If a client options file exists at this location, you are prompted to overwrite it, or choose a new directory.
7. Specify a name for the Tivoli Storage Manager Scheduler that will be used to perform the backups. Select **Use the Client Acceptor Daemon to manage the scheduler**, if you want the client acceptor daemon to manage the scheduler.
8. Specify the node name for the cluster node and the password that is used to log on to the Tivoli Storage Manager Server. By default, the option to have the password validated by the server is selected. Clear this option if you do not want the password validated.
9. Specify the account that the scheduler and client acceptor daemon services log on as when the services are started. Specify whether you want to start the service manually, or when the node is booted.
10. Specify the names and location of the client schedule log file and error log file. By default, event logging is enabled. Clear this option if you do not want to log events.

To ensure that any node can perform backups if any other node fails, the wizard copies the registry data to all nodes in the cluster.

Configure the web client in a cluster environment

To use the web client in a cluster environment, you must configure the backup-archive client GUI to run in a cluster environment.

See “Configuring cluster protection (Windows Server clients)” on page 71 for detailed information about installing and configuring the backup-archive client in a MSCS or VCS environment.

Configure the web client to process cluster disk resources

After installing and configuring the backup-archive client in a MSCS or VCS environment, there are some steps you must perform to process cluster disk resources.

Step 1: Identify the cluster groups to manage:

Use the Microsoft Cluster Administrator utility or VCS Configuration editor to determine which groups contain physical disk resources for the backup-archive client to process.

Register a unique node name on the backup server for each group.

For example, an MSCS cluster named **mscs-cluster** contains the following groups and resources:

- **group-a** - Contains physical disk **q:** (quorum), and physical disk **r:**

Note: VCS does not have quorum disk.

- **group-b** - Contains physical disk **s:**, and physical disk **t:**

In this example, the administrator registers two node names: **mcs-cluster-group-a** and **mcs-cluster-group-b**. For example, to register **mcs-cluster-group-a** the administrator can enter the following command:

```
register node mcs-cluster-group-a password
```

Step 2: Configure the client options file:

Configure the client options file (dsm.opt) for each cluster group. Locate the option file on one of the disk drives that are owned by the cluster group.

About this task

For example, the option file for **mcs-cluster-group-a** resides on either **q:** or **r:**.

Procedure

To configure the dsm.opt file for each cluster group, specify the following options:

nodename

Specify a unique name. For example: **mcs-cluster-group-a**

domain Specify the drive letters for the drives that are managed by the group. For example: **q: r:**

See “Frequently asked questions” on page 79 for information on how to add a cluster drive to an existing Tivoli Storage Manager cluster scheduler service resource for backup.

clusternode

Specify the Yes value. If you set the clusternode option to Yes, the client does the following actions:

1. Checks for a cluster environment (MSCS or VCS).
2. Uses the cluster name instead node name for file space naming and encryption. This action allows the use of one password file for all nodes in the cluster.
3. Builds a list of shared volumes and works only with shared volumes. Backing up local volumes is not permitted if the clusternode option set to yes.

Important: For the VCS, cluster database processing is skipped because VCS does not have a cluster database. VCS stores all cluster configuration information in an ASCII configuration file that is called **main.cf**, which is in the path pointed by **%VCS_HOME%conf/config** on each node in the cluster. If this file is corrupted, the cluster configuration is also corrupted. Be careful when you handle this file. The **VCS_HOME** environment variable points to the directory where VCS is installed on the node.

passwordaccess

Specify the generate value.

managedservices

(Optional). Specifies whether the Tivoli Storage Manager Client Acceptor service manages the scheduler, the web client, or both. The examples in this appendix assume that the client acceptor manages both the web client and the scheduler for each cluster group. To specify that the client acceptor

manages both the web client and the scheduler, enter the following option in the dsm.opt file for each cluster group:

```
managedservices webclient schedule
```

httpport

Specify a unique TCP/IP port number that the web client uses to communicate with the client acceptor service that is associated with the cluster group.

errorlogname

Specify a unique error log name.

Note: This file is not the same error log file that the client uses for other operations. Ideally, this file is stored on a cluster resource, but at minimum it should be stored in a location other than the client directory.

schedlogname

Specify a unique schedule log name. It is a best practice to specify a different log file name for each cluster group.

Note: This file is not the same schedule log file that the client uses for other operations. Ideally, this file is stored on a cluster resource, but at minimum it should be stored in a location other than the client directory.

Related reference:

"Clusternode" on page 345

"Domain" on page 371

"Errorlogname" on page 397

"Managedservices" on page 457

"Nodename" on page 472

"Passwordaccess" on page 480

"Schedlogname" on page 517

Step 3: Install a Client Acceptor Service and Client Agent:

Install a unique client acceptor service and client agent for each cluster group and generate a password file.

To install the Client Acceptor Service for **group-a** from workstation **node-1**, ensure that **node-1** currently owns **group-a** and issue the following command:

```
dsmcutil install cad /name:"tsm client acceptor: group-a"  
/clientdir:"c:\Program Files\tivoli\tsm\baclient" /optfile:  
q:\tsm\dsm.opt /node:mscs-cluster-group-a /password:nodepassword  
/validate:yes /autostart:yes /startnow:no httpport:1582 /cadschedname:  
"tsm scheduler service:group-a"
```

This installs the service on **node-1**.

To install the client agent service for **group-a** from workstation **node-1**, ensure that **node-1** currently owns **group-a** and issue the following command:

```
dsmcutil install remoteagent /name:"tsm client agent: group-a"  
/clientdir:"c:\Program Files\tivoli\tsm\baclient" /optfile:  
q:\tsm\dsm.opt /node:mscs-cluster-group-a /password:nodepassword  
/validate:yes /startnow:no /partnername:"tsm client acceptor: group-a"
```

This installs the remote client agent service on node1.

Note:

1. Do not use the `/autostart:yes` option.
2. Note that because the `/clusternode` and `/clustername` options are not allowed in this command at this level, it is possible that the password in the Windows Registry might need to be reset. After installing these three services for each cluster group, generate a Tivoli Storage Manager password for each cluster group node name. You need to identify the proper `dsm.opt` file for each cluster group node name you authenticate. For example: **`dsmc query session -optfile="q:\tsm\dsm.opt"`**
3. See "Frequently asked questions" on page 79 for information on what to do if a generic service resource for the cluster group fails because the CAD service has been removed.

Using the Microsoft Cluster Administrator utility or VCS Configuration Editor, move **group-a** to **node-2**. From **node-2**, issue the same commands to install the services on **node-2** and generate a password file. Repeat this procedure for each cluster group.

Step 4: Create a network name and IP address resource:

Add a network name and IP address resource for each group that is managed by the client, using the Microsoft Cluster Administrator or VCS Configuration Editor.

MSCS:

You must use the Microsoft Cluster Administrator utility to add an IP address resource to each cluster group managed by Tivoli Storage Manager.

About this task

Follow these steps to add an IP address resource:

Procedure

1. Select the **group-a** folder under the MSCS-Cluster\Groups folder and select **File > New > Resource** from the dropdown menu.
2. In the New Resource dialog, enter a unique name in the **Name** field. For example: IP address for GROUP-A. Enter a description in the **Description** field. Change resource type to IP address in the **Resource Type** field. Enter the group name in the **Group** field. Press **Enter**.
3. In the Possible Owner dialog, ensure that all cluster nodes appear as possible owners. Press **Enter**.
4. In the Dependencies dialog add all physical disk resources as Resource Dependencies. Press **Enter**.
5. In the TCP/IP Address dialog, enter appropriate values for address, subnetmask, and network. Press **Enter**.
6. Select the new resource from the Microsoft Cluster Administrator utility, and from the dropdown menu click **File** and then **Bring Online**.

Results

You must use the Microsoft Cluster Administrator utility to add a network name to each cluster group managed by Tivoli Storage Manager.

Follow these steps to add a network name:

1. Select the group-a folder under the MSCS-Cluster\Groups folder and select **File > New > Resource** from the dropdown menu.
2. In the New Resource dialog, enter a unique name in the **Name** field. For example: Network Name for GROUP-A. Enter a description in the **Description** field. Change resource type to Network Name in the **Resource Type** field. Enter the group name in the **Group** field. Press **Enter**.
3. In the Possible Owner dialog, ensure that all cluster nodes appear as possible owners. Press **Enter**.
4. In the Dependencies dialog add the IP address resource and all physical disk resources as Resource Dependencies. Press **Enter**.
5. In the Network Name Parameters dialog, enter a network name for GROUP-A. Press **Enter**.
6. Select the new resource from the Microsoft Cluster Administrator utility, and from the dropdown menu click **File** and then **Bring Online**.

The IP address and network name to backup the disks in the cluster group are now resources in the same group.

Repeat this procedure for each cluster group managed by Tivoli Storage Manager.

VCS:

You must use the VCS Configuration Editor to add a network name and IP address resource for each group that is managed by the client.

About this task

Follow these steps to add a network name and IP address resource:

Procedure

1. Open the VCS Configuration Editor. You are prompted with the Build a new configuration or modify existing configuration window which provides the following options: **New Config** - If you select this option you are prompted for the path for the types.cf file, and **Open Existing Config** - If you select this option, the configuration window opens. Click on RESOURCE GROUP you want to modify.
2. Click on the **Edit** button and select **Add resource**. The Add Resource window opens.
3. Enter the name you want to give the resource in **Resource Name** field.
4. Select the **Resource Type** as **IP**. The attributes of the IP resource type are displayed.
5. Click the **Edit** button to modify the resource attributes.
 - a. Select the **MACAddress** attribute and enter the MAC address of adapter you want the IP to be assigned to.
 - b. Select the **SubNetMask** attribute and enter the subnetmask.
 - c. Select the **Address** attribute and enter the IP address you want to make High-Availability.
6. When you are finished, close the window. The Configuration window prompts you whether to save the configuration; click **Yes**.

Step 5: Create a generic service resource for failover:

This topic guides you through the steps to create a generic service resource for failover.

Microsoft Cluster Server (MSCS):

To add a Generic Service resource to each cluster group managed by Tivoli Storage Manager, you must use the Microsoft Cluster Administrator utility.

Procedure

1. Select the **group-a** folder under the MSCS-Cluster\Groups folder and select **File > New > Resource** from the dropdown menu.
2. In the New Resource dialog, enter a unique name in the **Name** field. For example: TSM CLIENT ACCEPTOR SERVICE for GROUP-A. Enter a description in the **Description** field. Change resource type to **Generic Service** in the **Resource Type** field. Enter the group name in the **Group** field. Press **Enter**.
3. In the Possible Owner dialog, ensure that all cluster nodes appear as possible owners. Press **Enter**.
4. In the Dependencies dialog add all physical disk resources as Resource Dependencies. Press **Enter**.
5. In the Generic Service Parameters dialog, enter the service name you specified with the **dsmcutil** command, in the **Service Name** field. Leave the **Startup Parameters** field blank. Press **Enter**.
6. In the Registry Replication dialog, add the registry key corresponding to the Tivoli Storage Manager node name and server name. The format for this key is: `HKEY_LOCAL_MACHINE\SOFTWARE\IBM\ADSM\CurrentVersion\Nodes\nodename\TSM_server_instance_name`, where *nodename* is the name of your Tivoli Storage Manager node, and *TSM_server_instance_name* is the name of the Tivoli Storage Manager server that the node connects to. For example, if the node name is **mscs-cluster-group-a** and the Tivoli Storage Manager server name is **tsmsv1**, then you should enter the following registry key in the Registry Replication dialog: `HKEY_LOCAL_MACHINE\SOFTWARE\IBM\ADSM\CurrentVersion\Nodes\mscs-cluster-group-a\tsmsv1`. This entry should match an existing key in the Windows registry.

Results

The client acceptor service is now a resource in the same group. If the group is moved (failed) to the other nodes in the cluster, the service should correctly fail over between the cluster nodes and notify both cluster nodes of automatic password changes.

Note:

1. If you manually change the password, you must stop the remote agent and the client acceptor services, regenerate the password, and restart the client acceptor service (do not restart the remote agent). You can generate the password by running this command:

```
dsmc query session -optfile="q:\tsm\dsm.opt"
```
2. See "Frequently asked questions" on page 79 for information on what to do if a generic service resource for the cluster group fails because the CAD service has been removed.

Veritas Cluster Server (VCS):

To add a Generic Service resource to each cluster group managed by the backup-archive client, you must use the VCS Configuration Editor.

Procedure

1. Open the VCS Configuration Editor. You are prompted with the Build a new configuration or modify existing configuration window which provides the following options: **New Config** - If you select this option you are prompted for the path for the types.cf file, and **Open Existing Config**- If you select this option, the configuration window opens. Click on RESOURCE GROUP you want to modify.
2. Click on the **Edit** button and select **Add resource**. The Add Resource window opens.
3. Enter the name you want to give the resource in **Resource Name** field.
4. Select the Resource Type as **GenericService**. The attributes of the **GenericService** resource type are displayed.
5. Click the **Edit** button to modify the resource attributes.
6. Select the **ServiceName** attribute and enter the name of scheduler service that you want to make High-Availability.
7. When you are finished, close the window. The Configuration window prompts you whether to save the configuration; click **Yes**.

Results

Use the VCS Configuration Editor to configure the registry replication resource, as follows:

1. Open the VCS Configuration Editor. You are prompted with the Build a new configuration or modify existing configuration window which provides the following options: **New Config** - If you select this option you are prompted for the path for the types.cf file, and **Open Existing Config** - If you select this option, the configuration window opens. Click on RESOURCE GROUP you want to modify.
2. Click on the **Edit** button and select **Add resource**. The Add Resource window opens.
3. Enter the name you want to give the resource in **Resource Name** field.
4. Select the **Resource Type** as **RegRep**. The attributes of the **RegRep** resource type are displayed.
5. Click the **Edit** button to modify the resource attributes.
6. Select the **MountResName** attribute and enter the shared disk on which you want to store the registry keys.
7. When you are finished, close the window. The Configuration window prompts you whether to save the configuration; click **Yes**.

The client acceptor service is now a resource in the same group. If the group is moved (failed) to the other nodes in the cluster, the service should correctly fail over between the cluster nodes and notify both cluster nodes of automatic password changes.

Note:

1. If you manually change the password, you must stop the remote agent and the client acceptor services, regenerate the password, and restart the client acceptor

- service (do not restart the remote agent). You can generate the password by running this command: **dsmc query session -optfile="q:\tsm\dsm.opt"**
2. See "Frequently asked questions" for information on what to do if a generic service resource for the cluster group fails because the CAD service has been removed.

Step 6: Start the web client:

This topic guides you through the steps to start the web client to use cluster services.

Procedure

1. Start the Client Acceptor Service for each resource group on each node.
2. To start the web client, point your browser at the IP address and httpport specified for the Resource Group. For example, if you used an IP address of 9.110.158.205 and specified an httpport value of 1583, open the web address: <http://9.110.158.205:1583>.

Results

Alternatively, you can point your browser at the network name and httpport. For example, if you used a network name of **cluster1groupa** and specified an http port value of 1583, open the web address: <http://cluster1groupa:1583>.

Note that the web client connects to whichever workstation currently owns the resource group. The web client displays all of the local file spaces on that workstation, but to ensure that the files are backed up with the correct node name you should only back up the files for the resource group.

When failing back to the original node after a failover scenario, ensure that the remote agent service on the original workstation is stopped. The remote agent can be stopped manually, or it stops automatically after 20 to 25 minutes of inactivity. Because the remote agent is configured for manual startup, it will not start automatically if the workstation on which it was running is rebooted.

Frequently asked questions

This section contains some frequently asked questions and answers about using cluster services.

About this task

Q: How do you configure a shortcut for the backup-archive client GUI in a cluster environment?

A: To configure a backup-archive client GUI icon (for example on the Windows desktop) that you can use to manage operations for a cluster resource group on a Windows cluster, perform the following steps:

Procedure

1. Right-click on the desktop and select **New > Shortcut**.
2. In the window that appears, find the path to the dsm.exe executable (located by default in directory C:\program files\tivoli\tsm\baclient\). If you type the path in, instead of using the **Browse** button, the path should be enclosed in double quotation marks. For example: "C:\Program Files\tivoli\tsm\baclient\dsm.exe"

3. After you enter the path and executable in the text field, add the following information after the closing double quotation marks (add a space between the double quotation marks and the following): `-optfile="x:\path\to\cluster\dsm.opt"`. This identifies the proper Tivoli Storage Manager cluster options file you want to use. This example assumes that the cluster options file is located in the folder `"x:\path\to\cluster\"` and has the file name `dsm.opt`.
4. The complete line in the text field now should look similar to the following:
`"C:\Program Files\tivoli\tsm\baclient\dsm.exe" -optfile="x:\path\to\cluster\dsm.opt"`.
5. Click **Next** and give this shortcut a meaningful name, such as **Backup-Archive GUI: Cluster Group X**.
6. Click **Finish**. A desktop icon should now be available. The properties of this icon shows the following correct Target, as noted in step 4: `"C:\Program Files\tivoli\tsm\baclient\dsm.exe" -optfile="x:\path\to\cluster\dsm.opt"`.

Results

Q: How do you verify that a scheduler service setup in a cluster environment works?

A: Setting up a Tivoli Storage Manager scheduler service for a Microsoft clustered resource group can be time consuming, and can be lengthened by mistakes and errors in the syntax of the commands used to set them up. Carefully entering the commands and recording important information about the cluster setup can minimize setup time. To successfully set up a scheduler service for Microsoft cluster environments:

1. Carefully read the information in this appendix for correct syntax on setting up a scheduler service for a cluster group.
2. Ensure that the proper `dsm.opt` file(s) are used for the cluster. In a typical normal workstation, only one `dsm.opt` file is used. In a clustered environment, additional `dsm.opt` files are required. Each cluster group that is backed up must have its own `dsm.opt` file. A cluster group is any group listed under the **GROUPS** folder of the cluster tree within the Microsoft Cluster Administrator utility or VCS Configuration Editor.
3. Understand what the following `dsmcutil.exe` options mean, and when to use them. (1) `/clustername:clustername` - Specifies the name of the Microsoft cluster, where *clustername* is the name at the top level of the tree within the Microsoft Cluster Administrator utility or VCS Configuration Editor. Use this option with `dsmcutil.exe`, only when installing a scheduler service for a cluster group. Do not specify a clustername of more than 64 characters. If you specify more than 256 characters and you are using Veritas Storage Foundation with High Availability or a Microsoft Cluster Server configuration, you might not be able to install or start the Tivoli Storage Manager scheduler service, and (2) `/clusternode:yes` - Specifies that you want to enable support for cluster resources. Use this option in the `dsm.opt` file for each cluster group, and with `dsmcutil.exe` when installing a scheduler service for a cluster group.
4. Common mistakes are made in typing the syntax of the `dsmcutil.exe` command. An easy way to prevent such syntax problems is to create a temporary text file which is accessible to the cluster group (for instance, place it on a cluster drive belonging to that cluster group), and type the syntax into this file. When needed, cut and paste this syntax from the file to the DOS prompt and press the **Enter** key. It guarantees the consistency of the command syntax regardless of which computer you enter it on.

5. If the scheduler service is failing to restart after a failover of the cluster group occurs (using the MOVE GROUP option in Cluster Administrator, for example), there might be potential password synchronization problems between the two cluster workstations. To verify that the passwords are the same, browse to this registry key on each workstation and compare the encrypted password value: `HKEY_LOCAL_MACHINE\SOFTWARE\IBM\ADSM\CurrentVersion\Nodes\nodename\servername`.

If the encrypted keys for this node do not match between the two cluster workstations, there is a password mismatch on one or both of the two workstations. To correct this problem, use the `dsmc.exe` program to update the password manually on both workstations.

For example, assume that the Y: drive is part of the cluster group that is experiencing problems being backed up with a scheduler service. The Y:\tsm directory contains the `dsm.opt` file for this cluster group in the Y:\tsm directory. To update the password manually, enter the following command on both workstations: `dsmc -optfile=Y:\tsm\dsm.opt -clusternode=yes`, and enter the following command to receive the prompt for the node name and password: **`dsmc q se -optfile=Y:\tsm\dsm.opt -clusternode=yes`**.

Verify that the passwords are synchronized, and restart the scheduler service to verify if the password remains consistent. If password mismatching continues, it might be due to a syntax error in the original `dsmcutil.exe` command that was used to install the scheduler service. In this case, uninstall the scheduler service (using the `dsmcutil remove /name:schedule_name` command) and reinstall the scheduler service again (using the shared text file syntax as shown previously).

Q: How do you add a cluster drive to an existing cluster scheduler service resource for backup?

A: To add an additional cluster drive resource to an existing backup-archive client cluster scheduler service, the following components must be modified or updated to properly reflect this change:

1. The cluster drive resource, and any related resource shares, must exist and reside within the designated cluster group as defined in the Microsoft Cluster Administrator utility or VCS Configuration Editor. The designated cluster group must already contain the cluster scheduler service resource for which this new drive is added.
2. The `dsm.opt` file used by the designated cluster scheduler service resource must be modified to include the additional cluster drive resource on the `domain` option statement. For example, if you want to add the R:\ drive, and the `domain` statement currently identifies cluster drives Q: and S:, update the `domain` statement in your `dsm.opt` file as follows: `domain Q: S: R:.`
3. You must modify the cluster scheduler service resource properties to include this file in the list of dependent resources needed to bring this resource online. This ensures that the cluster drive resource being added is included in the new backups, and for backups which run after a failover occurs.

After the changes above are made, bring the cluster scheduler service resource offline, and back online. The schedule should now process this additional resource for backups.

Q: The client acceptor service has been removed and now the generic service resource for the cluster group is failing. How can this be corrected?

A: The client acceptor can be used to control the scheduler, the web client, or both for a cluster environment. If the client acceptor is removed without updating the generic cluster resource, the resource fails. To correct this:

1. Verify which scheduler service was controlled by the client acceptor.
2. Using the Microsoft Cluster Administrator utility or VCS Configuration Editor, go to the properties window of the service resource, select the Parameters tab, and enter the name of the correct scheduler service to use.
3. Repeat steps one and two for each cluster group that was managed by the specific client acceptor.
4. To test the updated service resource, initiate a failure of the resource. If the resource comes back online with no failures, the update has worked properly.

Note: To fully disable the client acceptor service, remove the `managedservices` option from the cluster group `dsm.opt` file or comment it out.

Configuring online-image backup support

If the online image feature is configured, the backup-archive client performs a snapshot-based image backup, during which the real volume is available to other system applications.

About this task

A consistent image of the volume is maintained during the online image backup.

To configure online image backup, perform the following steps:

Procedure

1. Select **Utilities > Setup Wizard** from the backup-archive client GUI main window. The Client Configuration Wizard panel appears.
2. Select **Help me configure Online Image Support** and click **Next**. The TSM Online Image Support Wizard panel appears.
3. Select the **VSS Snapshot Provider** and then click **Next**. To disable online image support, select **None**.
4. Click **Finish** button to complete the setup.
5. Complete each panel in the wizard and click the **Next** to continue. To return to a previous panel, click the **Back**. To display help information for a panel, click the help icon.

Results

To set preferences for open file support, use the Include-Exclude tab on the Tivoli Storage Manager Preferences editor. You can set these options for all volumes or for individual volumes using the `include.fs` option: `snapshotproviderfs`, `presnapshotcmd`, `postsnapshotcmd`.

Related concepts:

"Client options reference" on page 321

"Image backup" on page 153

Configuring Open File Support

You configure Open File Support (OFS) after you install the Window client.

About this task

If the Open File Support feature is configured, the backup-archive client performs a snapshot-based, file-level operation, during which the real volume is available to other system applications. A consistent image of the volume is maintained during the operation.

To configure OFS, perform the following steps:

Procedure

1. Start the Windows client Java GUI (run `dsm.exe`).
2. Select **Utilities > Setup Wizard**.
3. Select **Help me configure Online Image Support** and click **Next**.
4. Click **Next** again.
5. Select the **VSS** snapshot provider to enable Open File Support or select **None** to perform normal (non-snapshot) backups of the files on your volume; then click **Next**.
6. Click **Apply** and then click **Finish**.

Results

To set preferences for open file support, use the Include-Exclude tab on the Preferences editor. You can set these options for all volumes or for individual volumes using the `include.fs` option: `snapshotproviderfs`, `presnapshotcmd`, `postsnapshotcmd`

Related concepts:

“Client options reference” on page 321

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:
Allowed Capabilities: login-http-admin,api-*

```

- c. If the **security.passwd.firstlogin.enable** option for the user ID on the NetApp server is set to on, ensure that all groups have the **login-telnet** and **cli-passwd*** capabilities.

Tip: When **security.passwd.firstlogin.enable** option is enabled, the user ID is set to expired when created. The user cannot run any commands, including snapshot 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 Windows.

- Map the NetApp volumes by using CIFS.
 - Ensure the NetApp volumes have the NTFS 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 83 using the following steps:
 - a. Log on as the user with read/write access to the CIFS share.
 - 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 83.

newPassword

This value is the password for the user ID that you created in step 1 on page 83.

Related reference:

“Snapdiff” on page 529

“Snapdiffhttps” on page 534

“Createnewbase” on page 354

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 83.
- 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.

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).

Procedure

Complete the following steps on the remote machine where the backup-archive client is installed:

1. Set the cluster management filer credentials with 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.

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 management_filer_hostname  
storage_virtual_machine_hostname storage_virtual_machine_name
```

Where:

management_filer_hostname

The host name or IP address of the cluster management filer.

storage_virtual_machine_hostname

The host name or IP address of the storage virtual machine that is used to mount volumes to back up.

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. Map the volumes to drive letters. For example, enter the following command for each storage virtual machine:

```
net use y: \\storage_virtual_machine_hostname domain_name\CIFS_share_name
```

Where:

y: The drive to map the volume to.

storage_virtual_machine_hostname

The host name or IP address of the storage virtual machine.

domain_name\CIFS_share_name

The CIFS share that is defined on the filer on the volume being backed up.

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 Windows clients, enter the following command:

```
dsmc incr y: -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 Windows clients, enter the following command:

```
dsmc incr y: -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: netappl-v1.example.com
Storage virtual machine name: netappl-client1
CIFS share: demovol
Volume name: demovol
```

Infinite volume storage virtual machine

```
Hostname: netappl-v4.example.com
Storage virtual machine name: netappl-infiniteVolume1
CIFS Share: InfiniteVol
```

The user completes the following steps on the backup-archive client:

1. Set the management filer credentials with 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 netapplmgmt.example.com netappl-v1.example.com
netappl-client1
```

```
dsmc set netappsvm netapplmgmt.example.com netappl-v4.example.com
netappl-infiniteVolume1
```

3. Map remote volumes to drive letters for each storage virtual machine:

```
net use y: \\netappl-v1.example.com\demovol WINDC\domainuser
net use z: \\netappl-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 CIFS 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 CIFS 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 CIFS 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 `createnewbase=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.

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

Closed registration

With closed registration, a Tivoli Storage Manager administrator must register your workstation as a client node with the server. If your enterprise uses closed registration, you must provide some information to your Tivoli Storage Manager administrator.

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

The client options file, dsm.opt, must be in a non-Unicode format. However, if you are using a separate include-exclude file, it can be in Unicode or non-Unicode format.

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.opt, 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 client options file
 3. **Important:** Group your include-exclude options together in your client options file.
 4. Enter your include and exclude statements. The 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. All other include-exclude statements are processed from the bottom of the list up. Therefore, it is important to

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

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

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

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

5. Save the file and close it.
6. Restart the client and the scheduler and client acceptor services to enable your include-exclude list.

Related concepts:

“System files to exclude” on page 96

Chapter 9, “Storage management policies,” on page 267

Related reference:

“Incl excl” on page 429

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 7 on page 94 lists the options you can use to exclude file spaces and directories from processing.

Table 7. Options for excluding file spaces and directories

Option	Description
exclude.dir "Exclude options" on page 400	<p>Excludes a directory, its files, and all its subdirectories and their files from backup processing. For example, the statement <code>exclude.dir c:\test\dan\data1</code> excludes the <code>c:\test\dan\data1</code> directory, its files, and all its subdirectories and their files. Using the <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 <code>include</code> options to override an <code>exclude.dir</code> statement. Only use <code>exclude.dir</code> when excluding an entire directory branch.</p> <p>If you define an <code>exclude</code> statement without using a drive letter, such as <code>exclude.dir dirname</code>, this excludes from processing any directory named <code>dirname</code> on any drive.</p> <ul style="list-style-type: none"> The following examples illustrate valid <code>exclude.dir</code> statements: <p>Exclude directory <code>C:\MyPrograms\Traverse</code> and its files and subdirectories:</p> <pre>exclude.dir c:\MyPrograms\Traverse</pre> <p>Exclude all directories below <code>c:\MyPrograms\Traverse</code>. Note that directory <code>C:\MyPrograms\Traverse</code> and the files immediately below <code>C:\MyPrograms\Traverse</code> is eligible for backup.</p> <pre>exclude.dir c:\MyPrograms\Traverse*</pre> <p>Exclude all directories whose names begin with <code>temp</code>, and are located within directory <code>x:\documents and settings</code> and its subdirectories, where <code>x</code> is any drive.</p> <pre>exclude.dir "x:\documents and settings\...\temp"</pre> <p>Exclude all directories whose names begin with <code>temp</code>, regardless of the drive or directory in which they reside:</p> <pre>exclude.dir temp*</pre> <p>The following example is invalid because it ends with a directory delimiter:</p> <pre>exclude.dir c:\MyPrograms\Traverse\</pre> Use the following statements to exclude drive <code>x</code>: altogether from backup processing. Note that the drive root (<code>x:\</code>) is backed up, but all other files and directories on <code>x</code>: is excluded. An alternative method for excluding an entire drive from domain incremental backup is to use a domain statement to exclude the drive. For example: <pre>domain -x:</pre> <p>This alternative still permits selective and explicit incremental backup processing of files on <code>x</code>:. For example:</p> <pre>dsmc s x:\ -subdir=yes dsmc i x: dsmc i x:\MyPrograms\ -subdir=yes</pre>

Include-exclude statements for networked file systems

Include-exclude statements that involve networked file systems (remote drives) must be written in the UNC format.

In the following example Z: is a mapped drive to a remote file system on vista.example.com.

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

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

Here is an example of the new format using UNC:

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

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

Exclude files and directories from a journal-based backup

There are two methods of excluding files and directories from a journal-based backup.

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

Note: There is no correlation between the two exclude statements. 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 8 lists the options that you can use to control processing with include and exclude statements.

Table 8. Options for controlling processing using include and exclude statements

Option	Description	Page
Back up processing		
exclude	<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 400
exclude.backup		
exclude.file		
exclude.file.backup		
include	Use these options to include files or assign management classes for backup processing.	"Include options" on page 430
include.backup		
include.file		
include.fs	Use this option to set options on a file space-by-file space basis.	"Include options" on page 430
Archive processing		

Table 8. Options for controlling processing using include and exclude statements (continued)

Option	Description	Page
exclude.archive	Excludes a file or group of files from archive services.	"Exclude options" on page 400
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 430
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.dir statements. This option is for all Windows clients.	"Exclude options" on page 400
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 all Windows clients.	"Exclude options" on page 400
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 client options file (dsm.opt). See "Toc" on page 563 for more information. This option is valid for all Windows clients.	"Include options" on page 430
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 all Windows clients.	"Include options" on page 430
System state processing		
include.systemstate	Assigns management classes for backup of the Windows Server 2008 and Windows 7 system state. The default is to bind the system state object to the default management class.	"Include options" on page 430

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.

The implicitly generated statements can be seen in the lines of output of the **query inclexcl** command with the source "operating system".

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

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

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

Exclude files with UNC names

You can exclude remotely accessed files by specifying their universal naming convention (UNC) names in your exclude statement.

The following example assumes that local drive letter `g` is mapped to the remote share point:

```
\\remote\books
```

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

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

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

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

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 Windows, the names of directories and files should not contain the following symbols:

```
? * < > " / \ : |
```

However, they can contain the following symbols:

```
[ ]
```

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 C:\[dir2] from backup processing, enter the following in the dsm.opt file:

```
exclude.dir "C:\[[]dir2[\\]"
```

To exclude the single file C:\file[.txt from backup processing, enter the following in the dsm.opt file:

```
exclude.dir "C:\file[\\].txt"
```

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.opt file or the include-exclude file.

Related concepts:

"Wildcard characters" on page 632

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 9. Wildcard and other special characters

Character	Function
?	<p>The match one character matches any single character <i>except</i> the directory separator; it does not match the end of the string. For example:</p> <ul style="list-style-type: none"> The pattern ab?, matches abc, but does not match ab, abab, or abzzz. The pattern ab?rs, matches abfrs, but does not match abrs, or abllrs. The pattern ab?ef?rs, matches abdefjrs, but does not match abefrs, abdefrs, or abefjrs. The pattern ab??rs, matches abcdrs, abzzrs, but does not match abrs, abjrs, or abkkrs.
*	<p>The match-all character. For example:</p> <ul style="list-style-type: none"> The pattern ab*, matches ab, abb, abxxx, but does not match a, b, aa, bb. The pattern ab*rs, matches abrs, abtrs, abrsrs, but does not match ars, or aabrs, abrss. The pattern ab*ef*rs, matches abefrs, abefghrs, but does not match abefr, abers. The pattern abcd.*, matches abcd.c, abcd.txt, but does not match abcd, abcdc, or abcdtxt.

Table 9. Wildcard and other special characters (continued)

Character	Function
\...	<p>The match-<i>n</i> character matches zero or more directories.</p> <p>The following pattern specifies all files in the root directory of the C drive:</p> <p>c:*</p> <p>The following pattern specifies all files and all directories on the C drive:</p> <p>c:\...*</p>
[<p>The open character-class character begins the enumeration of a character class. For example:</p> <p>xxx[abc] matches xxxa, xxxb, or xxxc.</p>
-	<p>The character-class range includes characters from the first character to the last character specified. For example:</p> <p>xxx[a-z] matches xxxa, xxxb, xxxc, ... xxxz.</p> <p>This format should not be used to specify remote drives in an <i>exclude</i> statement.</p>
\	<p>The literal escape character. When used within a character class, it treats the next character literally. When used outside of a character class, it is not treated in this way. For example, if you want to include the ']' in a character class, enter [...\...]. The escape character removes the usual meaning of ']' as the close character-class character.</p>
]	<p>The close character-class character ends the enumeration of a character class.</p>
:	<p>The drive separator character separates a file specification. The character <i>before</i> the colon identifies a drive letter. The characters <i>after</i> the colon identify file specification or pattern. For example:</p> <p>d:\direct\file.nam</p>

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

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

If you are using UNC names, Table 10 shows how to correctly specify shared drives.

Table 10. Specifying a drive specification using wildcards

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

Related concepts:

“Wildcard characters” on page 632

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.

Table 11 shows how to use wildcard characters to include or exclude files.

Table 11. Using wildcard characters with include and exclude patterns

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

Related concepts:

"Examples using wildcards with include and exclude patterns"

Related reference:

"Exclude options" on page 400

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 12. Options for controlling compression and encryption processing

Option	Description	Page
Compression processing		

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

Option	Description	Page
exclude.compression	Excludes files from compression processing if compression= <i>yes</i> is specified. This option applies to backups and archives.	"Exclude options" on page 400
include.compression	Includes files for compression processing if compression= <i>yes</i> is specified. This option applies to backups and archives.	"Include options" on page 430
Encryption processing		
exclude.encrypt	Excludes files from encryption processing.	"Exclude options" on page 400
include.encrypt	Includes files for encryption processing. The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received. Important: The include.encrypt option is the only way to enable encryption on the Backup-Archive client. If no include.encrypt statements are used encryption will not occur.	"Include options" on page 430
Adaptive subfile backup processing		
exclude.subfile	Excludes files from adaptive subfile backup processing. This option does not apply to archive processing. This option is valid for all Windows clients.	"Exclude options" on page 400
include.subfile	Includes files for adaptive subfile backup processing. This option does not apply to archive processing. This option is valid for all Windows clients.	"Include options" on page 430

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 693

"Preview Backup" on page 694

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 client options file include-exclude list contains one or more `incl excl` options that specify include-exclude files, the include-exclude statements in these files are placed in the list position occupied by the `incl excl` option and processed accordingly.

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.dir` statements first, and removes the excluded directories and files from the list of objects available for processing.

After evaluating all `exclude.dir` statements, the client evaluates the include-exclude list from the bottom up and stops when it finds an include or exclude statement that matches the file it is processing. The order in which the include and exclude options are entered therefore affects which files are included and excluded.

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

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

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

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

Example 2

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

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

The file being processed is: c:\widg\copyit.bat. Processing follows these steps:

1. Rule 3 is checked and finds no match.
2. Rule 2 is checked and finds no match.
3. Rule 1 is checked and finds no match.
4. Because a match is not found, file c:\widg\copyit.bat is implicitly included and backed up.

Example 3

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

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

The current file being processed is: c:\lib\objs\printf.obj. Processing follows these steps:

1. Rule 3 is checked and finds no match.
2. Rule 2 is checked and finds no match.
3. Rule 1 is checked and a match is found.
4. Processing stops, the option is checked, and it is excluded.
5. File c:\lib\objs\printf.obj is not backed up.

Related concepts:

“Exclude file spaces and directories” on page 93

Chapter 11, "Processing options," on page 295

Related reference:

"Exclude options" on page 400

"Query Inclexcl" on page 712

Processing rules when using UNC names

When processing files with UNC names, there are rules that must be followed.

The backup-archive client uses the rules that are described in "Include and exclude option processing" on page 102. The rules in "Explicit use of UNC names for remote drives" also apply.

Explicit use of UNC names for remote drives

The backup-archive client recognizes explicit use of UNC names for remote drives.

For example, as shown in Table 13, the UNC name pattern can be substituted for the DOS pattern.

Assume local drive letter r: is mapped to remote share point \\remote\\c\$, s: is mapped to \\remote\\share4, and t: is mapped to \\remote\\share2.

Table 13. UNC name patterns and DOS patterns

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

Conversion of DOS pathnames for fixed and remote drives

The backup-archive client converts DOS path names that are mapped to remote share points.

For example, a remote share point that is mapped from r:\\test\\...\\exclude.out to \\remote\\share\\test\\...\\exclude.out is converted. Remote share points that are not mapped are not converted. Files on removable media are not converted.

Character-class matching examples

This topic shows examples of valid matches using character class.

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

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 110
- “Starting a command-line session” on page 111
- “Starting a web client session” on page 115
- “Start the client scheduler automatically” on page 116
- “Changing your password” on page 116

You can also complete the following tasks:

- “Sorting file lists using the backup-archive client GUI” on page 117
- “Displaying online help” on page 119
- “Ending a session” on page 119

Backup-archive client operations and security rights

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

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

Table 14 summarizes the user security rights needed for backup and restore operations. The information in the table assumes that the default privileges for the Microsoft Windows Administrators group, Backup Operators group, and Users group have not been altered.

Table 14. Required user security rights for Tivoli Storage Manager backup and restore services

Operating system	Account	What can I back up and restore?
Windows Clients	Member of Administrators group	<ul style="list-style-type: none">• Back up and restore all file and directory objects• Back up and restore system state• System state data (Backup Operators group cannot back up ASR writer data and cannot restore system state data)

Table 14. Required user security rights for Tivoli Storage Manager backup and restore services (continued)

Operating system	Account	What can I back up and restore?
Windows Clients	Member of Backup Operators group	<ul style="list-style-type: none"> • Back up and restore all file and directory objects • Back up system state, except for ASR Writer <p>Note: Backup Operator group members cannot restore system state.</p>
Windows Clients	Member of Users or other group	<ul style="list-style-type: none"> • Back up and restore all file and directory objects <p>Attention: Users must have the following Microsoft Windows security privileges in order to back up and restore files and directories:</p> <ul style="list-style-type: none"> – Back up files and directories – Restore files and directories <p>These privileges represent a potential security risk since they allow the user to back up any file, or restore any file for which a backup copy exists. The privileges should be granted only to trusted users. For more information about these privileges, see the Microsoft Windows documentation.</p> <p>Note: System state cannot be backed up or restored.</p>

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

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

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

Backup Operators group operations

The Backup Operators group allows users to back up and restore files regardless of whether they have read or write access to the files.

This group has a limited set of user rights, so some functions are not available to members of the Backup Operators group.

The following list contains the backup-archive client operations that a member of the Backup Operators can do:

- Back up and restore files (see Table 14 on page 105)
- Back up system state
You must be a member of the Administrators group to back up ASR writer data.
- Restore system state
You must be a member of the Administrators group to restore system state.
- Start the scheduler service

The following list contains the backup-archive client operations that a member of the Backup Operators cannot do:

- Start any other services (client acceptor, remote client agent, and journal service)
- Install and configure client services
- Use open file support (OFS)
- Back up and restore images
- Back up and restore Windows file shares

Considerations before you start using a Backup Operators group account

There are some items that you need to consider before you use a Backup Operators group account to back up, archive, restore, or retrieve your data.

Consider these items before using a Backup Operators group account to back up, archive, restore, or retrieve your data:

- If you have already been using the backup-archive client with an Administrators group account you might not be able to launch the client because you cannot open the log files (for example `dsmerror.log`). To alleviate this problem, you can grant the Backup Operators group Read and Write permissions to the log files or the directories containing these log files.
- If you have been creating system state backups with an account that is a member of the Administrators group, and if you want to start backing up or restoring data with an account that is a member of the Backup Operators group, you must delete the staging directory, `c:\adsm.sys`, before backing up or restoring the system state as a member of the Backup Operators group.
- If you have existing backups from a Tivoli Storage Manager Version 5.2 or earlier client and you attempt an incremental backup of an existing file space with a member of the Backup Operators group, all of the data appears as changed and it is resent to the Tivoli Storage Manager Server.

- Members of the Backup Operators group might not be able to back up or restore file data that was encrypted by an Administrator account using the Windows encrypting file system (EFS).
- Members of the Backup Operators group do not have the proper authority to update the last access time for files that is encrypted with the Windows encrypting file system (EFS). If EFS files are restored by a member of the Backup Operators group, the last access time will not be preserved.

Permissions required to restore files that use adaptive subfile backup

To restore files that were processed using adaptive subfile backup, you must either be the owner of the file or have read access.

These permissions are in addition to those required to perform a normal restore.

Related tasks:

“Performing a backup with limited bandwidth” on page 133

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

To back up, restore, archive, or retrieve data residing on Microsoft Cluster Server (MSCS) or Veritas Cluster Server cluster resources, your Windows account must belong to the Administrators or Domain Administrators group or Backup Operators group.

By default, Backup Operators do not have the user rights necessary to perform these tasks on a cluster node. However, Backup Operators can perform this procedure if that group is added to the security descriptor for the Cluster service. You can do that using Cluster Administrator or cluster.exe.

Tivoli Storage Manager client authentication

When using the 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 512

 `QUERY ADMIN` command

User account control

User Account Control (UAC) is a Windows security feature that helps prevent malware from compromising the operating system. UAC restricts programs to standard user privileges.

When UAC is enabled, programs that require elevated privileges cannot run without your permission.

The backup-archive client requires elevated privileges. If UAC is enabled when you run the client, a User Account Control dialog box is displayed. The dialog asks if you want to allow the program to run. If you are not logged in as an administrator, the dialog also asks for your account credentials.

Enabling client access to network shares when UAC is enabled

When Windows User Account Control (UAC) is enabled, the backup-archive client cannot access existing network share mappings. The solution is to map the network shares from an elevated command prompt before you start the client.

About this task

When you map a network share, the share is linked to your current Windows login access token. That token has only standard user privileges. Because the backup-archive client must run with elevated privileges, a different access token is

used. Since the network share is not linked to this other access token, the mapped network share is not visible to the client. The network share must be linked to the access token that has the elevated privileges to make the share visible to the client.

Procedure

Complete the following steps to enable the client to access data on network shares.

1. Create a desktop shortcut for the Windows command prompt. The default location of the command prompt executable file is `C:\Windows\System32\cmd.exe`.
2. Right-click the shortcut and select **Run as Administrator**. A UAC prompt is displayed with instructions that describe how to proceed.
 - If you are logged in as a member of the Administrators group, click **Yes** to allow the client to run with elevated privileges.
 - If you are not logged in as a member of the Administrators group, enter your credentials when prompted to do so, and then click **Yes** to allow the client to run with elevated privileges.

Perform the remaining steps in the elevated command prompt window that you just opened.

3. Use the Windows **net use** command to map the network shares. Contact your system administrator if you need help with the **net use** command.

Note: Do not use Windows Explorer to map the network share because Windows Explorer runs with the standard user rights token.

4. Change to the directory where the client is installed. The default installation directory is `C:\Program Files\Tivoli\TSM\baclient`.
5. Start the client GUI (`dsm.exe`) or the command-line client (`dsmc.exe`) and backup or restore data that is on network shares.

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
Windows	<p>To start the backup-archive client GUI on a Windows system, use one of the following methods:</p> <ul style="list-style-type: none"> • Click Start > Programs > Tivoli Storage Manager > Backup-Archive GUI. • Click Start > Run and enter the full path to the backup client <code>dsm.exe</code> file. • On the command line, change directory to the backup-archive client installation directory and enter dsm. <p>On Windows operating systems that have the User Account Control feature enabled, you might be prompted to allow the <code>dsm.exe</code> program to run. To allow the program to continue and start the backup-archive client GUI, provide administrative credentials.</p>

The backup-archive client locates and uses the options that are specified in the client options file (`dsm.opt`).

Related concepts:

Chapter 2, “Configure the Tivoli Storage Manager client,” on page 27

Related tasks:

“Configuring the language for displaying the Java GUI” on page 33

“Configuring the language for displaying the Java GUI” on page 33

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 116

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.

You can launch the setup wizard at any time to modify your client options file.

The client options file is `dsm.opt`.

Starting a command-line session

You can start a command-line session by invoking the **dsmc** command.

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

You can open the Windows **Start** menu and select **Programs > Tivoli Storage Manager > Backup-Archive Command Line**.

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:

“Backup-archive client operations and security rights” on page 105

“Options in interactive mode” on page 630

“Start and end a client command session” on page 628

Chapter 12, “Using commands,” on page 625

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 c:\myfiles\file1.dat
```

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 480

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>
```

On Windows 7, when you log on with an administrator ID, you can complete standard user tasks.. If you are not logged on before you begin a task from a command-prompt window, you are prompted to do so..

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 c:\myfiles\file1.dat
```

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.

Displaying Euro characters in a command-line prompt

This topic explains how to display the Euro character in the Windows command-line prompt (console window).

Procedure

1. Contact your Microsoft Representative for the 858 code page (the file name is `c_858.nls`). Copy the file into your Windows system32 directory (for example, `C:\WINNT\system32`).
2. Edit the Windows Registry key, using this command: `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Nls\CodePage\850`, and set it to value `c_858.nls`. **Any changes that you make to the Windows Registry editor cannot be undone.** Errors made in editing the Windows Registry can cause your system to malfunction, and you might not even be able to restart your system. **Be very careful** when editing the Windows Registry. If you are unfamiliar with using the Windows Registry editor, then ask someone else who is familiar with the Windows Registry editor to help you.
3. In your Regional Settings, select a Western European country (Germany, France, Italy, etc.) as your locale setting.
4. Exit and reboot the system.

Results

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

Use options on the DSMC command

This topic shows some examples of how to use options on the **dsmc** command.

About this task

For example, suppose you have one workstation with node name `galaxy1`, and another workstation with node name `galaxy2`, and you want to restore the data from `galaxy1` to the `galaxy2` system. To recover a file from one workstation (`galaxy1`) while at the other workstation (`galaxy2`), you must access `galaxy1`. Use the **set access** command to gain access.

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

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

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

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

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

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

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

Related reference:

“Fromnode” on page 420

“Set Access” on page 775

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 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.

For all supported Microsoft Internet Explorer browsers, the browser option to allow scripting of Java applets must be enabled. By default, this option is enabled in Internet Explorer.

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.

Note: Back up and restore activities that are running with a NAS server continue after the web client disconnects.

Related concepts:

“Web client configuration overview” on page 33

User privileges

If you plan to use the web client, ensure that you were assigned an administrative user ID with system privilege, policy privilege, client access authority, or client owner authority.

When a new node is registered with the server, by default it is given an admin ID of the same node name with client owner authority.

Note: You can use the `revokeremoteaccess` option to prevent a Tivoli Storage Manager administrator with client access privilege from performing client operations on your workstation through the web client. However, Tivoli Storage Manager administrators with client owner privilege, system privilege, or policy privilege can still perform client operations on your workstation through the web client.

Related concepts:

“Tivoli Storage Manager client authentication” on page 108

Related reference:

“Revokeremoteaccess” on page 512

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.

Tivoli Storage Manager supports remote network connections to the server. With a remote network connection, mobile users no longer need to dial-in to their company network when a backup is scheduled to run. Tivoli Storage Manager automatically establishes a connection before the scheduled backup occurs. If the connection fails, Tivoli Storage Manager reestablishes the connection before attempting the backup.

Related tasks:

“Setting the client scheduler process to run as a background task and start automatically at startup” on page 252

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. From the main window, open the **Utilities** menu and select **Change password**.
2. Enter your current and new passwords, and enter your new password again in the **Verify password** field.
3. Click **Change**.

Results

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

For UNIX, Linux, and Windows clients:

```
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 Windows systems:

Enclose the command parameters in quotation marks (").

Command line example:

```
dsmc set password "t67@#$$%^&" "pass2<w0rd"
```

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 116

Related reference:

"Password" on page 478

"Set Password" on page 781

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

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 428

"Pick" on page 481

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.
- From the **dsmc** command line: Enter the **help** command. The complete table of contents for the available help text is displayed.

Related reference:

“**Help**” on page 681

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 main window:
 - Click **File > Exit**.
 - Press Alt-X.
 - 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 Windows Task Manager to end the **dsmc** process. Do not press Ctrl-C because, while it ends the session, it can lead to unexpected results.

Related reference:

“**Loop**” on page 690

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 web client does not offer a Setup wizard, which is available in the backup-archive client GUI on Windows clients. The web client cannot browse network resources.

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

The client provides backup and archive services for all files on the following file systems: File Allocation Table (FAT), FAT 32, NTFS, and ReFS.

The following is a list of primary backup tasks.

- “Planning your backups (Windows)”
- “Pre-backup considerations (Windows)” on page 130
- “Incremental, selective, or incremental-by-date backups (Windows)” on page 136
- “Deleting backup data” on page 128
- “Backing up files from one or more file spaces for a group backup (Windows)” on page 144
- “Backing up Windows system state” on page 149
- “Backing up Automated System Recovery files” on page 150
- “Image backup” on page 153
- “Back up NAS file systems using Network Data Management Protocol” on page 159
- “Preparing the environment for full backups of VMware virtual machines” on page 167
- “Back up virtual machines on a Hyper-V system” on page 178
- “Backing up Net Appliance CIFS share definitions” on page 182

Planning your backups (Windows)

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 performing a backup.

Read the tasks listed in this table to determine whether you are ready to back up your data.

Table 16. Planning your backups

<input type="checkbox"/>	Decide whether you want to back up or archive files. See “When to back up and when to archive files” on page 129 for more information.
--------------------------	--

Table 16. Planning your backups (continued)

<input type="checkbox"/>	See “Pre-backup considerations (Windows)” on page 130 for important migration information, and how you might increase performance before backing up files and directories.
<input type="checkbox"/>	Create an include-exclude list to specify files and directories you want to exclude from backup services. See “Control processing with an include-exclude list” on page 134 for more information.
<input type="checkbox"/>	Decide what type of backup you want according to your needs. See the following sections for more information: <ul style="list-style-type: none"> • “Incremental, selective, or incremental-by-date backups (Windows)” on page 136 • “Backing up files from one or more file spaces for a group backup (Windows)” on page 144 • “Backing up Windows system state” on page 149 • “Backing up Automated System Recovery files” on page 150 • “Image backup” on page 153 • “Back up NAS file systems using Network Data Management Protocol” on page 159 • “Parallel backups of virtual machines” on page 174
<input type="checkbox"/>	For additional backup considerations, see “Backup (Windows): Additional considerations” on page 185.

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 267

“Management classes and copy groups” on page 268

Related tasks:

“Backing up Windows system state” on page 149

Related reference:

“Absolute” on page 322

Open file support for backup operations

The VSS snapshot provider is used for open file support.

VSS is the snapshot provider for Windows.

Some applications can create files and open these files in a way that denies access to all other processes on a Microsoft Windows operating system. Although this is not a common practice, it is sometimes used by database vendors or other applications that might want to limit access to certain files. By restricting access to these files, backup products are prevented from backing up the data. These locked files are not the same as files that are open, or in use. The backup-archive client, running without the open file support (OFS) feature, can back up open, or in use files, including files that are open for reading or writing, files that are changing during the backup, executable and dll files that are running, log files that are being appended to, and so on.

You can create OFS or online image backups on workstations with a single NTFS-based, or ReFS-based, C:\ drive.

The following is the error message that is seen in the dsmerror.log when the client encounters one of these locked files without OFS support enabled:

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

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

Do not use OFS for backing up locked Windows system files, such as the Windows system state. The client has advanced features for backing up data that is contained within these files. The backup of the system data that is contained in these files requires extra processing and must be backed up in a group to allow for a successful restore. These files are excluded from Tivoli Storage Manager file level backup.

For database applications that use certain files for transactional consistency (for example, a recovery log file), it might not be possible to back up and restore these files without database coordination. In these situations, do not back up these database files with the normal file level backup. You can exclude these files from backup processing by using an exclude or exclude.dir statement. A number of data protection clients (IBM Tivoli Storage Manager for Databases, IBM Tivoli Storage Manager for Mail, and so on) are available to provide this database coordination and backup along with other advanced database backup features. For a current list of data protection clients go to this website: <http://www.ibm.com/software/tivoli/products/storage-mgr/product-links.html>

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

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

If the time it takes to complete the file level backup is too long to have the open files offline (for example, having the database offline or holding up transactions), use the OFS feature to create a point-in-time snapshot of the volume. In this case, use the `presnapshotcmd` and `postsnapshotcmd` options to signal the database or application to coordinate with the backup of these open files. The snapshot, which occurs between the pre-snapshot command and post-snapshot command, generally takes only a few seconds to create. This allows the database or application to resume operations quickly while still allowing the client to perform a full incremental backup of the volume, including the locked files. There are other situations where these application-locked files can be safely backed up and restored on a file-by-file basis. In these situations, you can enable the OFS feature for that volume where the open files exist. The client then has access to these files and back them up using file level backup and archive operations.

For information about Open File Support restrictions and issues, see technote 1248971.

If open file support has been configured, the client performs a snapshot backup or archive of files that are locked (or "in use") by other applications. The snapshot allows the backup to be taken from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to the file system are not included in the backup. You can set the `snapshotproviderfs` parameter of the `include.fs` option to **none** to specify which drives do not use open file support.

To control an open file support operation, you can specify these additional options in your `dsm.opt` file or as values of the `include.fs` option: `snapshotproviderfs`, and `presnapshotcmd` and `postsnapshotcmd`.

Note:

1. You can use the `include.fs` option to set snapshot options on a per file system basis.
2. Open file support is provided for both backup and archive. For backup, this includes incremental, incremental by date, selective, incremental image, adaptive subfile backup, and journal-based backup.
3. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with FAT, FAT32, NTFS, or ReFS file systems. This support includes SAN-attached volumes that meet these requirements.
4. To enable OFS support in a cluster environment, all workstations in the cluster must have OFS configured. Set VSS as the snapshot provider on the `snapshotproviderfs` option.

Related concepts:

Chapter 11, "Processing options," on page 295

Related tasks:

“Backing up Windows system state” on page 149

“Configuring Open File Support” on page 83

Backing up data using the GUI

You can use the backup-archive client GUI to back up specific files, a group of files with similar names, or entire directories.

About this task

You can locate the files that you want to back up by searching or filtering. Filtering displays only the files that match the filter criteria for your backup. Files that do not match the filter criteria do not display.

To perform a GUI backup, use the following steps:

Procedure

1. Click **Backup** from the GUI main window. The **Backup** window appears.
2. Expand the directory tree by clicking the plus sign **+**. To display files in a folder, click the **Folder** icon. To search or filter files, click the **Search** icon from the toolbar.
3. Click the selection box for the objects that you want to back up.
4. Select the type of backup from the pull-down menu:
 - a. To run an incremental backup, select **Incremental (complete)**.
 - b. To run an incremental backup by date, select **Incremental (date only)**.
 - c. To run a selective backup, select **Always backup**.
 - d. To run an incremental backup without using the journal database, select **Incremental (without journal)**. If you installed the journal engine service and it is running, then by default the **Incremental** command automatically performs a journal-based backup on selected file systems that are being monitored by the journal engine service. This option performs a traditional full incremental backup, instead of the default journal-based backup.
5. Click **Backup**. The **Backup Task List** window displays the backup processing status. When processing completes, the **Backup Report** window displays processing details.

Results

The following are some items to consider when you use the GUI to back up your data.

- 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 that is used is either a default that is selected for you, or one that you assign to the file using an **include** option in the include-exclude options list. Select **Utilities** → **View Policy Information** from the backup-archive client or web client GUI to view the backup policies that are defined by the Tivoli Storage Manager server for your client node. Select **Edit** → **Client Preferences** from the backup-archive client or web client GUI and select the **Include-Exclude** tab in the Preferences editor to display your include-exclude list.
- To modify specific backup options, click the **Options** button. Any options that you change are effective during the current session only.

- To perform subsequent incremental backups, from the Tivoli Storage Manager main window, open the **Actions** menu and select **Backup Domain**.
- The web client GUI cannot browse network resources to perform a backup. No shares are listed if you expand the **Network** branch. It is possible to backup a network resource from the web client as long as the entire file is processed. To do this, the filesystem is specified using the domain option in dsm.opt. E.g. domain all-local \\server\share. To complete the backup, select **Backup Domain** from the **Action** menu. This processes all the filesystems that are specified with the domain option. Alternatively, you can use the GUI Client to perform the backup.

Related concepts:

Chapter 9, “Storage management policies,” on page 267

Related tasks:

“Restoring data by using the GUI” on page 195

“Setting the client scheduler process to run as a background task and start automatically at startup” on page 252

Specifying drives in your domain

When you start the client, it sets your default domain to the drives you specify with the domain option in the dsm.opt file.

About this task

If you do not set the domain option, the default domain is all local fixed drives (the drives on your workstation).

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

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

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

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

You can also select **Actions > Backup Domain** from the backup-archive client GUI to perform these backup functions.

Related reference:

“Domain” on page 371

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 commands to perform different tasks.

About this task

Table 17. Command line backup examples

Task	Command	Considerations
<i>Incremental backups</i>		
Perform an incremental backup of your client domain.	<code>dsmc incremental</code>	See “ Incremental ” on page 683 for more information about the incremental command. See “Full and partial incremental backup” on page 136 for detailed information about incremental backups.
Back up the g: and h: drives in addition to the c:, d:, and e: drives defined in your client domain.	<code>dsmc incremental -domain="g: h:"</code>	See “Domain” on page 371 for more information about the domain option.
Back up all local volumes defined in your client domain <i>except</i> for the c: drive and systemobject domain.	<code>dsmc incremental -domain="all-local -c: -systemobject"</code>	You cannot use the (-) operator in front of the domain keyword all-local. See “Domain” on page 371 for more information. For Windows clients you can also exclude the systemstate domain from backup processing in this way.
Back up all local volumes defined in your client domain <i>except</i> for the c: drive and systemstate domain.	<code>dsmc incremental -domain="all-local -c: -systemstate"</code>	You cannot use the (-) operator in front of the domain keyword all-local. See “Domain” on page 371 for more information.
Back up <i>only</i> the g: and h: drives.	<code>dsmc incremental g: h:</code>	None
Back up all files in the c:\Accounting directory and all its subdirectories.	<code>dsmc incremental c:\Accounting* -sub=yes</code>	See “Subdir” on page 548 for more information about the subdir option.
Assuming that you initiated a snapshot of the C: drive and mounted the snapshot as the logical volume \\florence\c\$\snapshots\snapshot.0, run an incremental backup of all files and directories under the local snapshot and manage them on the Tivoli Storage Manager server under the file space name C:.	<code>dsmc incremental c: -snapshot=\\florence\c\$\snapshots\snapshot.0</code>	See “Snapshotroot” on page 537 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 445 for more information about the incrbydate option.
<i>Selective backups</i>		

Table 17. Command line backup examples (continued)

Task	Command	Considerations
Back up all files in the d:\proj directory.	dsmc selective d:\proj\	Use the selective command to back up specific files, a group of files with similar names, or empty directories and their attributes regardless of whether those files or directories were backed up during your last incremental backup and without affecting the last incremental backup count from the backup server. You can use wildcards to back up multiple files at once. See “ Selective ” on page 771 for more information about the selective command.
Back up the d:\proj directory and all its subdirectories.	dsmc selective d:\proj\ -subdir=yes	See “Subdir” on page 548 for more information about the subdir option.
Back up the d:\hl.doc and d:\test.doc files.	dsmc selective d:\hl.doc d:\test.doc	You can specify as many file specifications as available resources or other operating system limits permit. Separate file specifications with a space. You can also use the filelist option to process a list of files. The backup-archive client opens the file you specify with this option and processes the list of files within according to the specific command. See “Filelist” on page 414 for more information.
Back up a list of files in the c: drive.	dsmc selective -filelist=c:\filelist.txt	Use the filelist option to process a list of files. See “Filelist” on page 414 for more information.
Assuming that you initiated a snapshot of the C: drive and mounted the snapshot as the logical volume \\florence\c\$\snapshots\snapshot.0, run a selective backup of the c:\dir1\sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name C:.	dsmc selective c:\dir1\sub1* -subdir=yes snapshot=\\florence\c\$\snapshots\snapshot.0	See “Snapshotroot” on page 537 for more information.

Related concepts:

“Backup (Windows): Additional considerations” on page 185
Chapter 12, “Using commands,” on page 625

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.

About this task

For example, you might need to delete sensitive data that was backed up (intentionally or unintentionally), and now needs to be removed from the server. Or you might need to delete files that were backed up, but were later found to contain viruses. To determine if you have the authority to delete individual backup copies from the Tivoli Storage Manager server without deleting the entire file space, select **File** → **Connection Information** from the backup-archive client GUI or web client main menu. Your authority status is provided in the **Delete Backup Files** field.

Important: When you delete backup files, *you cannot restore them*. Verify that the backup files are no longer needed before you delete them. Tivoli Storage Manager prompts whether you want to continue with the delete. If you specify *yes*, the specified backup files are immediately deleted and removed from Tivoli Storage Manager server storage.

To delete backup copies using the Tivoli Storage Manager GUI or web client:

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. 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. A directory is deleted only if you select **Delete All Objects**.

Results

To delete backup copies using the Tivoli Storage Manager command line client, use the **delete backup** command.

Related reference:

“Delete Backup” on page 674

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:

“Restore data from a backup set” on page 206

Pre-backup considerations (Windows)

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.

All Windows 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 392

“Lanfreecommmethod” on page 450

“Lanfreeshmport” on page 452

“Lanfreessl” on page 453

“Lanfreetcpport” on page 453

“Lanfreetcpserveraddress” on page 454

Unicode file spaces (Windows)

The Windows client is Unicode-enabled. However, client versions before Version 4.2 were not enabled for Unicode.

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

Related concepts:

“Considerations for Unicode-enabled clients” on page 430

Related reference:

“Autofsrename” on page 333

“Detail” on page 363

“Query Filespace” on page 706

“Restore” on page 730

“Retrieve” on page 766

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 462

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`**.
- 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.
- Reduce the number of files per file system by spreading the data across multiple file systems.

Related reference:

"Snapdiff" on page 529

"Exclude options" on page 400

"Incrbydate" on page 445

"Memoryefficientbackup" on page 462

"Resourceutilization" on page 508

Performing a backup with limited bandwidth

If you plan to perform a backup over a network device that has limited bandwidth, such as a modem, you can help reduce network traffic by using *adaptive subfile backup*.

About this task

An adaptive subfile backup sends only changed portions of a file to the server during successive backup operations.

Perform the following steps to use adaptive subfile backup:

Procedure

1. Ensure that the server allows this type of backup. Check with your system administrator.
2. Set the **subfilebackup**, **subfilecachepath** and **subfilecachesize** options in your client options file **dsm.opt**. If you set the **subfilebackup** option to *yes* and no **exclude.subfile** statements exist, the backup-archive client considers all files for adaptive subfile backup processing.
3. Add an **exclude.dir** option to your client options file, which excludes the directory containing the subfile cache information.

Results

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

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

Related concepts:

"Adaptive subfile backup processing" on page 403

"Determine compression and encryption processing" on page 100

Related reference:

“Exclude options” on page 400

“Subfilebackup” on page 549

“Subfilecachepath” on page 550

“Subfilecachesize” on page 551

Control processing with an include-exclude list

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

Use the include and exclude options in the client options file (dsm.opt) to define which files to include or exclude from incremental or selective backup processing. A file is eligible for backup unless excluded by an exclude option. It is not necessary to use an include option to include specific files for backup unless those files are in a directory that contains other files that you want to exclude.

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

Tivoli Storage Manager uses *management classes* to determine how to manage your backups on the server. Every time you back up a file, the file is assigned a management class. The management class is either a default that is chosen for you, or one you assign to the file by 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.

You can also add include-exclude statements in the backup-archive client GUI directory tree. You can use the **preview** command to see the resultant effects of the currently defined include-exclude list without need of running an actual backup operation.

Related tasks:

“Creating an include-exclude list” on page 92

“Setting the client scheduler process to run as a background task and start automatically at startup” on page 252

Related reference:

“Preview Backup” on page 694

Data encryption during backup or archive operations

For the strongest possible encryption, use 256-bit Advanced Encryption Standard (AES) data encryption, with the **encryptiontype** option. AES 128-bit encryption is currently the default.

The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received.

Attention: If the encryption key password is not saved in the Windows Registry, and you have forgotten the password, your data cannot be recovered.

The **include.encrypt** option is the only way to enable encryption on the Backup-Archive client. If no include.encrypt statements are used, encryption will not occur.

Encryption is not compatible with VMware virtual machine backups that use the incremental forever backup modes (**MODE=IFIncremental1** and **MODE=IFFull11**). If the client is configured for encryption, you cannot use incremental forever backup. However, you can use the full or incremental backup modes (**MODE=Full1** and **MODE=Incremental1**).

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

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

While restoring an encrypted file, you are prompted for the key password to decrypt the file in the following cases:

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

Related concepts:

“Backup (Windows): Additional considerations” on page 185

Related reference:

“Encryptiontype” on page 393

“Encryptkey” on page 394

“Exclude options” on page 400

“Include options” on page 430

Maximum file size for operations

The maximum file sizes for backup and restore, and archive and retrieve operations depends on the Windows file system that is used.

The following table shows the maximum file size, in bytes, for backing up, restoring, and retrieving data.

Table 18. Maximum file size

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

How Tivoli Storage Manager handles long user and group names

Tivoli Storage Manager can handle user and group names that are up to 64 characters without any issues. However, names longer than 64 characters require special handling.

Restriction: Do not exceed the 64-character limit for user and group names. Tivoli Storage Manager shortens the name to fall within this limit by using the following algorithm: Use the first 53 characters, append a forward slash (/), and then use 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.

Incremental, selective, or incremental-by-date backups (Windows)

Your administrator might set up schedules to automatically back up files. This section contains information about how to back up files without a schedule.

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

If you migrate files with Tivoli Storage Manager HSM for Windows, there can be consequences for backup operations.

Related concepts:

 Backup and restore of migrated files

Related tasks:

“Setting the client scheduler process to run as a background task and start automatically at startup” on page 252

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 drives, 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
 - Sparse, reparse point or encrypted file attributes.
 - NTFS or ReFS file security descriptors: Owner Security Identifier (SID), Group SID, Discretionary Access Control List (ACL), and System ACL.
 - Directory attributes

If only the following attributes change, the attributes are updated on the Tivoli Storage Manager server, but the file is not backed up:

- Read-only or read/write
- Hidden or not hidden
- Compressed or not compressed

The archive attribute is not examined by Tivoli Storage Manager in determining changed files.

- 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 directory modification time stamp changed since the last backup

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 `preserveLastAccessDate` 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 267

Related reference:

"Exclude options" on page 400

"PreserveLastAccessDate" on page 489

Journal-based backup

Journal-based backup is an alternate method of backup that uses a change journal maintained by the Tivoli Storage Manager journal service process.

Journal-based backup is supported for all Windows clients.

To support journal-based backup, you must configure the journal engine service using the **dsmcutil** command or the client GUI setup wizard.

A backup for a particular file system will be journal-based when the Tivoli Storage Manager journal service has been installed and configured to journal the particular file system, and a valid journal has been established for the file system.

The primary difference between traditional incremental backup and journal-based backup is the method used for backup and expiration candidates.

Traditional incremental backup obtains the list of backup and expiration candidates by building comprehensive lists of local objects, and lists of active server objects for the file system being backed up. The local lists are obtained by scanning the entire local file system. The server list is obtained by querying the entire server inventory for all active objects.

The two lists are compared, and candidates are selected according to the following criteria:

- An object is selected as a backup candidate if it exists in the local list but doesn't exist in the server list. The object is also a backup candidate if it exists in both lists, but differs according to incremental criteria (for example, attribute changes, date and size changes).

- An object is selected as an expiration candidate if it exists in the server list, but doesn't exist in the local list.

Journal-based backup obtains the candidates list of objects to backup and expire by querying the journal service for the contents of the change journal of the file system being backed up.

Change journal entries are cleared (marked as free) after they have been processed by the backup client and committed on the Tivoli Storage Manager server.

Journal-based backup is activated by configuring the journal service to monitor specified file systems for change activity.

Journal-based backup is enabled by successfully completing a full incremental backup.

The journal engine service does not record changes in specific system files, such as the registry, in the journal. Therefore, a journal-based backup will not back up this file. See the journal service configuration file, `tsmjbbd.ini`, in the client installation directory for excluded system files.

You can use journal-based backup when backing up file systems with small or moderate amounts of change activity between backup cycles. If you have many file changes between backup cycles, you will have very large change journals. Many changes to the journal-based backup file might pose memory and performance problems that can negate the benefits of journal-based backup. For example, creating, deleting, renaming, or moving very large directory trees can also negate the benefit of using journal-based backup instead of normal incremental backup.

Journal-based backup is not intended to be a complete replacement for traditional incremental backup. You should supplement journal-based backup with a full progressive incremental backup on a regular basis. For example, perform journal-based backups on a daily basis, and full incremental backups on a weekly basis.

Journal-based backup has the following limitations:

- Individual server attributes are not available during a journal-based backup. Certain policy settings such as copy frequency and copy mode might not be enforced.
- Other operating-system 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.

Note:

1. Multiple journal-based backup sessions are possible.
2. When using antivirus software, there are limitations to journal-based backup.
3. A journal-based backup might not fall back to the traditional incremental backup if the policy domain of your node is changed on the server. This depends on when the policy set within the domain was last updated and the

date of the last incremental backup. In this case, you must force a full traditional incremental backup to rebind the files to the new domain. Use the `nojurnal` option with the **incremental** command to specify that you want to perform a traditional full incremental backup, instead of the default journal-based backup.

Related tasks:

“Configuring the journal engine service” on page 45

Restore processing with journal-based backups (Windows):

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.

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

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

Files that were renamed after the last incremental backup, but otherwise remain unchanged, will not be backed up. Renaming a file does not change the

modification date and time of the file. However, renaming a file does change the modification date of the directory in which it is located. In this case, the directory is backed up, but not the files it contains.

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

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.

NetApp snapshot difference

For NAS and N-Series file servers that are running ONTAP 7.3.0, or later, you can use the `snappdiff` 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 **`dsmdc 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.

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 Windows systems, the snapshot directory is in `~snapshot`.

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 concepts:

“Comparing incremental-by-date, journal-based, and NetApp snapshot difference to full incremental and partial incremental backups” on page 141

Related tasks:

“Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 83

“Running a snapshot differential backup with an HTTPS connection”

Related reference:

“Snappdiffhttps” on page 534

“Snappdiff” on page 529

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 83.

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 **snapdiff** and **snapdiffhttps** options at the command-line interface.

For example, on a Windows system with a network share \\netappl.example.com\vol1, where netappl.example.com is a filer, issue the following command:

```
dsmc incr \\netappl.example.com\vol1 -snapdiff -snapdiffhttps
```

Related concepts:

“Snapshot differential backup with an HTTPS connection” on page 142

Related reference:

“Snapdiffhttps” on page 534

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 GUI” on page 125

Related reference:

“Selective” on page 771

Backing up files from one or more file spaces for a group backup (Windows)

Use the **backup group** command to create and back up a group from a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.

About this task

A *group backup* creates 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. Use the **include** option to bind a group to a management class.
- Existing **exclude** statements for any files in the group are ignored.
- All objects in the group are exported together.
- All objects in the group are expired together as specified in the management class. No objects in a group are expired until all other objects in the group are expired, even when another group they belong to gets expired.

A group backup can be added to a backup set.

You can perform a full or differential backup by using the mode option.

Procedure

Enter the **backup group** command to start a group backup.

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

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

Related concepts:

“Restore data from a backup set” on page 206

Related reference:

“Backup Group” on page 642

“Include options” on page 430

“Mode” on page 463

Backing up data with client-node proxy support (Windows)

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

Before you begin

The following considerations apply 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.
- You cannot use **asnodename** with the **backup nas** command.
- You cannot use **asnodename** with the **fromnode** option.
- If you use **asnodename** to backup and restore volumes that are in a cluster configuration, do not use **clusternode yes**.
- You cannot use **asnodename** to back up or restore system state.
- If an agent node restores data from a backup set, the system state object in the backup set is not restored.
- You can use **asnodename** with the **backup image** command, but you must specify the volume by UNC name. You cannot use the drive letter.
- If you use the same **asnodename** value to back up files from different machines, you need to keep track which files or volumes are backed up from each system so that you can restore them to the correct location.
- All agent nodes in a multiple node environment should be of the same platform type.
- Do not use target nodes as traditional nodes, especially if you encrypt your files before backing them up to the server.

About this task

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

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

Using an agent node to backup target nodes is useful when the workstation responsible for performing the backup can change over time, such as with a cluster configuration.

The `asnodename` option allows data to be restored from a different system than the one which performed the backup.

Use the `asnodename` option with the appropriate command to back up, archive, restore, and retrieve data under the target node name on the Tivoli Storage Manager server. This support is only available with Tivoli Storage Manager Version 5.3 and higher server and client.

Procedure

To enable this option, follow these steps:

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

Related reference:

"`Asnodename`" on page 324

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 operation on behalf of the target node. Each agent node must use the `asnodename` option within their schedule to perform multiple node backup for the agent node.

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>`

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

- The administrator registers all 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, by issuing 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, by issuing the following commands:
 - define schedule standard proxy1 description="NODE-A proxy schedule" action=incremental options="-asnode=NODE-Z" objects=C: startdate=05/21/2005 starttime=01:00
 - define schedule standard proxy2 description="NODE-B proxy schedule" action=incremental options="-asnode=NODE-Z" objects=D: startdate=05/21/2005 starttime=01:00
 - define schedule standard proxy3 description="NODE-C proxy schedule" action=incremental options="-asnode=NODE-Z" objects=E: startdate=05/21/2005 starttime=01:00

Note: Place the 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.

Start the schedules by either configuring a Tivoli Storage Manager schedule service, or by using the following client command: `dsmc sched`

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

Important:

- Each schedule can be started from a different workstation or LPAR.
- After running the schedules, any proxied client can query and restore all 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 324

"Session settings and schedules for a proxy operation" on page 326

 [DEFINE SCHEDULE command](#)

Associate a local snapshot with a server file space (Windows)

Use the `snapshotroot` option with the **incremental** and **selective** commands 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.

Related reference:

“Snapshotroot” on page 537

Backing up Windows system state

The backup-archive client uses VSS to back up all system state components as a single object, to provide a consistent point-in-time snapshot of the system state. System state consists of all bootable system state and system services components.

About this task

The client supports the Microsoft volume shadow copy service (VSS) on the supported Windows clients.

System state is represented by several VSS writers of type "bootable system state" and "system service". Of these, the System Writer is the largest part of the system state in terms of number of files and size of data. By default, the System Writer backup is incremental. You can use the `systemstatebackupmethod` option to perform full backups of the System Writer. For more information, about this option, see “Systemstatebackupmethod” on page 552. The client always backs up all other writers in full.

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

You must be a member of the Administrators or Backup Operators group to back up system state information.

To back up a system state object using the command line:

1. On the command line, use the **backup systemstate** command to back up all system state or system services components as a single object.
2. Use the **query systemstate** command to display information about a backup of the system state on the Tivoli Storage Manager server.

To back up a system state object using the GUI:

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

System and boot files are backed up as a group only if one of the members of the group (one of the files) changes. If the files have not changed since the last backup, the system and boot files are not redundantly backed up.

By default, system state backups are bound to the default management class. To bind them to a different management class, use the `include.systemstate` option; specify **all** as the pattern, and specify the name of the new management class.

You can use the `domain` option to exclude the entire system state from domain incremental backup processing.

The system `dllcache` directory is now included in the boot partition backup of Windows systems. When the `dllcache` files are not available when you restore a Windows computer, system recovery might require availability of the operating system installation media. By backing up the `dllcache` directory, you can avoid the need for installation media during system restores.

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

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

On Windows clients, **backup systemstate** also backs up ASR data.

Related tasks:

“Restoring Windows system state” on page 199

Related reference:

“Backup Systemstate” on page 650

“Domain” on page 371

“Exclude options” on page 400

“Include options” on page 430

“Query Systemstate” on page 720

“Restore Systemstate” on page 753

Backing up Automated System Recovery files

You can back up Automated System Recovery (ASR) files in preparation for recovering the Windows disk configuration information and system state in case a catastrophic system or hardware failure occurs.

About this task

The backup-archive client backs up ASR data when the backup-archive client backs up the Windows system state.

Procedure

To back up ASR files on Windows operating systems, use the **backup systemstate** command.

Results

The client generates the ASR files in the \adsm.sys\ASR staging directory on the system drive of your local workstation and stores these files in the ASR file space on the Tivoli Storage Manager server.

Related concepts:

“Preparation for Automated System Recovery”

Related tasks:

“Restoring Automated System Recovery files” on page 201

Related reference:

“Backup Systemstate” on page 650

Preparation for Automated System Recovery

Specific backups and media are required for Windows Automated System Recovery (ASR).

Creating a client options file for Automated System Recovery

Before you can recover a Windows computer by using Automated System Recovery (ASR), you must create an options file. The options file is unique for each computer.

About this task

This task assumes that you created a generic bootable WinPE CD or DVD. A generic bootable WinPE CD does not contain the client options file (dsm.opt) because the options file is unique for each computer. This task helps you create a computer-specific options file.

The Windows Preinstallation Environment (WinPE) requires particular options values.

Procedure

1. Find a copy of the client options file. You can find the file in several places:
 - There is an options file in the installation directory of an installed Tivoli Storage Manager client. The default installation location is C:\Program Files\Tivoli\TSM\baclient\dsm.opt. If you have the options file for the computer that you want to restore, this options file requires the fewest modifications.
 - There is a sample options file in the client installation package. The path in the package is TSM_BA_Client\program files\Tivoli\TSM\config\dsm.smp. Rename the file to dsm.opt.
2. Edit dsm.opt.
 - a. Enter a writable location for the error log. The backup-archive client creates several log files. Use the option errorlogname to specify the log file location. For example, in the dsm.opt file, specify errorlogname x:\dsmerror.log.

Note: This example uses x: because in WinPE mode, the default system drive is x:.
 - b. Enter the client node name with the nodename option.

- c. Optional: If you plan to restore the system state from files that are stored on a Tivoli Storage Manager server, enter the server connection information. Enter appropriate values for the `commethod` and `tcpserveraddress` options.
 - d. Optional: If you know the password for the node, enter the password with the `password` option.
3. Copy the `dsm.opt` file to media that the target computer can read during Automated System Recovery.
 4. Optional: Copy Tivoli Storage Manager client registry information to media that the target computer can read during Automated System Recovery. Use the **regedit.exe** utility to export Tivoli Storage Manager client registry entries from the `HKLM\SOFTWARE\IBM` key. For example, from a command prompt window, run this command:


```
regedit /e tsmregistry.out "HKEY_LOCAL_MACHINE\SOFTWARE\IBM"
```

Copy the `tsmregistry.out` file to media that the target computer can read during ASR.

During ASR, you can import the registry entries from the `tsmregistry.out` file. The backup-archive client can use the registry entries in the WinPE environment to access backup copies on the Tivoli Storage Manager server.

Note: Saving registry entries is optional because there are other ways to get access to the password-protected Tivoli Storage Manager server. You can access the server with the following methods:

- If you know the node password, you can type the password when prompted during recovery.
- Request the Tivoli Storage Manager administrator to change the node password and provide you with the new password at the time of recovery.
- Provide the password information in the `dsm.opt` file.

If the files you want to restore are included in a backup set on tape or on a CD or DVD, then you do not need to access the Tivoli Storage Manager server.

Results

You created an options file that contains client configuration information that is unique for each computer. This information complements the generic bootable WinPE CD.

Related tasks:

“Creating a bootable WinPE CD” on page 202

Backing up the boot drive and system drive for Automated System Recovery

Before you can recover your Windows computer by using Automated System Recovery (ASR), you must have a complete backup of the boot drive and system drive.

Procedure

1. Perform a full incremental backup of your system and boot drives. Assuming that your system and boot files are on the `c:` drive, enter the following command:


```
dsmc incremental c:
```
2. Back up the system state. To back up the system state, enter the following command:

```
dsmc backup systemstate
```

To verify that you backed up the system state, enter the following command:

```
dsmc query systemstate
```

You can specify `-showmembers=yes` to display file level detail.

Related concepts:

“Full and partial incremental backup” on page 136

Related tasks:

“Backing up Windows system state” on page 149

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.

These volumes can be formatted NTFS or ReFS, or unformatted RAW volumes. If a volume is NTFS-formatted, only those blocks that are used by the file system or smaller than the **imagegapsize** parameter are backed up.

Normally you cannot restore an image backup of the system drive over itself since an exclusive lock of the system drive is not possible. However, in a Windows pre-installation environment (WinPE), an image restore of the system drive is possible. For information about restoring data in a WinPE environment, see technote 7005028.

You cannot restore an image backup to the volume on which the client is running. Consider installing the backup-archive client on the system drive.

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

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

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.

If online image support is configured, the client performs an online image backup, during which the volume is available to other system applications. The snapshot provider, as specified by the `snapshotproviderimage` option, maintains a consistent image of a volume during online image backup.

You can use the `snapshotproviderimage` option with the **backup image** command or the `include.image` option to specify whether to perform an offline or online image backup.

Related tasks:

“Configuring online-image backup support” on page 82

Related reference:

“`Snapshotproviderimage`” on page 536

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.

- *To perform an offline or online image backup you must have administrative authority on the system.*
- You do not need more than one drive to perform an image backup.
- Ensure that no other application is using the volume when you run an offline image backup. To ensure a consistent image during backup processing, the client locks the volume, so that no other applications can write to it. If the volume is in use when the client attempts to lock the volume, the backup fails. If the client cannot lock a volume because it is in use, you can perform an online image backup.
- Use the `include.image` option to assign a management class to the volume image. If you do not assign a management class, the default management class is used for the image.

Note: If the `snapshotproviderimage` option is set to *none*, then the copy serialization parameters set by the management class is used.

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

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

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

Related concepts:

Chapter 9, “Storage management policies,” on page 267

Related reference:

“Exclude options” on page 400

“Include options” on page 430

“Snapshotproviderimage” on page 536

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 GUI” on page 125

“Performing an image backup using the GUI” on page 157

“Restoring an image using the GUI” on page 204

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 157

“Restoring an image using the GUI” on page 204

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 19. Comparing incremental image backup methods

Method 1: Using image backup with file system incremental	Method 2: Using image backup with incremental-by-date image backup
Files are expired on the server when they are deleted from the file system. On restore, you have the option to delete files which are expired on server from image.	Files are not expired on server. After the image incremental restore completes, all of the files that are deleted on the file system after the image backup are present after the restore. If file systems are running at or near capacity, an out-of-space condition could result.
Incremental backup time is the same as regular incremental backups.	Incremental image backup is faster because the client does not query the server for each file that is copied.
Restore is much faster compared to a full incremental file system restore.	Restore is much faster compared to a full incremental file system restore.
Directories deleted from the file system after the last image backup are not expired.	Directories and files deleted from the file system after the last full image backup are not expired.

Performing an image backup using the GUI

If the 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 client GUI image backup option, the backup operation is run according to the `snapshotproviderimage` setting in your client options file (`dsm.opt`). If the online image support is configured, the client performs an online image backup, during which the volume is available to other system applications.

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 an offline image backup, select **Image Backup** from the drop-down list.
- To perform an online image backup, select **Snapshot Image Backup** from the drop-down list.
- To perform an incremental-by-date image backup, select **Incremental image (date only)** from the drop-down list.

The following are some items to consider when you perform an online image backup:

- To modify specific backup options, click the **Options** button. The options you select are effective only during the current session.
- Because image backup allows you to back up only used blocks in a file system, the stored image size on the Tivoli Storage Manager server could be smaller than the volume size. For online image backups, the stored image can be larger than the file system based on the size of the cache files. To determine the actual stored image size, select **View > File Details**. The actual stored image size is noted in the Stored Size field.
- To modify specific backup options, click the **Options** button. The options you select are effective only during the current session.
- Because image backup allows you to back up only used blocks in a file system, the stored image size on the Tivoli Storage Manager server could be smaller than the volume size. For online image backups, the stored image can be larger than the file system based on the size of the cache files. To determine the actual stored image size, select **View > File Details**. The actual stored image size is noted in the Stored Size field.

Related reference:

"Snapshotproviderimage" on page 536

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.

You can use the snapshotproviderimage option with the **backup image** command or the include.image option in your dsm.opt file or on the command line to specify whether to perform an offline or online image backup.

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 644

"Mode" on page 463

"Restore Image" on page 748

"Snapshotproviderimage" on page 536

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`, `snapdiff`, `createnewbase`, and `diffsnapshot` options for more information.

After configuring NDMP support, the server connects to the NAS device and uses NDMP to initiate, control, and monitor each backup and restore operation. The NAS device performs outboard data transfer to and from the NAS file system to a locally attached library.

Filer to server data transfer is available for NAS devices that support NDMP Version 4.

The benefits of performing backups using NDMP include the following:

- LAN-free data transfer.
- High performance and scalable backups and restores.
- Backup to local tape devices without network traffic.

The following support is provided:

- Full file system image backup of all files within a NAS file system.
- Differential file system image backup of all files that have changed since the last full image backup.
- Parallel backup and restore operations when processing multiple NAS file systems.
- Choice of interfaces to initiate, monitor, or cancel backup and restore operations:
 - Web client
 - Backup-archive client command interface
 - Administrative client command line interface (backup and restore operations can be scheduled using the administrative command scheduler)
 - Administrative web client

The following functions are *not* supported:

- Archive and retrieve
- Client scheduling. Use server commands to schedule a NAS backup.
- Detection of damaged files.
- Data-transfer operations for NAS data stored by Tivoli Storage Manager:
 - Migration
 - Reclamation
 - Export
 - Backup set generation

Related concepts:

“NDMP support requirements (Extended Edition only)” on page 7

Related reference:

"Diffsnapshot" on page 364

"Incremental" on page 683

"Snapdiff" on page 529

"Snapshotroot" on page 537

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.

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 33

“Restore NAS file systems” on page 235

Related reference:

“Toc” on page 563

Back up NAS file systems using the command line

You can use the command line to back up NAS file system images.

Table 20 lists the commands and options that you can use to back up NAS file system images from the command line.

Table 20. NAS options and commands

Option or command	Definition	Page
domain.nas	Use the domain.nas option to specify the volumes to include in your default domain for NAS backups.	“Domain.nas” on page 375
exclude.fs.nas	Use the exclude.fs.nas option to exclude file systems on the NAS file server from an image backup when used with the backup nas command. This option is valid for all Windows clients.	“Exclude options” on page 400
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 Table of Contents (TOC) information is saved during a NAS file system image backup, using the <i>toc</i> option with the include.fs.nas option in your client options file.. This option is valid for all Windows clients.	“Include options” on page 430

Table 20. NAS options and commands (continued)

Option or command	Definition	Page
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 714
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 648
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 563
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 692
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 670
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 700
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 706
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 677

A NAS file system specification uses the following conventions:

- NAS nodes represent a new node type. The NAS node name uniquely identifies a NAS file server and its data to Tivoli Storage Manager. You can prefix the NAS node name to the file specification to specify the file server to which the include statement applies. If you do not specify a NAS node name, the file system you specify applies to all NAS file servers.
- Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.
- NAS file system designations on the command line require brace delimiters {} around the file system names, such as: {/vol/vol0}. Do not use brace delimiters in the option file.

Note: When you initiate a NAS backup operation 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 563

Methods for backing up and recovering data on NAS file servers accessed by CIFS protocol

The backup-archive client can process network-attached storage (NAS) file-server data that is accessed by using the Common Internet File System (CIFS) protocol.

You can use the following methods to back up and recover data on NAS devices:

- Use the backup-archive client to back up and restore data, by using CIFS to access files from the backup-archive client. Data can be stored on the Tivoli Storage Manager server with file-level granularity, by using the progressive-incremental backup method. The data is stored in the Tivoli Storage Manager storage hierarchy and can be migrated, reclaimed, and backed up to a copy storage pool.

This method increases processor usage when the client accesses individual files. The method requires that the data to flow through the client. This method also requires that the data flows through the Tivoli Storage Manager server unless a LAN-free configuration is used.

- Use the `snapdiff` option to mitigate the performance problems of CIFS backup. This option stores data with file-level granularity by using progressive incremental backup for CIFS.
- Use a backup-archive client that is running on the NAS device, if you can use external programs with the NAS operating system.

This method decreases processor usage of CIFS. Data can be stored on the Tivoli Storage Manager server with file-level granularity by using progressive-incremental backup. The data is stored in the Tivoli Storage Manager storage hierarchy and can be migrated, reclaimed, and backed up to a copy storage pool. This method requires that data flow through the backup-archive client. This method also requires that the data flows over a network and through the Tivoli Storage Manager server unless a LAN-free configuration is used.

- Use NDMP with the backup-archive client. File systems are backed up as full images (all files) or differential images (all files that changed since the last full backup). Backed up images are stored on a tape device that is accessed by the NAS file server. This method provides high performance because there is no data flow through a Tivoli Storage Manager client or server. Data that is backed up to the server by using NDMP cannot be migrated, reclaimed, or backed up to a copy storage pool.

The following limitations exist for NAS file server data when it is accessed by using CIFS:

- File and directory security information might be inaccessible when the Windows account that is performing the backup is not a member of the Domain Administrators group of the domain the NAS file server is a trusted member of. It is also possible that these security access failures might prevent the entire file or directory from being backed up.
- Performance degradation occurs because data is being accessed remotely.

- The mapped drives appear to the client as NTFS file systems, but they might not have full NTFS functionality. For example, the encryption attribute of a file is set, but when the client backs up the file, the backup fails because the volume-level encryption setting indicates that encryption cannot be used for the volume. ReFS file systems also appear to the client as NTFS file systems.

Tip: Use NDMP with the backup-archive client on a NAS file server to back up and restore volumes instead of backing up and restoring the volumes by using remote mapped drives.

Related reference:

“Snapdiff” on page 529

Support for CDP Persistent Storage Manager

Persistent Storage Manager (PSM) is the snapshot technology that is included with a number of Microsoft Server Appliance Kit-based NAS boxes that include the IBM TotalStorage NAS 200, 300, and 300G.

You can use the backup-archive client to back up the persistent images (PI) of a volume that is produced by PSM. You must first ensure that the volume has a label. You can then use PSM to schedule or create a persistent image with a specific image name, such as `snapshot.daily`, and set the number of images to save to 1. PSM overwrites the PI as needed and you can use the client to incrementally back up the PI. In this case, the client backs up only the files that changed between snapshots. One advantage of backing up a PSM PI rather than the actual volume, is that there are no open files in the PI.

Consider the following items before you use Persistent Storage Manager:

- By default, a PSM schedule uses a variable name (*snapshot.%i*) and keeps a number of images.

Important: Do not use the client with PSM in this manner. The client considers each image as unique and makes a complete copy of each image.

- The client requires that the volume used to make the PI has a label. If the volume does not have a label, the client does not back up its PI.
- You use the image backup function to back up the original volume that is used to create the PI. However, you cannot use the backup image function to back up the PI.
- To avoid backing up unnecessary files when you back up PSM, include the following entries in your client option file (`dsm.opt`):

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

Backing up 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 21 on page 165 lists the backup and restore operations for VMware virtual machines that the backup-archive client can implement on Windows platforms.

Restriction: You can complete VMware backup and restore operations with the backup-archive client only on 64-bit Windows operating systems. Backup-archive clients on Windows 32-bit operating systems cannot complete VMware backup and restore operations.

Table 21. Backup and restore capabilities for VMware virtual machines on Windows platforms

Capability	Comment
Full VM backup:	<p>Backs up a full virtual machine. The following types of full VM backup are possible:</p> <ul style="list-style-type: none"> • Create an image backup of the full virtual machine. • Create a backup of files or directories on the virtual machine. • Create an image backup of a Microsoft Hyper-V virtual machine.
Incremental VM backup:	<p>Requires the IBM Tivoli Storage Manager for Virtual Environments licensed product.</p> <p>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:</p> <ul style="list-style-type: none"> • Restore the most recent full backup. • Restore each incremental backup as required to restore your content to a specific point in time. <p>For more efficient backups, use the incremental-forever backup modes.</p>
Full VM incremental-forever backup:	<p>Requires the IBM Tivoli Storage Manager for Virtual Environments licensed product.</p> <p>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.</p> <p>You cannot use this backup mode to back up a VMware virtual machine if the client is configured to encrypt the backup data.</p>

Table 21. Backup and restore capabilities for VMware virtual machines on Windows platforms (continued)

Capability	Comment
Incremental-forever-incremental VM backup:	<p>Requires the IBM Tivoli Storage Manager for Virtual Environments licensed product.</p> <p>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.</p> <p>The incremental-forever-incremental backup mode provides the following benefits:</p> <ul style="list-style-type: none"> • Improves the efficiency of backing up virtual machines. • Simplifies data restore operations. • Optimizes data restore operations. <p>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.</p> <p>You cannot use this backup mode to back up a VMware virtual machine if the client is configured to encrypt the backup data.</p>
Item recovery for files and folders from a full backup of the virtual machine:	<p>Requires the IBM Tivoli Storage Manager for Virtual Environments licensed product.</p> <p>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.</p>
Full restore of the virtual machine:	Restores all of the file systems, virtual disks, and the virtual machine configuration.
File-level backup:	Uses the VMware vStorage API for Data Protection and VMware vStorage backup server to create backups of individual files in the virtual machine. You create include and exclude rules to identify the files to back up. For file-level backups, files are backed up as individual files rather than as an image backup.
File-level restore of the virtual machine:	<p>The restore approach depends on the type of backup of the virtual machine:</p> <ul style="list-style-type: none"> • If you have a license for Tivoli Storage Manager for Virtual Environments, you can restore files and directories from a full VM image backup. • Backup-archive client users can restore files and directories that are created by file-level backups of a virtual machine. You use the restore command to restore individual files from a file-level backup of a virtual machine, not the restore vm command.

For releases before Tivoli Storage Manager version 6.3, the backup-archive client was able to use VMware Consolidated Backup (VCB) to back up data on virtual machines. From Tivoli Storage Manager version 6.3, the use of VCB for backups is no longer available. The backup-archive client now backs up data by using the VMware vStorage API for Data Protection. You can use the **restore VM** command to restore backups that were created by using VCB. Using the vStorage API, the backup-archive client can create either full backups or file-level backups of VMware virtual machines. The vStorage API provides the following benefits:

- Less storage is needed on the vStorage backup server because a temporary directory is not needed to contain the backup.

- Backups that are created by using the vStorage API do not require the VMware converter tool to restore the virtual machine.

The file-level backup method uses the progressive incremental technique. You restore the files by using the normal file-level restore function that is provided by the backup-archive client. This backup approach is an alternative to backing up the full VMware virtual machine. Typically, backing up the full virtual machine is the preferred approach. Before you create file-level backups for a VMware virtual machine on Windows systems, first evaluate the benefits of backing up the full VMware virtual machine. Table 22 compares the different approaches for backing up VMware virtual machines on Windows systems.

Table 22. Comparison of full backups and file-level backups for VMware virtual machines on Windows systems

Capability	Full backups	File-level backups
VMware vStorage API for Data Protection functions:	Yes	Yes
Create backups for virtual machines that are running on an operating system other than Microsoft Windows:	Yes	No
Create incremental backups of the virtual machine:	Yes	Yes
Restore a complete virtual machine:	Yes	No
Restore individual files from the backup:	Yes	Yes
Use the same backup format as progressive incremental backup, and allow for file and folder restore by using the backup-archive client:	No	Yes
Define file-level include and exclude rules:	No	Yes

Related concepts:

“Parallel backups of virtual machines” on page 174

Related tasks:

“Preparing the environment for full backups of VMware virtual machines”

“Creating full backups for VMware virtual machines” on page 170

“Preparing the environment for file-level backups of VMware virtual machines” on page 171

“Creating file-level backups for VMware virtual machines” on page 173

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. When the backup-archive client runs on a vStorage backup server, this client configuration is called the Tivoli Storage Manager *data mover node*. A Windows system that is a data mover must have the 64-bit Windows client installed on it. A data mover node typically uses the SAN to back up and restore data. If you configure the data mover node to directly access the storage volumes, turn off automatic drive letter assignment. If you do not turn off letter assignments, the client on the data mover node might corrupt the Raw Data Mapping (RDM) of the virtual disks. If the RDM of the virtual disks is corrupted, backups fail. Consider the following conditions for restore configurations:

The data mover node is on a Windows Server 2008 or Windows Server 2008 R2 system:

If you plan to use the SAN to restore data, you must set the Windows SAN policy to **OnlineAll**. Run **diskpart.exe** and type the following commands to turn off automatic drive letter assignment and set the SAN policy to **OnlineAll**:

```
diskpart
  automount disable
  automount scrub
  san policy OnlineAll
exit
```

The backup-archive client is installed in a virtual machine on a Windows Server 2008 or Windows Server 2008 R2 system:

If you plan to use the hotadd transport to restore data from dynamically added disks, the SAN policy on that system must also be set to **OnlineAll**.

Whether the client uses the SAN or hotadd transport, the Windows SAN policy must be set to **OnlineAll**. If the SAN policy is not set to **OnlineAll**, restore operations fail, and the following message is returned:

```
ANS9365E VMware vStorage API error.
TSM function name: vddksdk Write
TSM file : vmvddkdsd.cpp (2271)
API return code : 1
API error message : Unknown error
ANS0361I DIAG: ANS1111I VmRestoreExtent(): VixDiskLib Write
FAILURE startSector=512 sectorSize=512 byteOffset=262144,
rc=-1
```

For a description of the vStorage transport settings and how you can override the defaults, see the following topic:

“Vmvstortransport” on page 618

- b. Install the backup-archive client on the vStorage backup server. At the custom setup page of the installation wizard, select **VMware vStorage API runtime files**.

Important: If you are moving the backup data by using backups that are not in a LAN, the SAN must have separate connections for tape and disk.

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.

Tip: On Windows systems, you can get the server full computer name by right-clicking on **My Computer**. Click the Computer Name tab and look at the name listed next to **Full computer name**.

- c. From the backup-archive client on the vStorage backup server, run the following command to register the node:

```
register node my_vm_name my_password
```

Where *my_vm_name* is the full name of the virtual machine that you are backing up.

4. If you back up a virtual machine where volumes are mounted to directories rather than drive letters, files might not be stored in the correct location. An error might be caused because the mount point does not correspond to the actual mount points of backed up files. An error is caused because the mount points for a virtual machine that is running Windows do not have a drive letter assignment. When you use the VMware vStorage APIs for Data Protection, a filespace name is created that includes a number assignment. The filespace names that are created for the mount point do not correspond to the actual mount points of the backed up file.

To back up or restore files to their original location, use the following steps:

- a. To restore files to their original location, map the drive or assign a drive letter to the mount point from the virtual machine.
 - b. If you restore a file that the vStorage API renamed, select a different restore location.
 - c. When using mount points without drive letter assignments, use an include or exclude statement for that volume. See the following example of an exclude statement:

```
exclude \\machine\3$\dir1\...\*.doc
```

Related tasks:

“Creating full backups for VMware virtual machines” on page 170

Related reference:

“Backup VM” on page 658

“Query VM” on page 725

“Restore VM” on page 758

“Vmchost” on page 583

“Vmcpw” on page 584

“Vmcuser” on page 586

“Vmlist” on page 598

“Vmvstortransport” on page 618

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 167
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
```

The default location of the restore is in the following directory:

c:\mnt\tsmvmbackup\my_vm_name\fullvm\
RESTORE_DATE_yyyy_mm_dd[hh_mm_ss].

- b. If errors occur in the restore processing, view the client error log for more information.

Tip: The error log is saved to the following file:

c:\Program Files\Tivoli\TSM\baclient\dsmerror.log

Related concepts:

“Parallel backups of virtual machines” on page 174

Related tasks:

“Preparing the environment for full backups of VMware virtual machines” on page 167

Related reference:

“Backup VM” on page 658

“Domain.vmfull” on page 381

“Query VM” on page 725

“Restore VM” on page 758

“Mode” on page 463

“Vmchost” on page 583

“Vmcpw” on page 584

“Vmcuser” on page 586

“Vmlist” on page 598

“Vmmc” on page 601

“Vmvstortransport” on page 618

Preparing the environment for file-level backups of VMware virtual machines

You can create file-level backups for VMware virtual machines that are running Microsoft Windows. Complete these steps to prepare the environment for file-level backups of virtual machines.

Procedure

1. Configure the vStorage backup server, by completing the following steps:
 - a. Install the backup-archive client with the custom installation option, including the **VMware vStorage API runtime files** component.
 - b. In the backup-archive client GUI, click **Edit > Client Preferences > VM Backup** and configure the following options in the dsm.opt file:

Option	Description
VMBACKUPTYPE FILE	Specifies that a file-level backup is created
VMCHOST	Specifies either the host name of each ESX server or the host name of the VMware Virtual Center. If you specify the VMware Virtual Center, you can create backups from any of the VMware servers that are managed by the VMware Virtual Center.
VMCPW	Specifies the password for the VMware Virtual Center or the ESX user ID that is specified with the vmcuser option.
VMCUSER	Specifies the user ID for the VMware Virtual Center or ESX user.

Option	Description
DOMAIN.VMFILE	Specifies the virtual machines to process for file-level backups.

- c. Optional: Define include and exclude options to control the files that are included in the backup.
- d. 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 data mover node can access the same disk volumes.
2. Prepare your VMware environment, by completing the following steps:
 - a. Ensure that the VMware Virtual Center and your vStorage backup server resolve host names to the same internet protocol (IP) addresses.
 - b. Optional: Install a backup-archive client within each virtual machine that you plan to use for file level restores. You do not need any special setup or options for the virtual machine that is installed by Tivoli Storage Manager.
3. Configure the Tivoli Storage Manager server for file-level VMware backups, by completing the following steps:
 - a. Register a single node for the vStorage backup server. For example:
register node DATAMOVER1 password.
 - b. For each virtual machine that is backed up, register a node on the Tivoli Storage Manager server. For best results, use the default client node name for the virtual machine. To find the node name, run the **hostname** command from the virtual machine. For example:

```
DSMADMC>register node DATAMOVER1 password
DSMADMC>register node VM1 password
DSMADMC>register node VM2 password
...
DSMADMC>register node VMn password
```

- c. Grant proxy access for each of the virtual machines you are backing up, by entering the following command:
grant proxy agent=*my_server_name* target=*my_guest_name*

For example:

```
grant proxy agent=DATAMOVER1 target=VM1
grant proxy agent=DATAMOVER1 target=VM2
...
grant proxy agent=DATAMOVER1 target=VMn
```

4. Configure include and exclude rules in the options file on the vStorage backup server to control the files that are included in the backup. The exclude list on the backup proxy can be set for all virtual machines or for a specific virtual machine. If incremental backups are created from within the virtual machine, the exclude list on the virtual machine must match the exclude list on the backup proxy. In most cases, the incremental backups run from the backup proxy.

The following examples demonstrate how to create exclude statements.

- To exclude files in the temp directory from the backup proxy for the single virtual machine vm1:
exclude "\\vm1\c\$\temp\...*"

- To exclude a file from backup proxy for all virtual machines:
exclude "*\X\$\Directory\File"
- When you use mount points without drive letter assignments, specify an include or exclude statement for that volume, for example:
exclude \\machine\3\$\dir1\...*.doc

Related tasks:

"Creating file-level backups for VMware virtual machines"

Related reference:

"Domain.vmfile" on page 378

"Query VM" on page 725

"Vmbackuptype" on page 581

"Vmchost" on page 583

"Vmcpw" on page 584

"Vmcuser" on page 586

"Vmlist" on page 598

"Vmvstortransport" on page 618

Creating file-level backups for VMware virtual machines

For VMware virtual machines that are running Microsoft Windows, you can use the backup-archive client to create file-level backups. You use include and exclude rules to determine the files that are included in the file-level backup.

About this task

The file-level backup of a virtual machine is implemented in the following sequence:

- A VMware snapshot is taken of the virtual machine to be backed up.
- The file systems of the virtual machine are remotely mapped to the vStorage backup server.
- A file-level progressive incremental backup is run for all of the file systems.
- The data is stored under the node name that matches the host name of the virtual machine. The data for each virtual machine is stored in the node that is associated with the virtual machine.
- The file system mappings are removed and the snapshot is removed.

Complete the following steps to create a file-level backup for a virtual machine that is running Windows.

Procedure

1. To configure the backup environment, complete the steps in the following topic:
"Preparing the environment for file-level backups of VMware virtual machines" on page 171
2. To start the file-level backup, open the command line of the vStorage backup server and run the **backup vm** command with the **vmbackuptype** option set to **file**. You can identify the virtual machine for backing up in several ways, for example:
 - To back up a list of virtual machines, run the **backup** command with the **vmlist** option:
dsmc backup vm -vmlist=vm1,vm2 -vmbackuptype=file

- To create file-level backups for all of the virtual machines defined in the DOMAIN.vmfile option, run the following command:
`dsmc backup vm -vmbackuptype=file`
 - To identify virtual machines by their internet protocol (IP) address, you must provide the Tivoli Storage Manager node name for each virtual machine. Use the following syntax in the command:
`dsmc backup vm 9.100.45.101[vm1], 9.100.45.102[vm2] -vmbackuptype=file`
3. Verify that you can restore the file-level backup for one of the virtual machines, by completing the following steps:
 - a. From the command line of the vStorage backup server, run the following command:
`dsmc restore \\my_vm_name\c$\my_docs\ c:\restore_temp\ -sub=yes
 -virtualnode=my_vm_name`
 - b. Verify that the files are restored for the virtual machine. By default, the files are restored in the following location:
`c:\mnt\tsmvmbackup\my_vm_name\fullvm\RESTORE_DATE_yyyy_mm_dd[hh_mm_ss]`
 - c. If the restore fails, view the error log for more information. The log file is saved to the following location:
`c:\Program Files\Tivoli\TSM\baclient\dsmerror.log`
 - d. For more information about restoring files from file-level backups, see the following topic:
 “Scenario: Restoring file-level VM backups” on page 219

Related concepts:

“Parallel backups of virtual machines”

Related tasks:

“Preparing the environment for file-level backups of VMware virtual machines” on page 171

Related reference:

“Domain.vmfile” on page 378

“Query VM” on page 725

“Vmbackuptype” on page 581

“Vmchost” on page 583

“Vmcpw” on page 584

“Vmouser” on page 586

“Vmlist” on page 598

“Vmvstortransport” on page 618

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 back ups 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 321. 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 `vm limitperhost` option limits how many virtual machines can be backed up, at any one time, per ESX server.
- The `vm limitperdatastore` option limits how many virtual machines can be backed up, at any one time, per datastore.

For backing up Hyper-V virtual machines, the `vmmaxparallel` option limits how many virtual machines can be backed up at any one time.

vmmaxparallel

The `vmmaxparallel` option specifies how many virtual machines can be backed up at any one time. The optimal value for `vmmaxparallel` 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 `vmmaxparallel`. 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 limitperhost

The `vm limitperhost` option specifies how many virtual machines can be backed up from a single ESX/ESXi host, at the same time. Like the value for `vmmaxparallel`, 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 limitperhost` 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 limitperdatastore

The `vm limitperdatastore` 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 176 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:

- `vmmaxparallel 3`
- `vm limitperhost 1`
- `vm limitperdatastore 1`

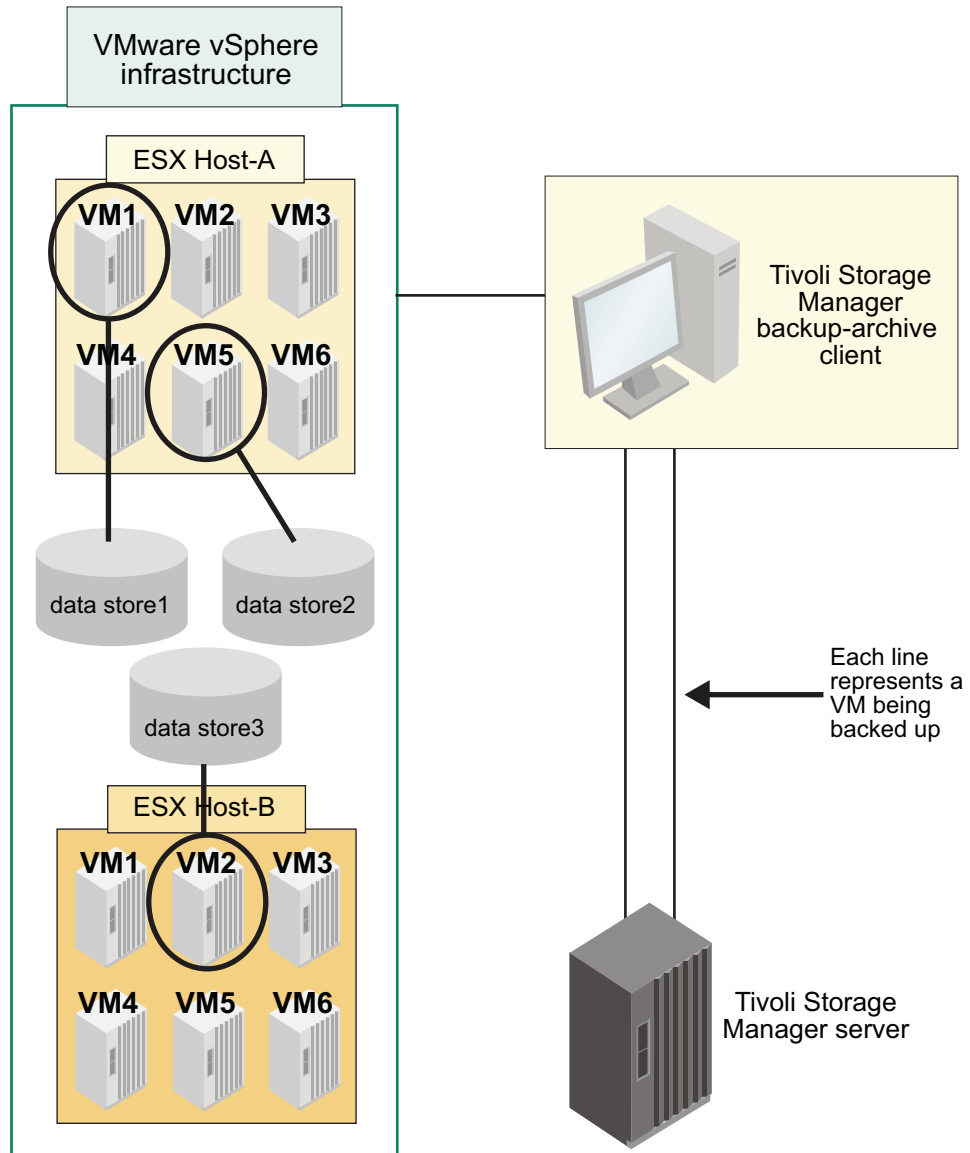


Figure 1. Virtual machines using unique datastores.

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 177 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

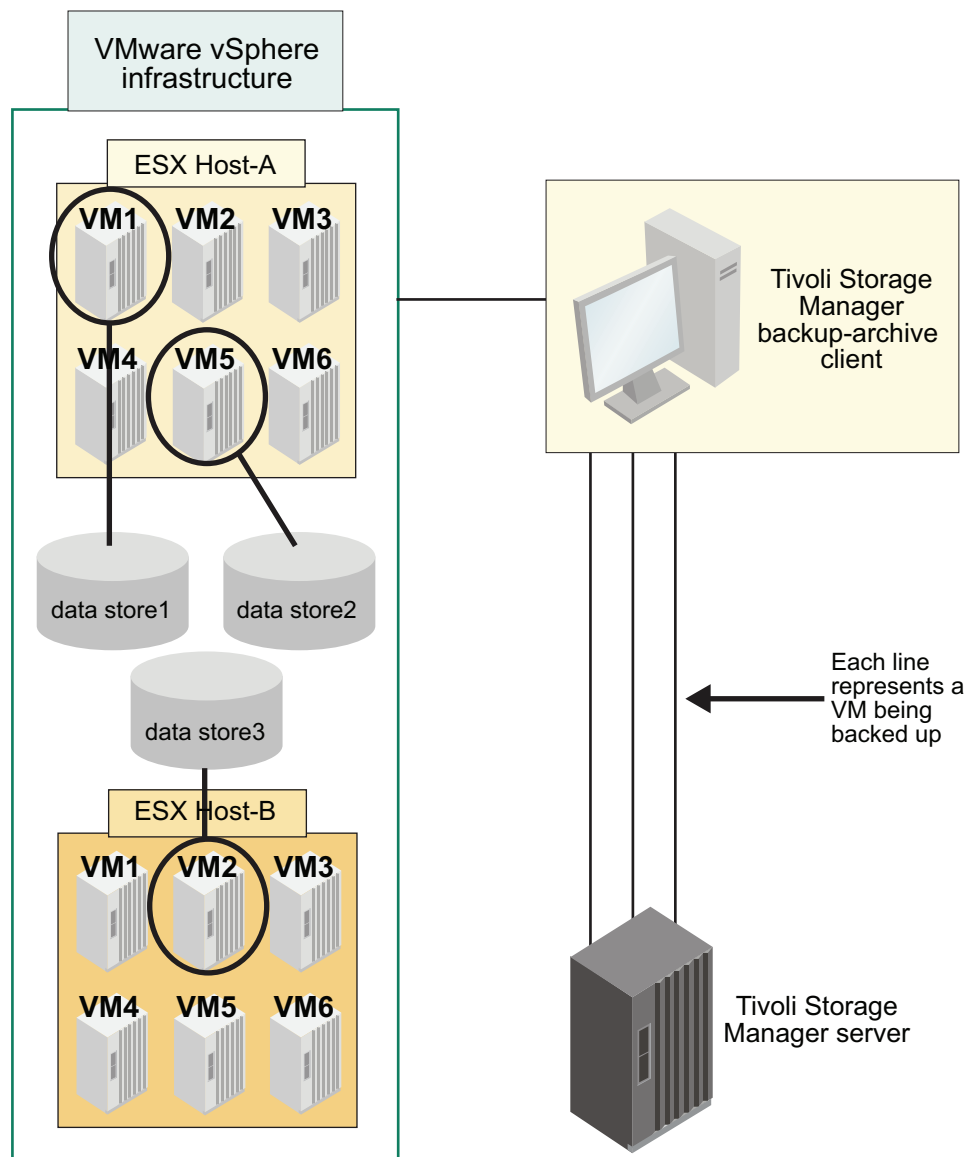


Figure 2. Virtual machines using unique datastores, with different option value for `vmlimitperhost`.

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 178 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.

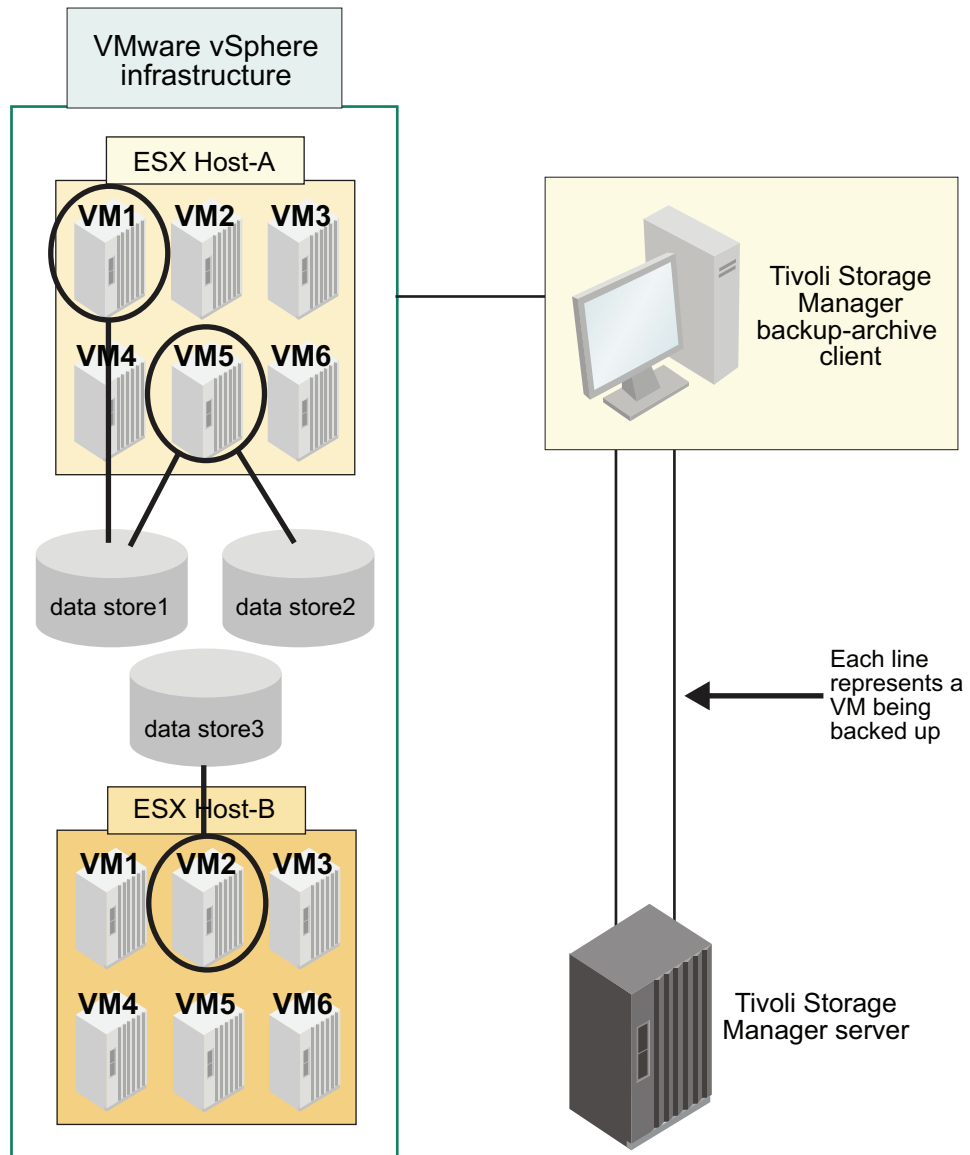


Figure 3. Virtual machines sharing a datastore.

Back up virtual machines on a Hyper-V system

You can use the backup-archive client to backup virtual machines that are managed by a Microsoft Hyper-V server.

Important: The following information applies only to Hyper-V backup and restore operations for Tivoli Storage Manager configurations that do not have a license file for the Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V product installed. If you have that license installed, use the Hyper-V protection procedures that are documented at IBM Tivoli Storage Manager for Virtual Environments, Data Protection for Microsoft Hyper-V.

You can use the backup-archive client to perform the following tasks on Hyper-V systems:

- Back up and restore snapshots of Hyper-V virtual machines by using the Microsoft Volume Shadow Copy Services (VSS) interface. You can restore a virtual machine from these files, which are created by VSS.
- Restore either individual virtual machines or a group of virtual machines that are running on a Hyper-V server.
- Back up data without stopping the virtual machine or stopping any running applications within the virtual machine.

Restriction: You can back up data this way only if the operating system of the virtual machine supports VSS. For more information about this restriction, see the "Planning for backup" article in the Hyper-V section of the Microsoft TechNet support site, which is at <https://technet.microsoft.com>.

- Add backup and restore support for guest operating systems without installing a backup-archive client on the guest VM.
- Back up and restore any guest operating systems that are hosted by the Hyper-V server on remote shares, regardless of whether the guest operating system is supported by Tivoli Storage Manager.

For both disaster recovery and long-term data backup support, perform the following steps:

1. Install the backup-archive client on the physical Hyper-V host server in the parent partition. The server must use Windows Server 2008 or Windows Server 2012, or newer operating system.
2. On the Tivoli Storage Manager server, define a single node for the Hyper-V system.
3. Back up a running virtual machine from the backup-archive client GUI or from the command line:
 - From **Backup** window in the backup-archive client GUI, expand the Hyper-V VMs container, and select the virtual machine to back up.
 - You can also enter `dsmc backup vm -vmlist=vm1 -vmbackuptype=HyperVFull` at a Tivoli Storage Manager command prompt, where **vm1** is the name of the virtual machine.

Important:

- a. If the virtual machine is not fully VSS-compliant (the operating system supports VSS, and all virtual hard disks (VHD) are configured to support VSS):
 - The virtual machine that is running pauses
 - A snapshot is taken
 - The virtual machine resumes
4. Restore a virtual machine:
 - From the **Restore** window in the backup-archive client GUI, expand the Hyper-V VMs container, and select the virtual machine to restore.
 - You can also enter: `dsmc restore vm vm1 -vmbackuptype=HyperVFull` at a client command prompt, where **vm1** is the name of the virtual machine.

Important: When you restore a virtual machine, the existing virtual machine is stopped, and all files that comprise the existing virtual machine (for example, VHD files) are deleted. If the virtual machine is a member of a Windows Server 2012 (or newer version) cluster, the virtual machine is taken offline from the

cluster, which stops the virtual machine. Then, the files are deleted and the virtual machine is restored from the Tivoli Storage Manager backup. When the backup snapshot is restored, the virtual machine is re-created, including any Hyper-V snapshots that existed when the files were backed up.

To schedule backups of Hyper-V virtual machines, create a schedule on the server to run the backups. When you define the schedule, specify ACTION=BACKUP, SUBACTION=VM, and add the -VMBACKUPTYPE=HYPERVFULL option to the **OPTIONS** parameter. The following shows a sample schedule definition for backing up Hyper-V virtual machines:

```
define sched standard HyperV-Backup type=client description="backup hyperv vms"  
    action=backup subaction=vm options="-vmbackuptype=hypervfull"
```

On the client system, add the virtual machines to be processed, when the scheduled backup is run, to the **VMLIST** option in the dsm.opt file. For example,
VMLIST VM1,VM2,VM99

Related reference:

“Backup VM” on page 658

“Query VM” on page 725

“Restore VM” on page 758

Hyper-V backup support limitations

Because of the tight integration of Microsoft Failover Clustering with Cluster Shared Volumes and Hyper-V, the following limitations apply when you are using the backup-archive client.

Important: The following limitations apply only to Hyper-V backup and restore operations for Tivoli Storage Manager configurations that do not have a license file for the Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V product installed. If you have that product and license installed, see the description of limitations at IBM Tivoli Storage Manager for Virtual Environments, Data Protection for Microsoft Hyper-V backup support limitations.

Microsoft Hyper-V support is provided for the following configurations:

- A dedicated Hyper-V system
- A Hyper-V configuration that is part of Microsoft Failover Clustering
- A Hyper-V configuration that is part of Microsoft Failover Clustering with Cluster Shared Volumes

For Windows Server 2008, virtual machine backups can be performed only by a backup-archive client that is installed on the cluster node that is the current owner of the virtual machine. To see which virtual machines are owned by the node, select **Server Manager > Roles > Hyper-V Manager**. This restriction does not exist for Windows Server 2012 hosts.

For Windows Server 2008, a virtual machine should be restored only by a backup-archive client that is installed on the cluster node that is the current owner of the virtual machine. To determine the virtual machine configurations of the cluster, select **Server Manager > Features > Failover Cluster Manager > cluster name > Services and Applications..** Use the **Services and Applications** window to change which node is the owner of the virtual machine before you restore it. This restriction does not exist for Windows Server 2012 hosts.

For Windows Server 2008, virtual machines that are currently owned by different nodes, and that have files on a common Cluster Shared Volume, cannot be simultaneously backed up, using the default system VSS provider. The second backup fails because the shared volume is owned by the first node. You can significantly reduce the amount of time that the first node owns the shared volume by using hardware snapshots through the VSS provider that is supplied by your storage hardware vendor. After the hardware snapshot is completed, the backup process can be initiated from the second node, therefore the backups can be running concurrently from both nodes. To determine which node owns the volume, select **ServerManager > Features > Failover Cluster Manager > *cluster name* > Cluster Shared Volumes**. The **Cluster Shared Volume** windows displays the status of the volume, including ownership, and if a backup is in progress (with redirected I/O enabled). This restriction does not exist for Windows Server 2012 hosts.

If your cluster environment typically has all of the virtual machines owned by a single node, and other nodes provide failover support only, the following applies:

- Consider setting CLUSTERNODE YES for backups from the active node and restore to failover nodes (if needed).
- Consider running the client scheduler as a failover service to periodically back up all virtual machines. For example: `dsmc backup vm *`
`-vmbackuptype=hypervfull`. A failover starts the scheduler on the failover node.

If your cluster environment typically has some virtual machines that are owned by one node, and other virtual machines that are owned by other nodes, the following applies:

- Consider using the “Asnodename” on page 324 option to allow for backup from multiple nodes and restore to failover nodes (if needed).
- Consider running the client scheduler as active on each node to periodically back up the virtual machines that are owned by each node. The schedules must be defined with non-overlapping start times to avoid contention for common shared volumes. If the scheduler on each node backs up with `-vmlist=*`, there is no need for the scheduler service to failover since the scheduler service on the other node will automatically pick up nodes that move over on the next backup.

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 8

“Configuring the client to back up and archive Tivoli Storage Manager FastBack data” on page 67

Related reference:

“Fbclientname” on page 407

“Fbpolicyname” on page 409

“Fbvolumename” on page 412

Backing up Net Appliance CIFS share definitions

Network Appliance (NetApp) CIFS share definitions include share permissions that are set on the file server.

About this task

The Windows client backs up the CIFS share definition under the root directory, the mapped CIFS share, or the UNC name. This support requires that the Net Appliance file server is running DATA ONTAP software, which presents CIFS shares to remote clients as ordinary remote NTFS shares.

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

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

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

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

Related concepts:

“Restore Net Appliance CIFS shares” on page 211

Related reference:

“Snapdiff” on page 529

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 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 23 lists some informational messages and meanings.

Table 23. Client command line informational messages

Informational message	Meaning
Directory-->	Indicates the directory that you back up.
Updating-->	Indicates that only the file meta data is sent, not the data itself.
Expiring-->	Indicates an object (file or directory) on the server that no longer exists on the client is expired and made inactive on the server.

Table 23. Client command line informational messages (continued)

Informational message	Meaning
Subfile-->	Indicates that a file is processed using adaptive subfile backup. The file size that is indicated in this message is the size of the entire file, not the amount of data sent.
Total number of objects inspected:	<p>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.</p> <p>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.</p>
Total number of objects backed up:	As indicated.
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 275 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 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 inspected:	Specifies the sum of the sizes of the files that are selected for the operation. For example, the total number of bytes inspected for this command is the number of bytes used in the directory C:\Users <code>dsmc.exe INCREMENTAL C:\Users* -su=yes</code>
Total number of bytes processed:	Specifies the sum of the sizes of the files that are processed for the operation.
Total number of subfile objects:	The number of files processed using adaptive subfile backup.
Data transfer time:	The total time to transfer data across the network. Transfer statistics might not match the file statistics if the operation was retried due to a communications failure or session loss. The transfer statistics display the bytes attempted to be transferred across all command attempts.

Table 23. Client command line informational messages (continued)

Informational message	Meaning
Network data transfer rate:	The average rate at which the network transfers data between the client and the server. This is calculated by dividing the total number of bytes transferred by the time to transfer the data over the network. The time it takes to process objects is not included in the network transfer rate. Therefore, the network transfer rate is higher than the aggregate transfer rate.
Aggregate data transfer rate:	<p>The average rate at which Tivoli Storage Manager and the network transfer data between the client and the server. This is calculated by dividing the total number of bytes transferred by the time that elapses from the beginning to the end of the process. Both Tivoli Storage Manager processing and network time are included in the aggregate transfer rate. Therefore, the aggregate transfer rate is lower than the network transfer rate.</p> <p>Note: On occasion, the aggregate data transfer rate might be higher than the network data transfer rate. This is because the backup-archive client can have multiple simultaneous sessions with the backup server. If you set the resourceutilization option, the client attempts to improve performance and load balancing by using multiple sessions when it backs up a volume or other set of files. When multiple sessions are open during backup, the data transfer time represents the sum of the times reported by all sessions. In this case, aggregate data transfer time is incorrectly reported as higher. However, when running with a single session, the aggregate data transfer rate should always be reported as lower than the network data transfer rate.</p>
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 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: $(1 - 10/100) \times 100 = 90\%$
Subfile objects reduced by:	<p>Expresses the savings achieved by adaptive subfile backup as a percentage of actual bytes sent versus actual file sizes. For example, if the actual sizes of the files backed up is 100 MB, but with adaptive subfile backup, only 25 MB were sent, then the reduction would be: $(1 - (25 \text{ MB}/100 \text{ MB})) \times 100 == 75\%$.</p> <p>Note: The calculation considers all files that were processed, not only those files that were processed by adaptive subfile backup. The calculation does not include the effect of compression.</p>
Elapsed processing time:	The active processing time to complete a command. This is calculated by subtracting the starting time of a command process from the ending time of the completed command process.
Total number of bytes transferred:	As indicated.
LanFree bytes transferred:	The total number of data bytes transferred during a lan-free operation. If the enablelanfree option is set to <i>no</i> , this line will not appear.
Total number of bytes inspected:	A sum of sizes of files selected for the operation.

Table 23. Client command line informational messages (continued)

Informational message	Meaning
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 (Windows): Additional considerations

This section discusses additional information to consider when backing up data.

Open files

Some files on your system might be in use when you try to back them up. These are called *open files* because they are locked by an application for its exclusive use.

It is not very common for files to be opened in locked mode. An application can open a file in this way to avoid other applications or users from reading or accessing the file, but it can prevent backup programs from reading the file for backup.

You might not always want to use the open file feature to back up open or locked files. Sometimes an application opens a file or group of files in this locked mode to prevent the access of these files in an inconsistent state.

To avoid the increase of processor usage when you create a volume snapshot for each backup, and on platforms where the open file feature is not available or is not in use, consider the following points:

- If the file is unimportant or can be easily rebuilt (a temporary file for example), you might not care if the file is backed up, and might choose to exclude it.
- If the file is important:
 - Ensure the file is closed before backing it up. If backups are run according to a schedule, use the `preschedulecmd` option to enter a command that closes the file. For example, if the open file is a database, issue a command to close the database. You can use the `postschedulecmd` option to restart the application that uses the file after the backup completes. If you are not using a schedule for the backup, close the application that uses the file before you start the backup.
 - The client can back up the file even if it is open and changes during the backup. This is only useful if the file is usable even if it changes during backup. To back up these files, assign a management class with *dynamic* or *shared dynamic* serialization.

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

Related concepts:

“Display information about management classes and copy groups” on page 269
 “Select a management class for files” on page 273

Ambiguous file space names in file specifications

If you have two or more file spaces such that one file space name is the same as the beginning of another file space name, then an ambiguity exists when you try to restore, retrieve, query, or do another operation that requires the file space name as part of the file specification.

For example, consider the following file spaces and the backup copies they contain:

File space name	File name
\\storman\home	amr\project1.doc
\\storman\home\amr	project2.doc

Notice that the name of the first file space, \\storman\home, matches the beginning of the name of the second file space, \\storman\home\amr. When you use the backup-archive command-line client interface to restore or query a file from either of these file spaces, by default the client matches the longest file space name in the file specification, \\storman\home\amr. To work with files in the file space with the shorter name, \\storman\home, use braces around the file space name portion of the file specification.

This means that the following query command finds project2.doc but does not find project1.doc:

```
dsmc query backup "\\storman\home\amr\*"
```

This is because the longer of the two file space names is \\storman\home\amr and that file space contains the backup for project2.doc.

To find project1.doc, enclose the file space name in braces. The following command finds project1.doc but does not find project2.doc:

```
dsmc query backup "{\\storman\home}\amr\*"
```

Similarly, the following command restores project1.doc but does not restore project2.doc:

```
dsmc restore {\\storman\home}\amr\project1.doc
```

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

Select **Utilities** → **View Policy Information** from the backup-archive client or web client GUI to view the backup policies defined by the Tivoli Storage Manager server for your client node.

Related concepts:

Chapter 9, "Storage management policies," on page 267

Related tasks:

“Setting the client scheduler process to run as a background task and start automatically at startup” on page 252

Deleted file systems

When a file system or drive has been deleted, or it is no longer backed up by Tivoli Storage Manager, the existing backup versions for each file are managed according to the following policy attributes: Number of days to keep inactive backup versions, and number of days to keep the last backup version (if there is no active version).

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

You can also use the **delete backup** command to delete individual backup versions, or the **delete filespace** command to delete the entire file space. Your Tivoli Storage Manager server administrator must give you "delete backup" authority to use these commands. If the file space also contains archive versions, you must also have delete archive authority to use **delete filespace**.

Use the **query session** command to determine whether you have delete backup and delete archive authority. Alternatively, you can ask your Tivoli Storage Manager server administrator to delete the file space for you.

When you delete a file system, it has no effect on existing archive versions. However, if you no longer require the archive versions, you can use the **delete archive** or **delete filespace** commands to delete archives.

Related concepts:

Chapter 9, “Storage management policies,” on page 267

Removable media backup

The backup-archive client backs up your removable media (such as tapes, cartridges or diskettes) based on the drive label, not the drive letter.

If a drive has no label, the backup does not occur. This use of drive labels permits you to perform such tasks as backing up different diskettes from the a: drive.

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

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

If the label of the d: drive does not match a file space name on the server, Tivoli Storage Manager cannot locate your files using the current d: drive label. However, the client can locate your files if you use the file space name based on the original drive label. A mismatch between a label and a file space name might

occur if you label your drives again, or if you access Tivoli Storage Manager from a different workstation than the one from which you backed up the files. If you have not relabeled the drive, and you are at the same workstation where the file was backed up, then you can use the drive letter as a shorthand version of the file space name (drive label).

Fixed drives

The backup-archive client can back up your fixed drives even if they do not have a label, including drive aliases created with the DOS **subst** command. This applies to both the drive alias and the underlying physical drive, because the alias name and the physical drive label are the same.

NTFS and ReFS file spaces

When you back up files on NTFS or ReFS partitions, the client also backs up file security information and file descriptors.

The following file descriptors are backed up:

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

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

Universal Naming Convention names

A Universal Naming Convention (UNC) name is a network resource name for a share point on a workstation.

The resource name includes the workstation name assigned to the workstation and a name you assign to a drive or directory so that it can be shared. The name you assign is also called a *share point name*.

Examples: UNC names in domain lists

This topic shows some examples of using UNC names to specify a domain list.

About this task

You must specify the following information:

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

Example 1: To specify drive a: containing removable media, enter
domain a: \\local\c\$

Example 2: To specify fixed drive c:, enter
domain c: \\remote\share1 \\remote\c\$

Examples: UNC name backup

You can back up shared files in a network through the use of a UNC name. Some examples of backing up UNC name files are shown.

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

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

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

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

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

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

```
\\ocean\wave
```

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

See the following table for examples showing selective backup of files using UNC names. In these examples, assume that:

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

Table 24. UNC examples

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

You can also specify UNC names for files in your include-exclude and domain lists.

Related tasks:

“Creating an include-exclude list” on page 92

Related reference:

“Domain” on page 371

Microsoft Dfs file protection methods

There are some methods that you can use to protect the data in your Microsoft Dfs environment.

About this task

Here are the methods you should use to protect your Microsoft Dfs data:

Procedure

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

Results

Note:

1. See the product README file for current limitations of this feature.

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

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

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

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

You can use the `dfsbackupmntpnt` option to specify whether the client sees a Dfs mount point as a Microsoft Dfs junction or as a directory.

Important: Restore the Dfs junction metadata first. This recreates the links. Then restore each junction and the data at each junction separately. If you do not restore the junction metadata first, the client creates a directory under the Dfs root using the same name as the junction point and restores the data in that directory.

The following example relates to method 1 above and illustrates how to use the client to back up and restore a Microsoft Dfs environment. Assume the existence of a domain Dfs environment hosted by the workstation `wkst1`:

Dfs root

\\wkst1\abc64test

Dfs link1

\\wkst1\abc64test\tools

Dfs link2

\\wkst1\abc64test\trees

Backup procedure:

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

```
dsmc inc \\wkst1\abc64test
```

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

```
dsmc inc \\wkst1\abc64test\tools
```

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

```
dsmc inc \\wkst1\abc64test\trees
```

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

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

Restore procedure:

1. Manually recreate shares at target workstations only if they no longer exist.
2. Manually recreate the Dfs root using the exact name as it existed at the time of back up.
3. Enter the following command to recover data from the tools link. This step is not necessary if the data still exists at the link target:

```
dsmc restore \\wkst1\abc64test\tools\* -sub=yes
```
4. Enter the following command to recover data from the trees link. This step is not necessary if the data still exists at the link target:

```
dsmc restore \\wkst1\abc64test\trees\* -sub=yes
```
5. Use the Distributed File System management console snap-in to reestablish replication for each link, if necessary.

The following limitations exist for restoring Microsoft Dfs data:

- The client does not restore root of Dfs. To recreate the Dfs tree, manually create the Dfs root first, then start restore to recreate the links.
- The client can back up the Dfs tree (both domain based Dfs and stand alone Dfs) hosted on local workstation only. You cannot back up Dfs if the Dfs host server is not your local workstation.
- The client cannot recreate shared folders on restore. For example, if you delete the junction and the shared folder the junction points to, restoring the Dfs root recreates the Dfs junction, but restoring a junction creates a local folder instead of creating the original backed up shared network folder.
- If a Dfs link is created with replica and the replica share is on a different server, the client does not display the replica data.

- If a Dfs root is added or modified, the client will not back it up. You must specify the Dfs root in the `domain` option in the client options file (`dsm.opt`) regardless of whether `DOMAIN ALL-LOCAL` is specified.

Chapter 5. Restoring your data

Use 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.

Note: When you restore a directory, its modification date and time is set to the date and time of the restore operation, and not to the date and time the directory had when it was backed up. This is because Tivoli Storage Manager restores the directories first, then adds the files to the directories.

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.

The following are the primary restore tasks:

- “Restoring files and directories” on page 195
- “Restoring Windows system state” on page 199
- “Restoring Automated System Recovery files” on page 201
- “Microsoft Dfs tree and file restore” on page 203
- “Restoring an image” on page 203
- “Restore data from a backup set” on page 206
- “Restoring data to a point in time” on page 233
- “Restore NAS file systems” on page 235
- “Authorizing another user to restore or retrieve your files” on page 229
- “Restoring or retrieving files from another client node” on page 231
- “Restoring or retrieving your files to another workstation” on page 232
- “Deleting file spaces” on page 232
- “Restoring data from a VMware backup” on page 212

Related tasks:

“Starting a web client session” on page 115

Duplicate file names

If you attempt to restore or retrieve a file whose name is the same as the short name of an existing file, a file name collision occurs (existence of duplicate file names).

An example is when the file *abcdefghijkl.doc* has a short name of *abcdef~1.doc*, and you attempt to restore or retrieve a file explicitly named *abcdef~1.doc* into the same directory. In this case, a collision occurs because the name of the file you are restoring conflicts with the short name for *abcdefghijkl.doc*.

A collision can occur even if the files are restored or retrieved to an empty directory. For example, files *abcdef~1.doc* and *abcdefghijkl.doc* might originally have

existed in the directory as *abcdefghijkl.doc* and *abcdef~2.doc*. During the restore, if *abcdefghijkl.doc* is restored first, it is assigned a short name of *abcdef~1.doc* by the Windows operating system. When you restore *abcdef~1.doc*, the duplicate file name situation occurs.

Tivoli Storage Manager handles these situations based on the value of the `replace` option. Use the `replace` option to specify whether to overwrite an existing file, or to prompt you for your selection when you restore or retrieve files.

If a file name collision occurs, you can do any of the following:

- Restore or retrieve the file with the short file name to a different location.
- Stop the restore or retrieve and change the name of the existing file.
- Disable short file name support on Windows.
- Do not use file names, such as *abcdef~1.doc*, that would conflict with the short file naming convention.

Related reference:

"Replace" on page 499

Universal Naming Convention names restore

Using a Universal Naming Convention (UNC) name permits you to restore specific shared files to a separate file space. This is useful if, for example, you or an administrator want to restore a portion of data that you would otherwise be unable to access.

Except for drives with removable media, every local drive letter is accessible using a local UNC name that includes the workstation name and a designation for the drive letter. For example, to enter a UNC name on drive `c:` for workstation `ocean`, enter:

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

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

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

```
\\ocean\wave
```

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

Active or inactive backup restore

Your administrator determines how many backup versions Tivoli Storage Manager maintains for each file on your workstation. Having multiple versions of a file permits you to restore older versions if the most recent backup is damaged.

The most recent backup version is the *active* version. Any other backup version is an *inactive* version. Every time Tivoli Storage Manager backs up your files, it marks the new backup version as the active backup, and the last active backup becomes an inactive backup. When the maximum number of inactive versions is reached, Tivoli Storage Manager deletes the oldest inactive version.

To restore a backup version that is inactive, you must display both active and inactive versions by clicking on the **View** menu → **Display active/inactive files** item. To display only the active versions (the default), click on the **View** menu →

Display active files only item. If you try to restore both an active and inactive version of a file at the same time, only the active version is restored.

On the Tivoli Storage Manager command line, use the `inactive` option to display both active and inactive objects.

Related reference:

“Inactive” on page 428

Restoring files and directories

You can locate the files you want to restore by searching and filtering.

Filtering displays only the files that match the filter criteria for your restore operation. Files that do not match the filter criteria do not display. The filter process searches files in the specified directory but does not include subdirectories.

Restoring data by using the GUI

You can use the client GUI to restore files and directories.

About this task

Restriction: The web client GUI cannot browse network resources for a restore operation. No shares are listed if you expand the **Network** branch. You can restore to a network resource from the web client if the entire file is processed. Specify the shared file system in the domain option in the `dsm.opt` options file. For example, domain `all-local \\server\share`. To complete the restore operation, specify **Network Share** in the **Restore Destination** dialog. This processes all file systems that are specified by the domain option. Alternatively, you can use the GUI Client to complete the restore operation.

Procedure

1. Click **Restore** on the main window. The Restore window appears.
2. Expand the directory tree by clicking the plus (+) sign or the folder icon next to an object in the tree. Select the object that you want to restore. To search or filter files, click the **Search** icon from the toolbar.
3. Click the selection box for the objects that you want to restore.
4. To modify specific restore options, click the **Options** button. Any options that you change are effective during the current session only.
5. Click **Restore**. The Restore Destination window appears. Enter the appropriate information.
6. Click **Restore**. The Restore Task List window displays the processing status.

Related tasks:

“Backing up data using the GUI” on page 125

Examples for restoring data using the command line

You can use the examples in this topic when you need to restore objects from Tivoli Storage Managerserver storage.

The following table shows how to use some restore commands to restore your objects from Tivoli Storage Manager server storage.

Table 25. Command-line restore examples

Task	Command	Considerations
Restore the most recent backup version of the c:\doc\h1.doc file, even if the backup is inactive.	<code>dsmc restore c:\doc\h1.doc -latest</code>	If the file you are restoring no longer resides on your workstation, and you have run an incremental backup since deleting the file, there is no active backup of the file on the server. In this case, use the 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 456 for more information.
Display a list of active and inactive backup versions of files from which you can select versions to restore.	<code>dsmc restore c:\project* -pick -inactive</code>	If you try to restore both an active and inactive version of a file at the same time, only the active version is restored. See "Pick" on page 481 and "Inactive" on page 428 for more information.
Restore all files with a file extension of .c from the c:\devel\projecta directory.	<code>dsmc restore c:\devel\projecta*.c</code>	If you do not specify a destination, the files are restored to their original location.
Restore the c:\project\doc\h1.doc file to its original directory.	<code>dsmc restore c:\project\doc\h1.doc</code>	If you do not specify a destination, the files are restored to their original location.
Restore the c:\project\doc\h1.doc file under a new name and directory.	<code>dsmc restore c:\project\doc\h1.doc c:\project\newdoc\h2.doc</code>	None
Restore the files in the e: drive and all of its subdirectories.	<code>dsmc restore e:\ -subdir=yes</code>	You must use the subdir option to restore directory attributes/permissions. See "Subdir" on page 548 for more information about the subdir option.
Restore all files in the c:\mydir directory to their state as of 1:00 PM on August 17, 2002.	<code>dsmc restore -pitd=8/17/2002 -pitt=13:00:00 c:\mydir\</code>	See "Pitdate" on page 482 and "Pittime" on page 483 for more information about the pitdate and pittime options.
Restore the c:\doc\h2.doc file to its original directory on the workstation, named star.	<p><code>dsmc restore c:\doc\h2.doc \\star\c\$\</code></p> <p>To restore the file to "star" which has been renamed "meteor", enter:</p> <p><code>dsmc restore \\star\c\$\doc\h2.doc \\meteor\c\$\</code></p> <p>You could also enter:</p> <p><code>dsmc restore \\star\c\$\doc\h2.doc c:\</code></p> <p>This example is valid because if the workstation name is not included in the specification, the local workstation is assumed ("meteor", in this case).</p>	For the purposes of this manual, the workstation name is part of the file name. Therefore, if you back up files on one workstation and you want to restore them to another workstation, you must specify a destination. This is true even if you are restoring to the same physical workstation, but the workstation has a new name.

Table 25. Command-line restore examples (continued)

Task	Command	Considerations
Restore a file that was originally backed up from the diskette labeled “workathome” in the a: drive, and restore it to a diskette in the a: drive labeled “extra”.	<code>dsmc restore {workathome}\doc\h2.doc a:\doc\h2.doc</code>	If you are restoring a file to a disk with a different label than the disk from which the file was backed up, you must use the file space name (label) of the backup disk instead of the drive letter.
Restore files specified in the c:\filelist.txt file to the d:\dir directory.	<code>dsmc restore -filelist=c:\filelist.txt d:\dir\</code>	See “Filelist” on page 414 for more information about restoring a list of files.
Restore all members of the virtfs\group1 group backup stored on the Tivoli Storage Manager server.	<code>dsmc restore group {virtfs}\group1</code>	See “Restore Group” on page 746 for more information.

Related concepts:

Chapter 12, “Using commands,” on page 625

Related reference:

“Restore” on page 730

Examples: Restoring large amounts of data

If you need to restore a large number of files, you get faster performance using the command line interface rather than the GUI interface. In addition, you improve performance if you enter multiple **restore** commands at one time.

About this task

For example, to restore all the files in your c: file space, enter:

```
dsmc restore c:\* -subdir=yes -replace=all -tapeprompt=no
```

However, if you enter multiple commands for the root directories in your c: file space, you can restore the files faster. For example, enter these commands:

```
dsmc restore c:\users\ -subdir=yes -replace=all -tapeprompt=no
dsmc restore c:\data1\ -subdir=yes -replace=all -tapeprompt=no
dsmc restore c:\data2\ -subdir=yes -replace=all -tapeprompt=no
```

Or, if you need to restore files for multiple drives, enter these commands:

```
dsmc restore c:\* -subdir=yes -replace=all -tapeprompt=no
dsmc restore d:\* -subdir=yes -replace=all -tapeprompt=no
dsmc restore e:\* -subdir=yes -replace=all -tapeprompt=no
```

You can also use the quiet option with the **restore** command to save processing time. However, you will not receive informational messages for individual files.

Note: If you already have the appropriate values set for the `subdir`, `replace`, `tapeprompt`, and `quiet` options in your client options file, it is not necessary to include these options in the commands.

When you enter multiple commands to restore your files, you must specify a unique part of the file space in each **restore** command. Do not use any overlapping file specifications in the commands.

To display a list of the root directories in a file space, use the **query backup** command. For example:

```
dsmc query backup -dironly -subdir=no c:\
```

As a general rule, you can enter two to four **restore** commands at one time. The maximum number you can run at one time without degrading performance depends on factors such as network utilization and how much memory you have. For example, if \users and \data1 are on the same tape, the restore for \data1 must wait until the restore for \users is complete. However, if \data2 is on a different tape, and there are at least two tape drives available, the restore for \data2 can begin at the same time as the restore for \users.

The speed at which you can restore the files also depends upon how many tape drives are available and whether your administrator is using collocation to keep file spaces assigned to as few volumes as possible. If your administrator is using collocation, the number of sequential access media mounts required for restore operations is also reduced.

Standard query restore, no-query restore, and restartable restore

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 **MAXNUMP** 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**, or **todate**, the client uses a *no-query restore* method for restoring files and

directories from the server. This method is called no-query restore because instead of querying the server for each object to be restored, a single restore request is sent to the server. In this case, the server returns the files and directories to the client without further action by the client. The client merely accepts the data that comes from the server and restores it to the destination named on the **restore** command.

Using the command-line client, an example of an unrestricted wildcard command would be:

```
c:\mydocs\2004\*
```

An example of a restricted wildcard file specification would be:

```
c:\mydocs\2004\sales.*
```

Restartable restore process:

If the restore process stops because of a power outage or network failure, the server records the point at which this occurred.

This record is known to the client as a *restartable restore*. It is possible to have more than one restartable restore session. Use the **query restore** command 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. 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.

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 508

“Restore” on page 730

Restoring Windows system state

The Microsoft Volume Shadowcopy Service (VSS) is supported on Windows backup-archive clients. The client uses VSS to restore the system state.

About this task

By default, all system state components are restored together.

Extra disk space is required to restore SFP files because they are restored to a temporary location on your disk.

Startable system state components can include the following:

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

System services components can include the following:

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

Important: Only restore system state during system recovery.

You must be a member of the Administrators group to restore Windows system state information.

On Windows clients, **restore systemstate** also restores ASR files when running in WinPE (Windows Preinstallation Environment).

To restore Windows system state information using the GUI, follow these steps:

Procedure

1. Click **Restore** from the GUI main window. The Restore window appears.
2. Expand the directory tree by clicking the plus sign +. To display files in a folder, click the folder icon.
3. Locate the system state node in the directory tree. You can expand the system state node to display the components.
4. Click the selection box next to the system state node to restore the entire system state. You can restore the startable system state node only as a single entity because of dependencies among the startable system state components. By default, all components are selected; you can restore individual system services components.
5. Click **Restore**. The Task List window displays the restore processing status.

Results

On the command line, use the **restore systemstate** command to restore a backup of system state.

Considerations:

- Your Windows client must be connected to a Tivoli Storage Manager Version 5.2.0 or higher server.
- If Active Directory is installed, you must be in Active Directory restore mode.
- See the related information for performing a system recovery to perform the following tasks:
 - Your operating system is still functioning, but a complete system restore is required.
 - A complete recovery is required, including an operating system re-installation.

Related concepts:

“Recovering a computer when the Windows OS is not working” on page 202

Related tasks:

“Restoring the operating system when the computer is working”

Related reference:

“Restore Systemstate” on page 753

Restoring Automated System Recovery files

You can restore Automated System Recovery (ASR) files to recover the Windows operating system volume configuration information and system state if a catastrophic system or hardware failure occurs.

Before you begin

You must be a member of the Administrators or Backup Operators group to back up and restore ASR files.

About this task

The backup-archive client restores ASR data when the backup-archive client restores the Windows system state.

Procedure

To restore ASR files on Windows operating systems, use the **restore systemstate** command.

Related concepts:

“Recovering a computer when the Windows OS is not working” on page 202

Restoring the operating system when the computer is working

If your computer is working, you can restore the operating system from backed up files.

About this task

If Active Directory is installed, you must be in Active Directory restore mode. When performing an operating system recovery including the system state, use the following restore order. Do not restart the computer between each step, even though you are prompted to do so.

Procedure

1. Restore the system drive. For example: `dsmc restore c:* -sub=yes -rep=all`.
2. Restore system state. For example: `dsmc restore systemstate`.

Recovering a computer when the Windows OS is not working

If the computer has a catastrophic hardware or software failure, you can recover a Windows operating system with Automated System Recovery (ASR).

Related tasks:

“Restoring the operating system when the computer is working” on page 201

Creating a bootable WinPE CD

Before you can recover a Windows computer by using Automated System Recovery (ASR), you must create a bootable Windows Preinstallation Environment (WinPE) CD or DVD.

Procedure

For instructions that describe how to create a bootable WinPE CD or DVD, see the following Tivoli Storage Manager Wiki articles:

For Windows Server 2008, Windows Server 2008 R2, and Windows 7

See *Best Practices for Recovering Windows Server 2008, Windows Server 2008 R2, Windows 7, and Windows Vista* at [https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli Storage Manager/page/Best Practices for Recovering Windows Server 2008, Windows Server 2008 R2, Windows 7, and Windows Vista](https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli%20Storage%20Manager/page/Best%20Practices%20for%20Recovering%20Windows%20Server%202008,%20Windows%20Server%202008%20R2,%20Windows%207,%20and%20Windows%20Vista).

For Windows Server 2012, Windows Server 2012 R2, Windows 8, and Windows

- 8.1 See *Best Practices for Recovering Windows Server 2012 and Windows 8* at [https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli Storage Manager/page/Best Practices for Recovering Windows Server 2012 and Windows 8](https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli%20Storage%20Manager/page/Best%20Practices%20for%20Recovering%20Windows%20Server%202012%20and%20Windows%208).

Restoring the Windows operating system with Automated System Recovery

You can restore the Windows operating system of a computer with Automated System Recovery (ASR).

Procedure

For instructions that describe how to restore a Windows system by using ASR, see the following Tivoli Storage Manager Wiki articles:

For Windows Server 2008, Windows Server 2008 R2, and Windows 7

See *Best Practices for Recovering Windows Server 2008, Windows Server 2008 R2, Windows 7, and Windows Vista* at [https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli Storage Manager/page/Best Practices for Recovering Windows Server 2008, Windows Server 2008 R2, Windows 7, and Windows Vista](https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli%20Storage%20Manager/page/Best%20Practices%20for%20Recovering%20Windows%20Server%202008,%20Windows%20Server%202008%20R2,%20Windows%207,%20and%20Windows%20Vista).

For Windows Server 2012, Windows Server 2012 R2, Windows 8, and Windows

- 8.1 See *Best Practices for Recovering Windows Server 2012 and Windows 8* at [https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli Storage Manager/page/Best Practices for Recovering Windows Server 2012 and Windows 8](https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli%20Storage%20Manager/page/Best%20Practices%20for%20Recovering%20Windows%20Server%202012%20and%20Windows%208).

What to do next

You can now restore other volumes.

Related tasks:

“Creating a bootable WinPE CD” on page 202

“Creating a client options file for Automated System Recovery” on page 151

Related reference:

“Restore” on page 730

“Restore Systemstate” on page 753

Microsoft Dfs tree and file restore

To restore Dfs junctions and the data for each junction, restore the Dfs junction metadata first and then restore each junction separately.

If the junction metadata is not restored, Tivoli Storage Manager creates a directory under the Dfs root using the same name as that of the junction point and restores the data in that directory.

Related tasks:

“Microsoft Dfs file protection methods” on page 190

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.
- The image restore operation overwrites the volume label on the destination volume with the one that existed on the source volume.
- Ensure that the volume to which you are restoring the image is at least as large as the image that is being restored.
- The file system or volume you are restoring to does not have to be the same type as the original. The volume does not even have to be formatted. The image restore process creates the appropriately formatted file system for you.
- 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 created an image of the system drive, you cannot restore the image to the same location because the client cannot have an exclusive lock of the system drive. Also, because of different system component configurations, the system image might not be consistent across components (such as Active Directory).

Some of these components can be configured to use different volumes where parts are installed on the system drive and others to non-system volumes.

- If you have run progressive incremental backups *and* image backups of your file system, you can perform an incremental image restore of the file system. The process restores individual files after the complete image is restored. The individual files restored are those backed up after the original image. Optionally, if files were deleted after the original backup, the incremental restore can delete those files from the base image.

Deletion of files 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 should run *chkdsk* to check for and repair any bad sectors (unless the restored volume is RAW).

You can use the *verifyimage* option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume. If bad sectors are detected on the target volume, the client issues a warning message on the console and in the error log.

If bad sectors present on the target volume, you can use the *imagetofile* option with the **restore image** command to specify that you want to restore the source image to a file. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume.

Related reference:

"Imagetofile" on page 428

"Verifyimage" on page 575

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 drive letter or mount point from which it was originally backed up. Alternatively, a different volume can be chosen for the restore location.
6. Click the **Restore** button to begin the restore. The **Task List** window appears showing the progress of the restore. The Restore Report window displays a detailed status report.

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. Because image backup allows you to back up only used blocks in a file system, the stored image size on the Tivoli Storage Manager server could be smaller than the volume size. For online image backups, the stored image can be larger than the file system based on the size of the cache files.
 - File system type
 - Backup date and time
 - Management class assigned to image backup
 - Whether the image backup is an active or inactive copy
- To modify specific restore options, click the **Options** button. Any options you change are effective during the current session *only*.
- In the Restore Options window, you can choose to restore the image only or the image and incremental directories files. If you choose **Image Only**, you 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.

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

If bad sectors are present on the target volume, you can use the **imagetofile** option with the **restore image** command to specify that you want to restore the source image to a file. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume.

Related reference:

“Imagetofile” on page 428

“Verifyimage” on page 575

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.
- You can perform an ASR restore using a backup set volume.

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.

The backup set can contain system state 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 26. 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 27. Backup set command-line restore restrictions

Data type in the backup set	Backup set location		
	Local (location=file or location=tape)	Tivoli Storage Manager Server (TOC available)	Tivoli Storage Manager Server (TOC not available)
file	Commands: restore restore backupset	Commands: restore restore backupset	Commands: restore backupset
image	Cannot be restored	Command: restore image	Cannot be restored
system state	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 456

“Query Backupset” on page 703

“Query Image” on page 710

“Restore” on page 730

“Restore Backupset” on page 739

“Restore Image” on page 748

“Restore Systemstate” on page 753

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.
- 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 456

“Restore” on page 730

“Restore Image” on page 748

“Restore Backupset” on page 739

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.

If you have upgraded from Tivoli Storage Manager Express®, some application data is also supported.

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. Restoring individual system state components.
 3. 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.
 4. Application data restore if the server was migrated from the Tivoli Storage Manager Express product.

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**
- **restore systemstate**

Use the appropriate command for the location of the backup set and the data in the backup set. For more information, see Table 27 on page 207.

Related reference:

"Query Backupset" on page 703

"Query Image" on page 710

"Restore" on page 730

"Restore Backupset" on page 739

"Restore Image" on page 748

"Restore Systemstate" on page 753

Restore Net Appliance CIFS shares

Restoring the share definition requires restoring the root directory of the share file space, which under most circumstances can be done as follows: `dsmc rest \\NetAppFiler\CifsShareName\ -dirsonly`.

The following output indicates that the root directory (and share definition has been restored):

```
Restoring          0 \\NetAppFiler\CifsShareName\ [Done]
```

If the CIFS share definition is deleted on the Net Appliance file server, the client is unable to directly restore the share definition because the share is no longer accessible.

The share definition can be restored indirectly by creating a temporary local share and restoring the share definition to the temporary share as follows:

```
md c:\tempdir net share tempshare=c:\tempdir
/remark:"Temporary Share for Restoring Deleted CIFS Share"
net use z: \\LocalMachineName\tempshare
dsmc res \\NetAppFiler\CifsShareName\ z:\ -dirsonly
```

This restores the original share definition (including permissions) on the file server.

Older versions of the Tivoli Storage Manager server might have a problem which prevents restoring the root directory and the CIFS share definition. If this problem occurs, it can be circumvented by using by one of the following methods:

1. Use the `DISABLENQR` testflag to restore the root directory as follows:

```
dsmc res \\NetAppFiler\CifsShareName\ -test=disablenqr -dirsonly
```

2. Use the command line client -pick option with a restore command and select the root directory:

```
dsmc res \\NetAppFiler\CifsShareName\ -dirsonly -pick
```

Related tasks:

“Backing up Net Appliance CIFS share definitions” on page 182

Restoring data from a VMware backup

You can use several methods for restoring data from backups to a VMware virtual machine. The restore method depends on the type of backup and on the version of the backup-archive client software that you use to run the restore.

Full VM restore

Use the **restore vm** command to restore an entire virtual machine from a full VM backup. When you restore a full VM backup, the restored image replaces the virtual machine or a new virtual machine is created. In a full VM restore, you restore all of the VMware files and the system state on Windows systems. If you have access to Tivoli Storage Manager recovery agent, you can restore individual files.

Depending on the version of the backup-archive client that is running on the VMware client, use the appropriate method to restore a full VM backup:

Versions of the backup-archive earlier than 6.2.2:

Restore the full VM backup by using VMware Consolidated Backup. For more information, see the following topic:

“Restoring full VM backups that were created with VMware Consolidated Backup” on page 222

Versions of the backup-archive client at 6.2.2 or later:

Restore the full VM backup by using the vStorage API. The Tivoli Storage Manager V6.2.2 or later client can restore full VMware backups that were created with versions of the client that is earlier than V6.2.2. For more information, see the following topic:

“Restoring full VM backups” on page 213

File-level restore

Use the **restore** command to restore individual files from a file-level VM backup. Use this method when you cannot practically restore an entire VMware image. The following restrictions apply to file-level restores:

- You can use the file-level restore method only if a file-level backup of the virtual machine exists.
- You cannot restore an entire virtual machine from file-level backups because the **restore** command does not re-create Windows system states.
- You cannot use this method to restore individual files from a full VM backup of a virtual machine.

Depending on the configuration of the virtual machine where you restore the files, use the appropriate method to restore files from a file-level backup:

The backup-archive client is not installed on the VM:

Restore the files from the vStorage backup server that backed up the virtual machine.

The backup-archive client is installed on the VM:

Restore the file from the backup-archive client that is installed on the virtual machine.

For more information, see the following topic:

“Scenario: Restoring file-level VM backups” on page 219

Restoring full VM backups

You can restore a full VMware backup to re-create all of the files for a VMware virtual machine (VM) directly to the VMware server. This method replaces the deprecated method of restoring backups that were created by using the VMware Consolidated Backup (VCB) tools. This restore method does not require you to use the VMware converter tool before you restore the backup to the VMware server. You cannot use this restore method to restore individual files from a full VM backup.

Before you begin

To restore a full VMware backup that was created by using VCB tools in Tivoli Storage Manager Version 6.2.0 or earlier, see the following topic:

“Restoring full VM backups that were created with VMware Consolidated Backup” on page 222

Procedure

1. Depending on the target location for the restore, complete the appropriate step:
 - If the restore of the full VM backup is going to overwrite the existing VMware virtual machine, delete the existing virtual machine.
 - If you restore the full VM backup to a new virtual machine, you do not need to delete the existing virtual machine. You can delete the existing virtual machine if you prefer, otherwise proceed to the next step.
2. Query the virtual machine for VMware backups, by completing the following steps:
 - a. From the off-host backup server, run the following command:

```
dsmc q vm *
```

The command lists the available backups, for example:

#	Backup Date	Mgmt Class	Type	A/I	Virtual Machine
1	12/03/2009 03:05:03	DEFAULT	VSTORFULL	A	vm_guest1
2	09/02/2010 10:45:09	DEFAULT	VSTORFULL	A	vm_guest11
3	09/02/2010 09:34:40	DEFAULT	VSTORFULL	A	vm_guest12
4	09/02/2010 10:10:10	DEFAULT	VSTORFULL	A	vm_guest13
5	12/04/2009 20:39:35	DEFAULT	VSTORFULL	A	vm_guest14
6	09/02/2010 11:15:18	DEFAULT	VSTORFULL	A	vm_guest15
7	09/02/2010 02:52:44	DEFAULT	VSTORFULL	A	vm_guest16
8	08/05/2010 04:28:03	DEFAULT	VSTORFULL	A	vm_guest17
9	08/05/2010 05:20:27	DEFAULT	VSTORFULL	A	vm_guest18
10	08/12/2010 04:06:13	DEFAULT	VSTORFULL	A	vm_guest19
11	09/02/2010 00:47:01	DEFAULT	VSTORFULL	A	vm_guest7
12	09/02/2010 01:59:02	DEFAULT	VSTORFULL	A	vm_guest8
13	09/02/2010 05:20:42	DEFAULT	VSTORFULL	A	vm_guest9

ANS1900I Return code is 0.
ANS1901I Highest return code was 0.

- b. From the results that are returned by the query command, identify a virtual machine to restore.

3. Restore the full VMware backup, by using the **restore vm** command. To restore the backup to a virtual machine with a new name, use the **-vmname** option. For example, in the following command the virtual machine is restored and a new name is specified for the restored virtual machine:

```
dsmc restore vm my_old_vmname -vmname=new_vm_name -datastore=myPath
```

4. When the restore is complete, the virtual machine is powered off. Start the virtual machine from the VMware vCenter.

Related reference:

“Query VM” on page 725

“Restore VM” on page 758

Scenarios for running full VM instant access and full VM instant restore from the backup-archive client command line

Full VM instant access and full VM instant restore operations require a license for Tivoli Storage Manager for Virtual Environments. You can perform either of these operations from the backup-archive client command line. Instant access and instant restore operations and options are supported only for VMware virtual machines that are hosted on VMware ESXi 5.1 servers, or later versions.

The following scenarios demonstrate the full VM instant access or full VM instant restore operations that you might perform. Before you can complete the operations that are described in the following text, you must configure at least one data mover node on the vStorage backup server so it can protect the virtual machines by starting off host backup and restore operations. The steps for setting up the data mover nodes are described in Setting up the data mover nodes in a vSphere environment.

Scenario: You want to perform a full VM instant access to verify the integrity of a backed up image of a VMware virtual machine, without actually restoring the virtual machine or disks to the ESXi host

The purpose of this goal is to verify that a backed up virtual machine image can be used to successfully restore a system if the virtual machine is deleted or its disks and data are corrupted or otherwise unusable.

For this scenario, assume that an ESX server has a virtual machine named Orion running on it. You want to verify that the backed up image that is stored by the Tivoli Storage Manager server can be used to restore this virtual machine if the current virtual machine fails.

You perform a VM instant access operation, you use the **restore vm** command with inventory location options specified to identify the location for the restored virtual machine. All inventory location options, such as **vmname**, **datacenter**, **host**, and **datastore** can be used in combination with the instant access option (**-VMRESToretype=INSTANTAccess**) to specify the location for the restored (instant access) virtual machine.

Because the Orion virtual machine does exist in the inventory and is running, you must provide a new name for a temporary virtual machine by adding the new name to the **vmname** option. You must also add the **-VMRESToretype=INSTANTAccess** option to the command line to indicate that this is an instant access restore operation.

Entering the following command prepares a virtual machine named “Orion_verify” so it is available for instant access. You can use this virtual machine to verify that the backed-up image can be restored.

```
dsmc restore vm Orion -vmname=Orion_verify -Host=esxi.example.com  
-datacenter=mydataCenter -VMRESToretype=INSTANTAccess -VMAUTOSTARTvm=YES
```

The `-VMAUTOSTARTvm=YES` option indicates that the virtual machine is started when it is restored. By default, the new virtual machine is not automatically started. With this default setting, you can reconfigure the virtual machine before you start it.

You can also list the versions of a virtual machine that were backed up by using the **inactive** or **pick** options or the **pittime** or **pitdate** options to select an inactive or active backup, from a particular date or time. For example, to display a list of backed up versions of the Orion virtual machine, by using the following command:

```
dsmc restore vm Orion -pick
```

For a virtual machine that is restored by using the `-VMRESToretype=INSTANTAccess` option, temporary data that is created by this virtual machine is stored in a VMware snapshot.

After you restore the temporary virtual machine (Orion_verify), run verification tools on it to verify the integrity of the disks and data. Use a utility such as **chkdsk**, or a utility or application of your choosing, to verify the virtual disks and data. If the temporary virtual machine passes the integrity checks, you can remove the temporary resources that were created to support the instant access restore operation.

Scenario: You want to determine whether any temporary (instant access) virtual machines exist, so you can run a clean-up operation to free the resources associated with them

Use the **query vm** command with one of the following options that you specify on the command line:

```
-VMRESToretype=INSTANTAccess  
-VMRESToretype=ALLtype
```

Where:

-VMRESToretype=INSTANTAccess

Displays all temporary virtual machines that are running in instant access mode, created by a `restore vm -VMRESToretype=INSTANTAccess` operation.

-VMRESToretype=ALLtype

Displays all virtual machines with active instant access or instant restore sessions that were started by a **restore vm** command that uses either the `-VMRESToretype=INSTANTAccess` or `VMRESToretype=-INSTANTRestore` options.

The following examples show the syntax for the various options:

```
query vm * -VMREST=INSTANTA  
query vm * -VMREST=ALL
```

You can add a **-Detail** option to each of the **query vm** commands shown to display more information about each of the temporary virtual machines.

```
query vm vmname -VMREST=INSTANTA -Detail
```

To remove the resources that were created for a temporary virtual machine named "Orion_verify", run the following command:

```
dsmc restore vm Orion -vmname=Orion_verify -VMRESToretype=VMCleanup
```

The **-VMRESToretype=VMCleanup** option deletes the temporary virtual machine from the ESXi host, unmounts any iSCSI mounts that were mounted, and clears the iSCSI device list from the ESX host. All temporary data for the temporary virtual machine is deleted from the VMware snapshot.

Scenario: You want to start an instant restore operation to restore a failed virtual machine to an ESX host, from a backup image created by Tivoli Storage Manager

The advantage of a full VM instant restore, as opposed to a classic full VM restore, is that an instant restore operation makes the virtual machine ready for immediate use, as soon as it is started. You do not have to wait for all data to be restored before you can use the virtual machine. During an instant restore operation, the virtual machine uses iSCSI disks until its local disks are fully restored. When the local disks are restored, the virtual machine switches I/O from the iSCSI disks to the local disks, without noticeable interruption of service.

Restore a virtual machine named Orion by using the following command:

```
dsmc restore vm Orion -Host=esxi.example.com -datacenter=mydatacenter  
-VMTEMPDatastore=temp_datastore -VMRESToretype=INSTANTRestore  
-datastore=mydatastore
```

This command specifies the name of the virtual machine to restore, the host and data center to restore it to, and the restore type (-VMRESToretype=INSTANTRestore). The **VMTEMPDatastore** option is a mandatory parameter for instant restore operations.

The temporary datastore is used by vMotion to store the configuration of the restored virtual machine during the instant restore process. The name that you specify must be unique. It cannot match the name of any of the original datastores that were used by the virtual machine when it was backed up, and it cannot be the same as the name specified on the optional **-datastore** option. If the **-datastore** option is omitted, the virtual machine files are restored to the datastores that they used when the virtual machine was backed up.

By default, virtual machines that are instantly restored are provisioned with thick disks. You can change this behavior and provision thin disks by adding the **-VMDISKProvision=THIN** option to the command line, or in the client options file.

Important: For instant restore operations, ensure that both the temporary datastore that you specify with the **vmtempdatastore** option and the VMware datastore that is specified by the **datastore** option on the **restore VM** command have enough free storage to save the virtual machine that you are restoring, and the snapshot file that contains changes that were made to the data. If you are restoring a virtual machine and you specify thin or thick provisioning (**-vmdiskprovision=thin** or **-vmdiskprovision=thick**), the datastore that you restore the VM to must have enough free space to accommodate the total capacity of the VM disk, and not just the amount of disk that is used. For example, if a VM has 300 GB total capacity for its disk, you cannot restore that VM to a datastore that has less than 300 GB available, even if only a portion of the total capacity is being used.

Full VM instant restore cleanup and repair scenarios

When an instant restore operation fails after the VM is powered on, manual cleanup and repair tasks are required.

An instant restore operation that fails with storage vMotion running creates either of the following situations:

- The instant restore operation generates an error message.
- The instant restore operation suspends indefinitely and the VM is not responsive.

To determine the cause of the problem, perform a detailed query of the VM by using the following command:

```
dsmc q vm * -vmrestoretype=instantrestore -detail
```

In the output that is produced by this command, for each VM in the output, look for the line that contains Action Needed. Use the following *Action Needed* paragraphs to recover from failed instant restore operation, depending on the Action Needed status.

Action Needed: Cleanup

In the output of the query `vm * -vmrestoretype=instantrestore -detail` command, verify that the storage vMotion status is successful (vMotion Status: Successful) and that all VM disks are physical disks (Disk Type: Physical). This status confirms that the VM was restored and cleanup of orphaned components, such as iSCSI mounts, is needed.

This type of failure occurs as a result of either of the following situations:

- The instant restore failed and Storage vMotion is running. VMware vSphere continues the vMotion process.
- Storage vMotion finished successfully, but the automatic cleanup of the iSCSI mounts fails.

To clean up any orphaned components, run the **restore vm** command with the **-VMRESToretype=VMCleanup** parameter. For example:

```
dsmc restore vm original_vmname -vmname=new_vm_name -VMRESToretype=VMCleanup
```

Action Needed: Repair

In the output of the query `vm * -vmrestoretype=instantrestore -detail` command, verify that the iSCSI device that is attached to the VM is dead (status is Disk Path: Dead).

This type of failure occurs as a result of one of the following three situations:

- The VM that is used as a data mover or the physical data mover machine failed.
- A network failure occurred between the data mover and the ESX host or the data mover and the Tivoli Storage Manager server.
- The Data Protection for VMware Recovery Agent Service failed.

The iSCSI device must be returned to an active state before any other instant operation is attempted.

To attempt to recover from a data mover failure, complete the following steps:

1. Investigate that cause of the failure and restart the data mover machine if it does not start automatically. This action starts an automatic recovery of the mounted iSCSI disks.
2. In the output of the query `vm * -vmrestoretype=instantrestore -detail` command, verify that the VM disks are active (Disk Path: Active). This status means that the VM was restored and is available for use.
3. Restart storage vMotion in the vSphere client and monitor its progress in the vSphere client status bar.
4. If storage vMotion processing completed successfully, run the **restore vm** command with the **-vmrestoretype=VMCleanup** parameter to clean up the iSCSI disks. For example:

```
dsmc restore vm original_vmname -vmname=new_vm_name -VMRESToretype=VMCleanup
```

To attempt recovery after a network failure, complete the following steps:

1. Repair the network issue so that communication between the data mover and the ESX host, and the data mover and the Tivoli Storage Manager server resumes.
2. In the output of the query `vm * -vmrestoretype=instantrestore -detail` command, verify that the VM disks are active (Disk Path: Active). This status means that the VM was restored and is available for use.
3. If the network failure did not cause storage vMotion to time out, no action is required.
4. If the network failure caused storage vMotion to time out, and the error message indicates that the source disk is not responding, restart storage vMotion in the vSphere client. When storage vMotion processing completes, run the **restore vm** command with the **-vmrestoretype=VMCleanup** parameter to clean up the iSCSI disks. For example:

```
dsmc restore vm original_vmname -vmname=new_vm_name -VMRESToretype=VMCleanup
```

To attempt recovery after a Data Protection for VMware Recovery Agent service failure, complete the following steps:

1. Investigate that cause of the failure and restart the Data Protection for VMware Recovery Agent service if it does not start automatically. This action starts an automatic recovery of the mounted iSCSI disks.
2. In the output of the query `vm * -vmrestoretype=instantrestore -detail` command, verify that the VM disks are active (Disk Path: Active). This status means that the VM was restored and is available for use.
3. If the Data Protection for VMware Recovery Agent service failure did not cause storage vMotion to time out, no action is required.
4. If the Data Protection for VMware Recovery Agent service failure caused storage vMotion to time out, and the error message indicates that the source disk as not responding, restart storage vMotion in the vSphere client. When storage vMotion processing completes, run the **restore vm** command with the **-vmrestoretype=VMCleanup** parameter to clean up the iSCSI disks. For example:

```
dsmc restore vm original_vmname -vmname=new_vm_name -VMRESToretype=VMCleanup
```

Full cleanup

If you are not able to recover from a failure and want to remove the VM and its components, run the **restore vm** with the **-vmrestoretype=VMFULLCleanup** parameter. For example:

```
dsmc restore vm original_vmname -vmname=new_vm_name -VMRESToretype=VMFULLCleanup
```

A **VMFULLCleanup** operation forces removal of the VM and all of its components, regardless of the state of the virtual machine. Do not start a full clean up operation while vMotion is still migrating a virtual machine.

Recovering from non-standard error conditions

Problems with iSCSI devices can prevent you from performing an instant access or instant restore operation.

About this task

When an ESX server cannot access a datastore on an iSCSI disk, a VMware message is issued to indicate that a "permanent device loss" error occurred. You should be offered an option to either cancel or retry the iSCSI connection attempt. Choose the option to try the operation again to see whether the error is transient and if recovery is possible. If the retry is not successful, try the following troubleshooting steps. If they are successful, then try the instant restore or instant access operation again.

Procedure

1. Examine the ESX server Task and Event log for an All Paths Down (APD) error. It can take time for this error to display in the logs, but it must be present before you continue to the next steps. If you do not wait for the error before you attempt more troubleshooting, you might bring the ESX server down.
2. Power off the virtual machine.
3. Rescan the HBA. Rescanning the HBA on the ESX server might reactivate the failed device. If VMware kernel locks prevent you from rescanning the HBA, perform the following steps:
 - a. In the vCenter interface, select the ESX host.
 - b. Click **Configuration**.
 - c. Right click **iSCSI Software Adapter** and select **Properties**.
 - d. Click **Static Discovery**.
 - e. Delete any static addresses and click **Close**.
 - f. Rescan the HBA.

Scenario: Restoring file-level VM backups

On Microsoft Windows systems, you can restore specific files from a file-level backup of a VMware virtual machine. A file-level restore is useful for restoring individual files that might be lost or damaged. You cannot use this method to restore files that were part of a full VM backup. Before you can restore files from the off-host backup server onto the VMware virtual machine, the off-host backup server must be configured as a proxy server.

Before you begin

For prerequisite information, see the following topics:

- For information about how to configure the off-host backup server as a proxy server, see the following topic:
"Preparing the environment for file-level backups of VMware virtual machines" on page 171.
- For information about how to create a file-level VM backup, see the following topic:
"Creating file-level backups for VMware virtual machines" on page 173

Important: Use the **restore** command to run a file-level restore. Do not use the **restore vm** command.

The following assumptions are made for this scenario of a file-level restore:

- The goal is to restore files that were previously backed up to the Tivoli Storage Manager server.
- The files were previously backed up on a VMware virtual machine called Orion, with the host name orion. For this scenario, the Orion VM fails and some of the files must be restored.
- Files on Orion were backed up to file spaces that match the lowercase form of the computer host name. The file space names are expressed in Universal Naming Convention (UNC) format, for example:
 - Files that are backed up from the C: drive on Orion, are stored in the \\orion\c\$ file space.
 - If Orion has a D: drive, files that are backed up from that drive are stored in the \\orion\d\$ file space.
- In this scenario, the files are restored from the C:\mydocs directory that was on Orion to the C:\restore_temp directory on a different computer. The computer that you restore file to can be another VMware virtual machine or a physical computer.
- The computer that runs the restore has a different host name and node name than the virtual machine Orion. During the restore, you must specify the source file specification in the complete UNC format and use one of the following parameters to access Orion:

-virtualnodename

Specifies the client node for which you are restoring a backup. Use this parameter if you are restoring files to the computer where you are currently logged on.

-asnodename

Specifies the client node for which you are restoring a backup. Use this parameter if you are restoring files to a computer for which you have proxy authority.

Complete the following steps to run a file-level restore for the computer Orion:

Procedure

1. Query the Tivoli Storage Manager server to determine the file spaces that are registered for Orion:
`dsmc query filespace -virtualnode=orion`
2. Restore files for the Orion file space, by running one of the following commands:

Restore files to the computer where you are currently logged on:

Assume that you are currently logged on to the computer called Orion. Run one of the following commands:

- a. If you know the password for the node that you are restoring, use the **-virtualnodename** option in the restore command. For example, run the following command to restore the files to Orion:

```
dsmc restore \\orion\c$\mydocs\ c:\restore_temp\ -sub=yes  
-virtualnodename=orion
```

- b. If you have proxy authority, you can restore files on behalf of the target node. Proxy authority must be granted from the agent node, in other words the node of the computer that the restore is run

from. You must know the password for the agent node so that you can access the target node. For example, run the following command to restore the files to Orion:

```
dsmc restore \\orion\c$\mydocs\ c:\restore_temp\ -sub=yes
-asnodename=orion
```

Table 28. Components for the restore command when you restore files to the same computer

Command component	Description
\\orion\c\$\mydocs\	Source file specification on the Tivoli Storage Manager server. This location contains the backed up files that you are restoring. The files are backed up for the orion VM, so the file specification must be in UNC format.
c:\restore_temp\	Destination file specification on the computer where you are currently logged on. The files are restored to this location.
-sub=yes	Specifies that all subdirectories in the source file specification are included when you run the restore operation.
-virtualnodename=orion	Notifies the Tivoli Storage Manager server that the backup is running from the node orion.
-asnodename=orion	Notifies the Tivoli Storage Manager server that the backup is running from the node orion.

Restore files to a different computer:

To restore the files from the Tivoli Storage Manager server to a computer other than the one you are logged on to, run the following command. You can use this command only if you are logged in with authority to write to the remote computer as controlled by the operating system.

```
dsmc restore \\orion\c$\mydocs\ \\orion\c$\restore_temp\ -sub=yes
-virtualnode=orion
```

Table 29. Components for the restore command when you restore files to a different computer

Command component	Description
\\orion\c\$\mydocs\	Identifies the source file specification on the Tivoli Storage Manager server. This location contains the backed up files that you are restoring. The files are backed up for the orion VM, so the file specification must be in UNC format.
\\orion\c\$\restore_temp\	Identifies the destination file specification on a computer other than the computer where you are logged on. The You are restoring the files to the orion VM over the network, by using a Microsoft feature that identifies network locations in UNC notation.
-sub=yes	Specifies that all subdirectories in the source file specification are included when you run the restore operation.

Table 29. Components for the restore command when you restore files to a different computer (continued)

Command component	Description
-virtualnodename=orion	Notifies the Tivoli Storage Manager server that the backup is running from the node orion.

Related concepts:

“Restoring data from a VMware backup” on page 212

Related tasks:

“Restoring full VM backups that were created with VMware Consolidated Backup”

“Restoring full VM backups” on page 213

Related reference:

“Query Filespace” on page 706

“Restore” on page 730

Restoring full VM backups that were created with VMware Consolidated Backup

You can restore a full VMware backup to re-create all of the files for a VMware virtual machine (VM). Complete these steps to restore full VM backups that were created by using VMware Consolidated Backup (VCB) running on Tivoli Storage Manager Version 6.2.0 or earlier.

Before you begin

To restore a full VMware backup that was created by using Tivoli Storage Manager Version 6.2.2 or later, see the following topic:

“Restoring full VM backups” on page 213

Procedure

1. Depending on the target location for the restore, complete the appropriate step:
 - If the restore of the full VM backup is going to overwrite the existing VMware virtual machine, delete the existing virtual machine.
 - If you restore the full VM backup to a new virtual machine, you do not need to delete the existing virtual machine. You can delete the existing virtual machine, otherwise proceed to the next step.
2. Query the virtual machine for full VMware backups, by completing the following steps:
 - a. From the off-host backup server, run the following command:

```
dsmc q vm *
```

The command lists the available backups, for example:

#	Backup Date	Mgmt Class	Type	A/I	Virtual Machine
1	12/03/2009 03:05:03	DEFAULT	VMFULL	A	vm_guest1
2	09/02/2010 10:45:09	DEFAULT	VMFULL	A	vm_guest11
3	09/02/2010 09:34:40	DEFAULT	VMFULL	A	vm_guest12
4	09/02/2010 10:10:10	DEFAULT	VMFULL	A	vm_guest13
5	12/04/2009 20:39:35	DEFAULT	VMFULL	A	vm_guest14
6	09/02/2010 11:15:18	DEFAULT	VMFULL	A	vm_guest15
7	09/02/2010 02:52:44	DEFAULT	VMFULL	A	vm_guest16
8	08/05/2010 04:28:03	DEFAULT	VMFULL	A	vm_guest17
9	08/05/2010 05:20:27	DEFAULT	VMFULL	A	vm_guest18
10	08/12/2010 04:06:13	DEFAULT	VMFULL	A	vm_guest19
11	09/02/2010 00:47:01	DEFAULT	VMFULL	A	vm_guest7
12	09/02/2010 01:59:02	DEFAULT	VMFULL	A	vm_guest8
13	09/02/2010 05:20:42	DEFAULT	VMFULL	A	vm_guest9

ANS1900I Return code is 0.
ANS1901I Highest return code was 0.

- b. From the results that are returned by the query command, identify a virtual machine to restore.
3. Restore the full VMware backup, by using the **restore vm** command. To restore a virtual machine from a specific point in time, include the **-pitdate** and **-pittime** options, for example:

```
dsmc restore vm my_vm_name destination -pitdate=date -pittime=hh:mm:ss
```

Where:

my_vm_name

Name of the virtual machine that you are restoring.

destination

Directory location for the restored vmdk file.

-pitdate

Date that the backup was created.

-pittime

Time that the backup was created.

4. When the restore is completed, the following message is returned. Enter Y.
Virtual Infrastructure Client or VMware Converter tool
can be used to redefine virtual machine to the VMware Virtual Center Inventory.

Would you like to launch VMware Converter now? (Yes (Y)/No (N))

Tip: If you enter N, the command-line returns without opening the VMware Converter. However, you must convert the image before the image can be restored.

5. To convert the restored VCB image into a virtual machine on a VMware server by using the VMware vCenter Converter tool, complete following steps:
 - a. From the Windows Start menu, open the Converter tool.
 - b. From the Converter tool, click **Convert Machine**.
 - c. In the **Virtual machine file** field, enter the location of the restored .vmx file.

Tip: The .vmx file is restored to the directory specified by the **vmbackdir** option of the **restore vm** command.

- d. Follow the remaining steps in the wizard to convert the full VM backup.
6. When the restore is complete, the virtual machine is powered off. Start the virtual machine from the VMware vCenter.

Related reference:

“Query VM” on page 725

“Restore VM” on page 758

Restore Windows individual Active Directory objects

You can restore individual Active Directory objects to recover from accidental corruption or deletion of Active Directory objects without requiring a shutdown or restart of the Active Directory server.

On the Windows Server client, use the **restore adobjects** command to restore local, deleted Active Directory objects (tombstone objects). You can also restore individual Active Directory objects from system state backups on the Tivoli Storage Manager server.

Related tasks:

“Restoring Windows system state” on page 199

Related reference:

“Restore Adobjects” on page 738

Reanimate tombstone objects or restoring from a system state backup

Tombstone reanimation is a process to restore an object that had been deleted from the Active Directory. When an object is deleted from Active Directory, it is not physically erased, but only marked as deleted. It is then possible to reanimate (restore) the object.

When an object is reanimated, not all object attributes are preserved. When an object becomes a tombstone object, many attributes are automatically stripped from it, and the stripped attributes are lost. It is possible, however, to change the Active Directory schema so that more attributes are preserved when the object is deleted.

User-group links are not preserved in tombstones. For example, when a user object is reanimated, the user account is not a member of any group. All of this information must be recreated manually by the Active Directory administrator.

When an Active Directory object is restored from a system state backup on the Tivoli Storage Manager server, virtually all of its attributes and its group membership are restored. This is the best restore option using a Windows Server domain controller. When an object is restored from the server:

- The Active Directory database is extracted from a system state backup and restored into a temporary location.
- The restored database is opened.
- Select which objects you want to restore. For each object:
 - A search for the matching tombstone is performed. The Globally Unique Identifier (GUID) of the restored object is used to search for the tombstone.
 - If the matching tombstone is found, it is reanimated. In this case, the restored object retains the original Globally Unique Identifier (GUID) and the Security Identifier (SID).
 - If the matching tombstone is not found, a new object is created in the database. In this case, the new object has a new GUID and a new SID that are different than the original object.
- Missing attributes are copied from the backup into the reanimated or recreated object. Existing attributes that have been changed since the backup was taken

are updated to match the value in the backup. New attributes that have been added since the backup was taken are removed.

- Group membership is restored.

Although all attributes that can be set and the group links are recreated, the restored objects might not be immediately available after the restore operation. An Active Directory administrator might have to manually update the restored objects in order to make them available. Make sure to read “Restrictions and limitations when restoring Active Directory objects” on page 226 before performing the restore.

Related concepts:

“Preserve attributes in tombstone objects” on page 227

Chapter 5, “Restoring your data,” on page 193

“Restrictions and limitations when restoring Active Directory objects” on page 226

Related tasks:

“Restoring Windows system state” on page 199

Related reference:

“Restore Adobjects” on page 738

Restoring Active Directory objects using the GUI and command line

To restore individual Active Directory objects, you must run the backup-archive client on a domain controller and your user account must be a member of the Administrators group. The Active Directory objects are not displayed in the directory tree if your user account is not a member of the Administrators group.

You can restore active directory objects or tombstone objects using either the GUI or the command line.

To restore individual objects from the GUI:

1. Click **Restore** in the IBM Tivoli Storage Manager window. The Restore window opens.
2. Expand the directory tree if necessary. To expand an object in the tree, click the plus sign (+) next to the object.
3. Locate the Active Directory node in the directory tree. Expand it to reveal **Local Deleted Objects**. The Server object is also available.
 - To restore tombstone objects, expand **Local Deleted Objects**, navigate to the tombstone objects that you want to restore, and select the tombstone objects.
 - To restore Active Directory objects that are backed up to the Tivoli Storage Manager server:
 - a. Expand the **Server** object. A window opens displaying a list of system state backups (with different time stamps) on the server.
 - b. Select a system state backup from the list. The Active Directory database from that system state is restored in the background, and the tree is populated with Active Directory objects.
 - c. Navigate to the Active Directory objects that you want to restore and select the Active Directory objects.

Tip: To see the attributes for an Active Directory object, keep expanding each Active Directory object in the tree until you reach the one you want. The

attributes for an object are displayed in the display area that is adjacent to the tree. You can search or filter the tree for an Active Directory object based on its name.

4. Click **Restore** to begin the restore operation. The Task List window opens and shows the progress of the restore operation.

On the command line, use the **query adobjects** command to query and the **restore adobjects** command to restore individual Active Directory objects.

Related reference:

"Query Adobjects" on page 696

"Restore Adobjects" on page 738

Restrictions and limitations when restoring Active Directory objects

There are some restrictions and limitations to be aware of when restoring Active Directory objects.

Understand the following restrictions before restoring objects:

- Do not restore the Active Directory as part of a system-state restore operation, unless it is intended to be used for a disaster recovery-level restore operation of the full Active Directory. This type of restore operation requires the Active Directory Server to be stopped and restarted.
- You cannot perform a point-in-time restore of tombstone objects. You can perform a point-in-time restore of Active Directory objects that are backed up to the server.
- You cannot restore Active Directory objects from backup sets.

Understand the following limitations before restoring objects:

- Restoring Active Directory objects from the Tivoli Storage Manager server requires temporary space on your local hard disk drive. You can use the `stagingdirectory` option to specify a directory on your local hard disk for storing temporary data from the server. Depending on the size of the temporary data, network bandwidth, and both client and server performance, this operation can take anywhere from 20 seconds to over an hour. There might be a delay in refreshing the Restore window when displaying the Active Directory tree.
- User passwords cannot be restored by default. A restored user object is disabled until the administrator resets the password and re-enables the account. Also, if an account was deleted from the domain and is then restored by the backup-archive client, it must be manually joined to the domain after the restore operation. Otherwise, users on the target computer cannot log on to the domain. In order to have a user or a computer object fully operational after restore, you must modify schema attribute *Unicode-Pwd* as described in **Preserve attributes in tombstone objects**.
- The Active Directory schema is not recreated when the Active Directory object is restored. If the schema was modified after the backup, the restored object might no longer be compatible with the new schema, and some Active Directory object attributes might no longer be valid. The client issues a warning message if some attributes cannot be restored.
- Group Policy Objects and their links to organizational units (OU) cannot be restored.
- Local policies for restored Active Directory objects are not restored.

- When you restore an object from the Tivoli Storage Manager server, if the target object already exists in the Active Directory and you replace it with its backup version, the object is not deleted and recreated. The existing object is used as a base, and its attributes are overwritten by the backup version. Some attributes, such as the GUID and the SID, stay with the existing object and are not overwritten by the backup version.
- If there are multiple tombstone objects for the same container, reanimate them from the backup-archive client command line using the object GUID, in which case the command-line client only reanimates the container object and not its children. In the backup-archive client GUI, the entire container can be selected to reanimate.
- When you restore an object from the Tivoli Storage Manager server, if the live Active Directory object exists and has the *prevent deletion* bit on, the client can modify the attributes of the object. However, if there is a tombstone object of the same name but a different object GUID, the Directory Services returns the *access denied* error.
- When you restore an object from the Tivoli Storage Manager server and the container of the object has been renamed, the client recreates the container using the original name at the time of the backup. When restoring a tombstone object, the client restores it to the renamed container because the *lastKnownParent* attribute of the tombstone object has been updated to reflect the new container name.

Related concepts:

“Preserve attributes in tombstone objects”

Chapter 5, “Restoring your data,” on page 193

Related reference:

“Restore Adobjects” on page 738

“Stagingdirectory” on page 547

Preserve attributes in tombstone objects

To specify an attribute to be preserved in the tombstone object, first locate this attribute in the Active Directory schema, then update the *searchFlags* attribute of the schema object.

There is vendor-acquired software (for example, ADSI Edit) that allows you to update the *searchFlags* attribute of the schema object.

Usually none of the bits in the *searchFlags* bit mask are set (the value is 0). Set *searchFlags* to 8 (0x00000008) if you want Active Directory to save the particular attribute in the tombstone object when the original object is deleted.

Related concepts:

Chapter 5, “Restoring your data,” on page 193

Related reference:

“Restore Adobjects” on page 738

Modifying the client acceptor and agent services to use the web client

You cannot restore individual Active Directory objects using the web client by default. The web client services (CAD and agent) run under the Local System account by default. The Local System account does not have enough privileges to restore Active Directory objects.

To enable this restore operation in the web client, follow these steps:

1. Modify the client acceptor (CAD) and agent services to use an administrative account such as *Administrator* when logging on to Windows.
2. You can edit the properties for the CAD and agent services (typically called TSM Client Acceptor and TSM Remote Client Agent) in the Control Panel.
3. Modify the CAD and the agent services in the **Login Options** page of the Tivoli Storage Manager configuration wizard when you set up the web client

If the web client is already set up, follow these steps:

1. Click **Start**.
2. Click **Control Panel** → **Administrative Tools** → **Services**.
3. Select the scheduler service from the list of Windows services.
4. Click the **Log On** tab.
5. Click **This Account** in the Login As section.
6. Enter an administrative account, or click **Browse** to locate the domain account.
7. Enter the password for the domain account.
8. Click **OK** and then click **Start**.

Related reference:

“Restore Adobjects” on page 738

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 61.

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 C:\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 60

Related tasks:

“Determining the status of replicated client data” on page 65

Authorizing another user to restore or retrieve your files

You can authorize a user on another node to restore your backup versions or retrieve your archive copies. In this way, you can share files with other people or with other workstations that you use with a different node name.

About this task

You can also authorize other nodes to access the automated system recovery (ASR) file space.

Another node can be used to create the ASR diskette so that the workstation can be recovered using ASR and the backup-archive client. Use the other node if a problem occurs with the workstation and the ASR diskette of the workstation is not available.

To authorize another node to restore or retrieve your files:

Procedure

1. Click **Utilities** → **Node Access List** from the main window.
2. In the **Node Access List** window, click the **Add** button.
3. In the **Add Access Rule** window, select an item in the **Permit Access** field to specify the type of data that the other user can access. You can select either **Backed up Objects** or **Archived Objects**.
4. Type the node name of the user in the **Grant Access to Node** field. Type the node name of the host workstation of the user in the **Grant Access to Node** field.
5. Type the user ID on the host workstation in the **User** field.
6. In the **Filespace and Directory** field, select the file space and the directory that the user can access. You can select one file space and one directory at a time. If you want to give the user access to another file space or directory, you must create another access rule.
7. If you want to limit the user to specific files in the directory, type the name or pattern of the files on the server that the other user can access in the **Filename** field. You can make only one entry in the **Filename** field. It can either be a single file name or a pattern that matches one or more files. You can use a wildcard character as part of the pattern. Your entry must match files that have been stored on the server.
8. If you want to give access to all files that match the file name specification within the selected directory including its subdirectories, click **Include subdirectories**.
9. Click **OK** 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 **OK**. If you do not want to save your changes, click **Cancel** or close the window.

Results

For example, to give the node user2 access to all backup files and subdirectories under the d:\user1 directory, create a rule with the following values:

```
Permit Access to: Backed up Objects
Grant Access to Node: user2
Filespace and Directory: d:\user1
Filename: *
Include subdirectories: Selected
```

The node you are authorizing must be registered with your Tivoli Storage Manager server.

On the command line of the client, use the **set access** command to authorize another node to restore or retrieve your files. You can also use the **query access** command to see your current list, and **delete access** to delete nodes from the list.

Related reference:

“Delete Access” on page 671

“Query Access” on page 695

“Set Access” on page 775

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 for another user on the server, restore the backup versions of files for another user, or retrieve the archive copies for another user to your local file system, by following these steps:

Procedure

1. Click **Utilities** from the main window.
2. Click **Access Another Node**.
3. Type the node name of the host workstation of the user in the **Node name** field and click **Set**.

Results

If you are using commands, use the `fromnode` option to indicate the node. You must also use the file space name, rather than the drive letter, to select the restore-retrieve drive that you want to access. Include the file space name in braces and specify it as you would specify a drive letter. For example, to restore the files from the cougar node \projx directory on the d-disk file space to your own \projx directory, enter:

```
dsmc restore -fromnode=cougar \\cougar\d$\projx\* d:\projx\
```

Use the **query filespace** command to display a list of file spaces. For example, to display a list of the file spaces of cougar, enter:

```
dsmc query filespace -fromnode=cougar
```

Important: The backup-archive client can use file space information when restoring files. The file space information can contain the name of the computer from which the files were backed up. If you restore files from another client node and do not specify a destination for the restored files, the client uses the file space information to restore the files. In this case, the client attempts to restore the files to the drive on the original computer. If the restoring computer has access to the drive of the original computer, you can restore files to the original drive. If the restoring computer cannot access the drive of the original computer, the client returns a network error message. If you want to restore the original directory structure but on a different computer, specify only the target drive when you restore the files. This is true when restoring files from another node and when retrieving files from another node.

Related reference:

“Fromnode” on page 420

“Restore” on page 730

“Retrieve” on page 766

Restoring or retrieving your files to another workstation

When you are using a different workstation, you can restore or retrieve files you backed up from your own workstation.

Your backup versions and archive copies are stored according to your node, not your specific workstation. Your Tivoli Storage Manager password protects your data.

To restore or retrieve files to another workstation, use the **virtualnodename** option to specify the node name of the workstation from which you backed up the files. You can use the **virtualnodename** option when starting Tivoli Storage Manager or place the option in your client options file, `dsm.opt`, on the workstation. If you are using a workstation other than your own, use the **virtualnodename** option with the **dsm** command. For example, if your node name is cougar, enter:

```
start dsm -virtualnodename=cougar
```

You can then restore or retrieve files as if you were working from your original workstation.

You can also use **virtualnodename** option on commands. For example, to restore your `\projx` files to your local `c:\myfiles` directory, enter:

```
dsmc restore -virtualnodename=cougar \\cougar\d$\projx\*. * c:\myfiles\
```

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

Restoring or retrieving files to another type of workstation

You can restore or retrieve files from one system type to another. This is called *cross-client restore*.

Restriction: You must have the appropriate permissions to access the file space of the other workstation.

NTFS and ReFS drives permit file and directory names that are longer than those permitted on FAT drives. If you are recovering files to a FAT drive with long file names, specify a destination file specification for each file.

When you use the Windows client to recover files with long names to an NTFS or ReFS file system, the long names are preserved, even if you are recovering the file to a different type of drive than the source drive.

Related tasks:

“Authorizing another user to restore or retrieve your files” on page 229

“Restoring or retrieving files from another client node” on page 231

Deleting file spaces

If your Tivoli Storage Manager administrator grants you authority, you can delete entire file spaces from the server.

About this task

You cannot delete individual backup copies that are kept on the server. When you delete a file space, you delete all the files, both backup copies and archive copies,

that are contained within the file space. For example, if you delete the file space for your C drive, you are deleting every backup copy for every file on that disk and every file that you archived from that disk.

Attention: Carefully consider what you are doing before you delete a file space.

You can delete file spaces using the GUI or the command-line client. To delete network-attached storage (NAS) file spaces, use the web client or command-line client.

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

Procedure

1. From the main window, click **Utilities** → **Delete Filespaces**.
2. Select the file spaces you want to delete.
3. Click **Delete**. 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 341

“Delete Filespace” on page 677

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 system state data, a file space, a directory, or a 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:

Related reference:

"Backup Image" on page 644

Restore NAS file systems

You restore NAS file system images using the web client or command line interface.

You can restore full or differential NAS file system images that were backed up previously. If you restore a differential image, Tivoli Storage Manager automatically restores the full backup image first, followed by the differential image. It is not necessary for a client node to mount a NAS file system to perform backup or restore operations on that file system.

Related concepts:

"Web client configuration overview" on page 33

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.

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.

Restoring NAS files and directories using the web client

You can use the `toc` option with the `include.fs.nas` option in your client options file to specify whether the client saves Table of Contents (TOC) information for each file system backup.

About this task

If you save TOC information, you can use web client to examine the entire file system tree and select files and directories to restore. Creation of a TOC requires that you define the `TOCDESTINATION` attribute in the backup copy group for the management class to which this backup image is bound. Note that TOC creation requires additional processing, network resources, storage pool space, and possibly a mount point during the backup operation. If you do not save TOC information, you can still restore individual files or directory trees using the `RESTORE NODE` server command, provided that you know the fully qualified name of each file or directory and the image in which that object was backed up.

To restore NAS files and directories:

Procedure

1. Click **Restore** from the main window. The Restore window appears.
2. Expand the directory tree if necessary. To expand a node in the tree, click the plus sign (+) next to an object in the tree. 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 appear 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 display the **File Level** object.
4. Expand the **File Level** object to display the volumes, directories, and files that were last backed up. When you expand the volume object, and complete TOC information is available on the server for the latest backup, the Load Table of Contents dialog appears. If complete TOC information is not available for the latest backup, no objects appear below the volume object. The next step explains how to display objects from backups other than the latest backup. Complete TOC information is provided if you performed either of the following operations: (1) A differential image backup with TOC information and its corresponding full image backup with TOC information, or (2) A full image backup with TOC information.

5. Click the selection boxes next to the directories or files that you want to restore.
 - a. If you want to restore files from a NAS image that was backed up on a particular date or display files from several older versions, highlight the volume you want to restore and click the **Point In Time** button.
 - b. If you select **Use a Point in Time Date** in the Point in Time Restore windows, files from the image backed up on that date, and if it is a differential image, files from its corresponding full image appear under the **File Level** object.
 - c. If you click **Use Selected Images** in the Point in Time Restore window, the Selected Images window appears for you to select images. The contents of the selected images appear in the **File Level** object.
6. Click **Restore**. The Restore Destination window appears. Enter the information in the Restore Destination window. If you choose to restore to a different destination, you can only restore one volume at a time to a different destination.
7. Click **Restore**. The NAS Restore Task List window displays the restore processing status and progress bar. If there is a number next to the progress bar, it indicates the size of the restore, if known. After the restore completes, the NAS Restore Report window displays processing details. 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 (either workstation or NAS).
- To view information about objects in a NAS node, highlight the object and select **View > File Details** from the menu.
- To delete NAS file spaces, select **Utilities > Delete Filespaces**. You can delete both workstation and remote objects.

Related reference:

"Toc" on page 563

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 30. 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 714

Table 30. NAS options and commands (continued)

Option or command	Definition	Page
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 700
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 706
restore nas	Restores the image of a file system belonging to a Network Attached Storage (NAS) file server.	"Restore NAS" on page 751
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 692
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 670
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 677

A NAS file system specification uses the following conventions:

- Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.
- NAS file system designations on the command line require brace delimiters {} around the file system names, such as: {/vol/vol0}.

Note: When you initiate a NAS restore operation using the command line client or the web client, the server starts a process to initiate, control, and monitor the operation. It might take several moments before you notice progress at the 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.

Chapter 6. Archive and retrieve your data (Windows)

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.

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

All the primary archive and retrieve procedures also apply to the web client, except for the following functions:

- Preferences editor
- Setup wizard

You can complete the following primary archive and retrieve tasks:

- “Archiving data with the GUI” on page 240
- “Archive data examples by using the command line” on page 241
- “Deleting archive data” on page 244
- “Retrieving archives with the GUI” on page 245
- “Retrieve archive copies by using the command line” on page 245

Related concepts:

“When to back up and when to archive files” on page 129

Related tasks:

“Starting a web client session” on page 115

Archive files

To archive files, select the files that you want to archive. You can select the files by name or description, or select them from a directory tree.

Your administrator might set up schedules to automatically archive certain files on your workstation. The following sections contain information about how to archive files without using a schedule.

You must assign an archive description for all archived files. An archive description identifies data through a meaningful description that you can use later to identify files and directories. You can enter as many as 254 characters to describe your archived data. If you do not enter a description, the following default archive description is assigned:

Archive Date: mm/dd/yyyy

where mm/dd/yyyy is the current date.

When you select the archive function from the backup-archive GUI, a list of all previously used archive descriptions are displayed. You can assign these archive descriptions to future archives.

Incremental backup might recall migrated files, while selective backup and archive always recall migrated files, if you do not use the skipmigrated option.

Related concepts:

➡ Options for backing up migrated files: skipmigrated, checkreparsecontent, stagingdirectory

Related tasks:

“Setting the client scheduler process to run as a background task and start automatically at startup” on page 252

Snapshot backup or archive with open file support

If open file support has been configured, the backup-archive client runs a snapshot backup or archive of files that are locked (or "in use") by other applications.

The snapshot allows the archive to be taken from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to the file system are not included in the archive. You can set the snapshotproviderfs parameter of the include.fs option to **none** to specify which drives do not use open file support.

Note:

1. You can use the include.fs option to set snapshot options on a per file system basis.
2. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with FAT, FAT32, NTFS, or ReFS file systems. This support includes SAN-attached volumes that meet these requirements.
3. If the client is unable to create a snapshot, failover to non-OFS backup occurs; the same backup support that would be done if the OFS feature was not installed.
4. To enable open file support in a cluster environment all workstations in the cluster should have the OFS feature configured.
5. When using the open file support feature with VSS, the client adds the snapshot volume name to the path of the objects being processed. The snapshot volume name can be up to 1024 bytes. The complete path (snapshot volume name plus object path) can be 8192 bytes or less.

For information about open file support restrictions and issues, search for *TSM Client Open File Support (OFS)* at the IBM support website.

Related concepts:

Chapter 11, “Processing options,” on page 295

Related tasks:

“Configuring Open File Support” on page 83

Archiving data with the GUI

You can archive specific files or entire directories from a directory tree. You can also assign a unique description for each group of files you archive (archive package).

About this task

To archive your files, complete the following steps:

Procedure

1. Click **Archive** in the GUI main window. The Archive window displays.

2. Expand the directory tree by clicking the plus sign (+) or a folder icon in the tree. To search or filter files, click the **Search** icon from the toolbar.
3. Enter a description, accept the default description, or select an existing description for your archive package in the **Description** field.
4. To modify specific archive options, click **Options**. Any options that you change are effective during the current session only.
5. Click **Archive**. The Archive Status window displays the progress of the archive operation.

Archive data examples by using the command line

You can archive data when you want to preserve copies of files in their current state, either for later use or for historical or legal purposes.

You can archive a single file, a group of files, or all the files in a directory or subdirectory. After you archive a file, you can delete the original file from your workstation. Use the **archive** command to archive files.

The following table shows examples of how to use the **archive** command to archive objects.

Table 31. Command-line archive examples

Task	Command	Considerations
Archive all files in the c:\plan\proj1 directory with a file extension of .txt.	dsmc archive c:\plan\proj1*.txt	Use wildcards to archive more than one file at a time.
Archive all files in the c:\small\testdir directory and delete the files on your workstation.	dsmc archive c:\small\testdir* -deletefiles	Retrieve the archived files to your workstation whenever you need them again. For more information about the deletefiles option, see "Deletefiles" on page 361.
Archive the c:\proj1\h1.doc file and the c:\proj2\h2.doc file	dsmc archive c:\proj1\h1.doc c:\proj2\h2.doc	You can specify as many files to be archived as the resources and operating system limits permit. Separate the files to be archived with a space. For more information about the filelist option, see "Filelist" on page 414.
Archive a list of files in the c:\filelist.txt file.	dsmc archive -filelist=c:\ filelist.txt	Use the filelist option to process a list of files. For more information about the filelist option, see "Filelist" on page 414.
Archive the a:\ch1.doc file and assign a description to the archive.	dsmc archive a:\ch1.doc -description="Chapter 1, first version"	If you do not specify a description with the archive command, the default is Archive Date:x, where x is the current system date. For more information about the description option, see "Description" on page 362.
Archive all the files in the d:\proj directory and its subdirectories.	dsmc archive d:\proj\ -subdir=yes	For more information about the subdir option, see "Subdir" on page 548.
Use the v2archive option with the archive command to archive only files in the c:\relx\dir1 directory.	dsmc archive c:\relx\dir1\ -v2archive	Tivoli Storage Manager archives only files in the c:\relx\dir1 directory. Directories that exist in the path are not processed. For more information about the v2archive option, see "V2archive" on page 570.

Table 31. Command-line archive examples (continued)

Task	Command	Considerations
Use the archmc option with the archive command to specify the available management class for the policy domain to which you want to bind your archived files.	<code>dsmc archive -archmc=RET2YRS c:\plan\proj1\ budget.jan*</code>	For more information about the archmc option, see “Archmc” on page 323. For more information about management classes, see Chapter 9, “Storage management policies,” on page 267.
Assume that you initiated a snapshot of the C:\ drive and mounted the snapshot as the logical volume \\florence\c\$\snapshots\ snapshot.0. You archive the c:\dir1\sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name C:\.	<code>dsmc archive c:\dir1\sub1* -subdir=yes -snapshotroot=\\florence\c\$\snapshots\snapshot.0</code>	For more information, see “Snapshotroot” on page 537.

Related reference:

“Archive” on page 633

Associate a local snapshot with a server file space (Windows)

You can associate the data on the local snapshot with the real file space data that is stored on the Tivoli Storage Manager server.

To associate the data on the local snapshot with the real file space data on the Tivoli Storage Manager server, use the **snapshotroot** option with the **archive** command, with a vendor-acquired application that provides a snapshot of a logical volume.

The **snapshotroot** option cannot provide any facilities to take a volume snapshot, it can manage only data that is created by a volume snapshot.

Related reference:

“Snapshotroot” on page 537

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.

About this task

This is useful when the workstation responsible for performing the archive can change over time, such as with a cluster. The **asnodename** option also allows data to be restored from a different system than the one that performed the backup. Use the **asnodename** option with the appropriate command to back up, archive, restore, and retrieve data under the target node name on the Tivoli Storage Manager server.

Tivoli Storage Manager FastBack clients are also backed up using client node proxy.

To enable this option, follow these steps:

1. Install the backup-archive client on all nodes in a shared data environment.

2. Register each node with the Tivoli Storage Manager server, if it does not exist. Register the common target node name to be shared by each of the agent nodes used in your shared data environment.
3. Register each of the nodes in the shared data environment with the Tivoli Storage Manager server. This is the agent node name that is used for authentication purposes. Data is not stored using the node name when the `asnodename` option is used.
4. The Tivoli Storage Manager administrator must grant proxy authority to all nodes in the shared environment to access the target node name on the Tivoli Storage Manager server, using the **GRANT PROXYNODE** server command.
5. Use the **QUERY PROXYNODE** administrative client command to display the client nodes of the authorized user, granted by the **GRANT PROXYNODE** command.

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

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

Follow these steps to enable multinode operation from the GUI:

1. Verify that the client node has proxy authority to a target node (or authorized to act as the target node) using the **QUERY PROXYNODE** administrative client command.
2. Select **Edit > Preferences** to open the preferences window.
3. Select the **General** tab and fill in the **As Node Name** field with the name of the proxy authorized target node.
4. Click **Apply** and then **OK** to close the preferences window.

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

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

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

Considerations for 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.
- All agent nodes in the multiple node environment must be of the same platform type.
- Do not use target nodes as traditional nodes. Use them only for multiple node processing.
- You cannot perform a system object or system state backup or restore.
- You cannot access another node (either from GUI drop down or use of the `fromnode` option).
- You cannot use the `clusternode` option.
- You cannot perform NAS backup or restore.

Related reference:

“Asnodename” on page 324

“Session settings and schedules for a proxy operation” on page 326

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 672

Retrieve archives

Select the **Retrieve** function to recover an archive copy of a file or a directory.

Note: When you retrieve a directory, its modification date and time is set to the date and time of the retrieve, not to the date and time the directory had when it was archived. This is because a retrieve operation retrieves the directories first, and then adds the files to the directories.

You can also retrieve archive copies from the directory tree, filter the directory tree, and retrieve archive copies of files owned by someone else. To do any of these, click the **Retrieve** button on the main window of the backup-archive client GUI and follow the directions provided in the task help of the GUI.

Important: When you retrieve a file without any specifications, and more than one version of the archive copy exists on the server, all of the copies are retrieved. 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:

“Duplicate file names” on page 193

Retrieving archives with the GUI

You can retrieve your archived files with the backup-archive client GUI.

Procedure

1. Click **Retrieve** on the 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 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 of 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. You can retrieve files to a directory or drive other than the one from where they were originally archived. You can also select how much of the parent directory structure is re-created at the retrieve location.
11. Click **Retrieve**. The Retrieve Status window displays the processing status.

Retrieve archive copies by using the command line

You retrieve a file when you want to return an archive copy from the server to your workstation. Some examples of how to retrieve archived files by using the command line are shown.

You can retrieve a single file, a group of files, or all the files in a directory or subdirectory. When you retrieve a file, the Tivoli Storage Manager server sends you a copy of that file. The archived file remains in storage.

Use the **retrieve** command to retrieve files. The following table shows examples of using the **retrieve** command.

Table 32. Command line examples of retrieving archives

Task	Command	Considerations
Retrieve the c:\doc\h2.doc file to its original directory.	dsmc retrieve c:\doc\h2.doc	If you do not specify a destination, the files are retrieved to their original location.
Retrieve the c:\doc\h2.doc file under a new name and directory.	dsmc retrieve c:\doc\h2.doc c:\proj2\h3.doc	None
Retrieve all files that are archived with a specific description to a directory named retr1 at a new location	dsmc retrieve c:* d:\retr1\ -sub=yes -desc="My first archive"	None
Retrieve all files from the c:\projecta directory that end with the characters .bak to the c:\projectn directory.	dsmc retrieve c:\projecta*.bak c:\projectn	None
Use the pick option display a list of archives from which you can select files to retrieve.	dsmc retrieve c:\project* -pick	For more information about the pick option, see "Pick" on page 481.
Retrieve a file that is originally archived from the diskette that is labeled <i>workathome</i> on the a: drive, to a diskette in the a: drive labeled <i>extra</i> .	dsmc retrieve {workathome}\doc\ h2.doc a:\doc\h2.doc	If you are retrieving a file to a disk that has a different label other than the disk from which the file was archived, use the file space name (label) of the archive disk rather than the drive letter.
Retrieve the c:\doc\h2.doc file to its original directory on the workstation, named <i>star</i> .	dsmc retrieve c:\doc\h2.doc \\star\c\$\ To retrieve the file to <i>star</i> , which was renamed <i>meteor</i> , enter: dsmc retrieve \\star\c\$\ doc\h2.doc \\meteor\c\$\ You can also enter: dsmc retrieve \\star\c\$\ doc\h2.doc c:\ This example is valid because if the workstation name is not included in the specification, the local workstation is assumed (<i>meteor</i> , in this case).	For the purposes of this manual, the workstation name is part of the file name. Therefore, if you archive files on one workstation and you want to retrieve them to another workstation, you must specify a destination. This requirement is true even if you are retrieving to the same physical workstation, but the workstation has a new name.

Related reference:

"Retrieve" on page 766

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, imagebackup, imagerestore, command, and macro. The scheduled action that is most frequently used is incremental with the **objects** parameter left undefined. With this setting, the Tivoli Storage Manager client performs a domain incremental backup of all drives defined by the client domain option. A schedule definition using the **command** action allows an operating system command or shell script to be executed. When automating tasks for Tivoli Storage Manager for Data Protection clients, you must use **command** action schedule definitions, which invoke the command-line utilities for 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='c:\home\proj1\Some file.doc'
objects='c:\home\proj1\Some file.doc' "c:\home\Another file.txt"
c:\home\noblanks.txt'
objects='c:\home\My Directory With Blank Spaces\'
objects='c:\Users\user1\Documents\Some file.doc'
objects='c:\Users\user1\Documents\Some file.doc'
"c:\Users\user5\Documents\ Another file.txt" c:\Users\user3\Documents\noblanks.txt'
objects='c:\Users\user1\My Directory With Blank Spaces\'
```

This syntax ensures that a file specification containing a space, such as c:\home\proj1\Some file.doc, is treated as a single file name, and not as two separate files (c:\home\proj1\Some, and file.doc)

The following examples show how to delimit schedule **options** parameters when file specifications contain space characters:

```
options='-preschedulecmd="c:\home\me\my files\bin\myscript"
-postschedulecmd="c:\home\me\my files\bin\mypostscript" -quiet'
options='-presched="c:\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 114

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
managedservices	X		
maxcmdretries	X		SET MAXCMDRETRIES command
maxschedsessions		X	
postschedulecmd, postnschedulecmd	X		
preschedulecmd, prenschedulecmd	X		
querschedperiod	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.opt file. The Tivoli Storage Manager server can also define some options in a client options set, or as part of the options parameter of the schedule definition. The Tivoli Storage Manager server can also set some options globally for all clients. By default, the client setting for these options is honored. If the global override on the Tivoli Storage Manager server is set, the client setting for the option is ignored. Defining client options as part of the schedule definition is useful if you want to use specific options for a scheduled action that differ from the option settings normally used by the client node, or are different for each schedule the node 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.exe command might exit with return code 3, if widget.exe ran successfully. Therefore, it is possible that the preschedulecmd or postschedulecmd command might end with a nonzero return code and 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 should exit with return code 0 if the command was successful; otherwise it should exit with a nonzero return code. The logic for a script that runs widget.exe might look like this example:

```

run 'widget.exe'
    if lastcc == 3
        exit 0
    else
        exit 1
    fi

```

Related reference:

“Postschedulecmd/Postnschedulecmd” on page 484

“Preschedulecmd/Prenschedulecmd” on page 487

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.

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 `managedservices` stanza set).

On Windows platforms, the scheduler and the client acceptor run as services. You can create and manage these services by using either the setup wizard or the Tivoli Storage Manager Client Service Configuration Utility, `dsmcutil.exe`.

- To start the setup wizard, select **Utilities > Setup Wizard** in the backup-archive GUI and select a **Help me configure** option for the appropriate service. Follow the prompts to install, configure, and start the service.
- To start the Client Service Configuration Utility, open a command prompt window and issue the following command to change to the directory that contains `dsmcutil.exe`:

```
cd /d "c:\program files\tivoli\tsm\baclient"
```

Use **dsmcutil** to manage the client acceptor service or the scheduler service. Full documentation on how to use **dsmcutil** is available by entering `dsmcutil help`.

The client scheduler can be managed by the client acceptor. When setting up scheduler services to run with client acceptor management, two services must be created: the scheduler service and the client acceptor service. When you install the client acceptor service with **dsmcutil.exe**, use the `/cadschedname:` parameter to identify which scheduler service the client acceptor manages. If you use the setup wizard to install the scheduler, you can select the **Use the CAD to manage the scheduler** check box, which automatically creates both services and associates them.

Using the Client Service Configuration Utility, you can use either of the following methods:

Client acceptor-managed method

1. In your client options file (`dsm.opt`), either set the `managedservices` option to **schedule** or **schedule webclient**.
2. In your client options file (`dsm.opt`), set the `passwordaccess` option to **generate**.
3. Create the scheduler service:

```
dsmcutil inst /name:"TSM Client Scheduler" /node:tsmclient1  
/password:secret /autostart:no /startnow:no
```
4. Create the client acceptor and associate scheduler service with the client acceptor:

```
dsmcutil inst CAD /name:"TSM Client Acceptor" /cadschedname:  
"TSM Client Scheduler" /node:tsmclient1 /password:secret /autostart:yes
```
5. Manually start the client acceptor service:

```
net start "TSM Client Acceptor"
```

Traditional method

1. In your client options file (`dsm.opt`), either remove the `managedservices` entirely (it defaults to **webclient**) or set it to **webclient**.
2. In your client options file (`dsm.opt`), set the `passwordaccess` option to **generate**.
3. Create the scheduler service:

```
dsmcutil inst /name:"TSM Client Scheduler" /node:tsmclient1  
/password:secret /autostart:yes
```

To increase the reliability of the client scheduler service on Windows, set the services to automatically recover from a failure as follows:

- Start the Windows services management console (**Start > Settings > Control Panel > Administrative Tools > Services**)
- Right-click the **TSM Client Scheduler** service and select **Properties**.
- Click the **Recovery** tab.
- Define the recovery action as **Restart the service** for first, second, and subsequent failures.

If you are using the client acceptor to manage the scheduler, you must set the recovery properties for the **TSM Client Acceptor** service, but leave the recovery settings for the **TSM Client Scheduler** service as **Take No Action** for the first, second, and subsequent failures. The same recovery settings can also be defined to increase the reliability of the **TSM Journal Service**.

Related reference:

"Cadlistenonport" on page 338

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 33 on page 255 displays sample classic **query schedule** output.

Table 33. Sample classic query schedule output

Schedule Name:	DAILY_INC
Description:	Daily System-wide backup
Schedule Style:	Classic
Action:	Incremental
Options:	QUIET
Objects:	
Priority:	1
Next Execution:	30 minutes
Duration:	4 Hours
Period:	1 Day
Day of Week:	Any
Month:	
Day of Month:	
Week of Month:	
Expire:	Never
Schedule Name:	WEEKLY_INC
Description:	Weekly backup for project files
Schedule Style:	Classic
Action:	Incremental
Options:	QUIET
Objects:	e: f:
Priority:	1
Next Execution:	60 minutes
Duration:	8 Hours
Period:	7 Days
Day of Week:	Friday
Month:	
Day of Month:	
Week of Month:	
Expire:	Never

The schedule name, **WEEKLY_INC**, starts a weekly incremental backup on the e: and f: drives.

The schedule name, **DAILY_INC**, starts a daily incremental backup. The next incremental backup 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 34 on page 256 displays sample enhanced **query schedule** output.

Table 34. 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.

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
```

When you run the **schedule** command as a service, output from scheduled commands displays in the application event log. Output is also directed to the `dsmsched.log` file in the current directory unless you change the path and file name using the `schedlogname` option. The amount of detail is determined by whether *verbose* or *quiet* is set in the `dsm.opt` file. The scheduler service also posts messages to the Windows event log.

After scheduled work is performed, check the schedule log to verify that all work completed successfully.

When a scheduled command is processed the schedule log contains the following entry:

```
Scheduled event eventname completed successfully
```

If the scheduled event does not complete successfully, you 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” on page 260

Examples: event logs

The scheduler service logs information into the application event log and provides an event identification (event ID) number for each event in the log. This topic shows examples of events that are logged to the application event log.

Scheduler service

Event 4097 (informational message)

Example 1:

Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4097
Date: 10/31/2002
Time: 8:29:57 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
TSM 515 Scheduler halted.

Example 2:

Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4097
Date: 10/31/2002
Time: 8:29:57 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Scheduler Terminated, service ending.

Example 3:

Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4097
Date: 10/31/2002
Time: 8:29:56 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
TSM Client Scheduler 'TSM 515 Scheduler'
Started.

Example 4:

Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4097
Date: 10/31/2002
Time: 8:29:56 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Starting Scheduler.

Example 5:

Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4097
Date: 10/30/2002
Time: 8:06:09 PM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Incremental backup of volume '\\MIKEDILE\C\$'

Event 4098 (warning message)

Example 1:

Event Type: Warning
Event Source: AdsmClientService
Event Category: None
Event ID: 4098
Date: 10/31/2002
Time: 8:29:56 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Error Initializing TSM Api, unable to verify
Registry Password, see dserror.log.

Example 2:

Event Type: Warning
Event Source: AdsmClientService
Event Category: None
Event ID: 4098
Date: 9/20/2002
Time: 6:20:10 PM
User: DILE\Administrator
Computer: MIKEDILE
Description:
ANS1802E Incremental backup of '\\mikedile\
c\$' finished with 3 failure

Event 4099 (error message)

Example 1:

Event Type: Error
Event Source: AdsmClientService
Event Category: None
Event ID: 4099
Date: 9/17/2002
Time: 6:53:13 PM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Scheduler exited with a result code of 4.

Example 2:

Event Type: Error
Event Source: AdsmClientService
Event Category: None
Event ID: 4099
Date: 9/17/2002
Time: 6:27:19 PM
User: DILE\Administrator
Computer: MIKEDILE
Description:

ANS4987E Error processing '\\mikedile\e\$\
tsm520c\client\winnt\mak \dsmwin32.ncb':
the object is in use by another process

Event 4100 (scheduler command message)

Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4100
Date: 10/31/2002
Time: 8:29:56 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Next Scheduled Event Obtained from Server
SNJEDS1 (MVS):

Schedule Name: NIGHTLY_BACKUP
Action: Incremental
Objects: (none)
Options: (none)
Server Window Start: 19:00:00 on 10/31/2002

Event 4101 (backup or archive statistics)

Displays backup and archive statistics, which might be useful in
determining the success or failure of a command.

Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4101
Date: 10/30/2002
Time: 8:29:21 PM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Backup/Archive Statistics for Schedule Backup
NIGHTLY_BACKUP :

Total number of objects inspected: 158,688
Total number of objects backed up: 2,486
Total number of objects updated: 0
Total number of objects rebound: 0
Total number of objects deleted: 0
Total number of objects expired: 12
Total number of objects failed: 0
Total number of bytes transferred: 1.15 GB
Data transfer time: 104.35 sec
Network data transfer rate: 11,564.84 KB/sec
Aggregate data transfer rate: 866.99 KB/sec
Objects compressed by: 100%
Elapsed processing time: 00:23:11

Event 4103 (backup-archive client service startup parameters)

Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4103
Date: 10/31/2002
Time: 8:29:56 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Backup/Archive Client Service Startup
Parameters:

Service Name : TSM 515 Scheduler

```
Last Update : Oct 14 2002
Client PTF Level : 5.1.5.2
Service Directory : D:\Program Files\
Tivoli\TSM515\baclient
Client Options File : E:\users\mikedile\
logfiles\dsm.opt
Client Node : MIKEDILE
Comm Method : (default or obtained from
client options file)
Server : (default or obtained from client
options file)
Port : (default or obtained from client
options file)
Schedule Log : E:\users\mikedile\logfiles\
dsmsched.log
Error Log : E:\users\mikedile\logfiles\
dsmerror.log
MS Cluster Mode : (default or obtained
from client options file)
```

Journal based backup service events

```
4097: Informational message
4098: Warning message
4099: Error message
4100: Journal Based Backup service file monitor parameters
4101: Journal Based Backup service database parameters
4102: Journal Based Backup Service configuration parameters
```

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 309

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 513

“`Schedrestretrdisbled`” on page 521

“`Srvprepostscheddisbled`” on page 540

“`Srvprepostsnapdisabled`” on page 541

Change processing options used by the scheduler service

When you configure the Tivoli Storage Manager central-scheduling services (the scheduler, the client acceptor, or the remote client agent), some of the processing options that you specify are defined in the Windows registry.

The following options can also be specified in the client options file (`dsm.opt`).

- `nodename`
- `httpport`
- `tcpserveraddress`
- `tcpport`
- `webports`

When the client scheduler runs as a foreground process using the **`dsmc sched`** command, the options in the client options file are used. However, when the scheduler runs as a Windows service, the options in the registry are used instead. If you are using the scheduler service and change an option in the `dsm.opt` file, you must update the corresponding value in the registry as well.

To update the Windows registry value:

Use the Setup wizard in the client GUI. For more information, see “Configuring the scheduler” on page 36.

Alternatively, you can use the `dsmcutil` utility to change the registry value. For example: `dsmcutil update scheduler /name: <service name> /node: <new node name> /password: <new node password>`.

Note: After updating the registry, you must restart the scheduler service for the changes to take effect. If you are using client acceptor daemon-managed scheduling this is not necessary because the scheduler is restarted by the client acceptor daemon for each backup.

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 drives on your client system must be backed up daily, and that one drive containing critical data must be backed up hourly. In this case, you would need to define two schedules to handle this requirement. To avoid conflict between the hourly and daily backup schedule, the *starttime* of each schedule needs to be varied.

In certain cases, it is necessary to run more than one scheduler process on a system. Multiple processes require a separate options file for each process and must contain the following information:

- Define a unique node name for each process
- Specify unique schedule and error logs for each process
- When running in prompted mode, you must use the `tcpclientport` option to specify a unique port for each process.

Note: When the scheduler runs as a service, processing options specified in the Windows registry override the same options specified in the client options file.

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.

You must create a separate service for each schedule process. If you are using the client acceptor daemon to manage the scheduler, a client acceptor daemon service and schedule service are required for each schedule. The following is an example of setting up two schedule processes to be managed by the client acceptor daemon:

```
dsmcutil inst /name:"TSM Client Scheduler1" /optfile:"c:\tsm\dsm.opt1"  
/node:tsmcli_sched1 /password:secret /autostart:no /startnow:no
```

```
dsmcutil inst CAD /name:"TSM Client Acceptor1" /optfile:"c:\tsm\dsm.opt1"  
/cadschedname:"TSM Client Scheduler1" /node:tsmcli_sched1 /password:secret  
/autostart:yes
```

```
dsmcutil inst /name:"TSM Client Scheduler2" /optfile:"c:\tsm\dsm.opt2"  
/node:tsmcli_sched2 /password:secret /autostart:no /startnow:no
```

```
dsmcutil inst CAD /name:"TSM Client Acceptor2" /optfile:"c:\tsm\dsm.opt2"  
/cadschedname:"TSM Client Scheduler2" /node:tsmcli_sched2 /password:secret  
/autostart:yes
```

Unique option files are required for each schedule instance, and must be identified at the time of service creation:

Option file #1 (c:\tsm\dsm.opt1)

```
tcps          tmserv1.example.com
nodename      tsmcli_sched1
passwordaccess generate
schedlogname  c:\tsm\dsmsched1.log
errorlogname  c:\tsm\dsmerror1.log
schedmode     prompted
tcpclientport 1507
domain        h:
managedservices schedule
```

Option file #2 (c:\tsm\dsm.opt2)

```
tcps          tmserv1.example.com
nodename      tsmcli_sched2
passwordaccess generate
schedlogname  c:\tsm\dsmsched2.log
errorlogname  c:\tsm\dsmerror2.log
schedmode     prompted
tcpclientport 1508
domain        i:
managedservices schedule
```

Related concepts:

“Change processing options used by the scheduler service” on page 261

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 35 describes the return codes and their meanings.

Table 35. 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. <ul style="list-style-type: none">• 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 271.
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 drive from being processed yields return code 12.

Table 35. Client return codes and their meanings (continued)

Code	Explanation
<i>other</i>	<p>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.</p> <p>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.</p>

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 c:\MyTools\* -subdir=yes
incremental c:\MyPrograms\TestDriver\* -subdir=yes
archive e:\TSM\* -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.

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.

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.

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 273

“Retention grace period” on page 276

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 36 shows the default values for the backup and archive copy groups in the standard management class.

Table 36. 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.

Journal-based incremental backup differs from the traditional full incremental backup because Tivoli Storage Manager does not enforce non-default copy frequencies (other than 0).

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.

Note: For all supported Windows 32-bit versions: During an image backup, the static copy serialization value is no longer controlled by the server management class, but is instead controlled directly from the client, using the `snapshotproviderimage` option.

- **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.

Important: For all supported Windows 32-bit versions: During an image backup, the dynamic copy serialization value is no longer controlled by the server management class, but is instead controlled directly from the client, using the `snapshotproviderimage` option.

- **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 concepts:

“Open file support for backup operations” on page 123

Related tasks:

“Configuring Open File Support” on page 83

Related reference:

“Snapshotproviderimage” on page 536

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.
- The attributes, except for the archive attribute, are different.
- If only the metadata changes (such as access permissions), the client might back up only the metadata.

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 322

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 53

Related reference:

“Deduplication” on page 360

“Enablededupcache” on page 389

“Exclude options” on page 400

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 271

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 `incl exc1` 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 c:\adsm\proj2\costs\* budget
```

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

```
include ?:\...\* managall
```

The following examples show how to assign a management class to files:

```
exclude ?:\...\*.sno
include c:\winter\...\*.ice      mcweekly
include c:\winter\december\*.ice mcdaily
include c:\winter\january\*.ice  mcmonthly
include c:\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
```

Related reference:

“Dirmc” on page 367

“Include options” on page 430

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 c:\plan\proj1\budget.jan
```

Select a management class for directories

If the management class in your active policy set containing the longest "Retain only version" (REONLY) setting meets your backup requirements for directories, it might not be necessary to take any action to associate directories with that management class. The management class association is done automatically when it backs up your directories.

If there is more than one management class with the longest REONLY 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 778

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. Tivoli Storage Manager Client Service Configuration Utility

The following client services can be installed when you install the backup-archive client or by using the Tivoli Storage Manager Client Service Configuration Utility after the backup-archive client is installed: (1) Backup-Archive Scheduler Service, (2) Client Acceptor Service (CAD), (3) Remote Client Agent Service, (4) Journal Engine Service, and (5) VMware Backup Tools Service.

For more information about using the Tivoli Storage Manager Client Service Configuration Utility to install client services, see the related information about using the **dsmcutil** command.

Related concepts:

“**dsmcutil** command” on page 283

Install the backup-archive scheduler service

You can use either the backup-archive client GUI or the TSM Client Service Configuration Utility to install the scheduler.

About this task

- From the backup-archive client GUI, click **Utilities**, and then click **Setup Wizard**. Select the **Help me configure the TSM Client Scheduler** option.
- If you have an account that belongs to the Administrator/Domain Administrator group, you can use the **TSM Client Service Configuration Utility** to configure client services on both local and remote Windows workstations.

Using the Client Service Configuration Utility (Windows)

This section provides the steps for using the Client Service Configuration Utility to automate backups, manage existing scheduler services, create a new scheduler, and associate a client acceptor to manage the scheduler.

About this task

This example illustrates the use of the Tivoli Storage Manager scheduler.

When the backup-archive client is registered with a Tivoli Storage Manager server, the procedure involves the following steps:

Procedure

1. **On the server:**
 - a. Define a schedule for the policy domain to which the backup-archive client is registered.
 - b. Associate the backup-archive client node to the defined schedule.
2. **On the backup-archive client:**
 - a. Install the scheduler as a Windows service for the backup-archive client.
 - b. Start the scheduler service installed for the backup-archive client.

Examples: Automating backups

This section provides examples of how to automate your backups.

About this task

The example below uses the following assumptions:

- The backup-archive client is registered to a Tivoli Storage Manager server with a node name of **mars** and a password of **marspswd** in policy domain **bacliwnt**.
- The event to be scheduled is a daily incremental backup of file systems on client workstations. The backup begins between 9:00 and 9:15 pm.
- The backup-archive client is installed to the `c:\program files\tivoli\tsm\baclient` directory.
- The communication parameters in the backup-archive client options file (`dsm.opt`) are appropriate for the Tivoli Storage Manager server.

On the server:

Procedure

1. Enter the following command on the server console or from an administrative client to define the schedule: `def sched bacliwnt wnt_daily_incr desc="Daily Incremental Backup" priority=2 starttime=21:00 duration=15 durunits=minutes period=1 perunits=days dayofweek=any`. The administrative client does not have to be running on the same system as the Tivoli Storage Manager server. The following message is displayed:

```
ANR2500I Schedule WNT_DAILY_INCR defined in policy domain BACLIWNT.
```

2. To associate the backup-archive client to this schedule, issue the following command: `define association bacliwnt wnt_daily_incr mars`. The following message is displayed:

```
ANR2510I Node MARS associated with schedule WNT_DAILY_INCR in policy domain BACLIWNT.
```

At this point, a schedule that performs an incremental backup is defined on the Tivoli Storage Manager server. The schedule starts around 9:00 pm. The schedule is re-executed once a day and can start on any day of the week. If you want to confirm that the schedule and association are set correctly, you can use the **Query Schedule** command.

Results

On the backup-archive client:

This example assumes that you installed the backup-archive client in the `c:\program files\tivoli\tsm\baclient` directory. It also assumes that the options files in each of these directories are updated so that the communication parameters point to the Tivoli Storage Manager server.

1. Login using an account with administrative privileges.
2. Open a command prompt window and issue the following command: `cd /d "c:\program files\tivoli\tsm\baclient"`. If the path contains a space, for example `c:\program files\tivoli\tsm\baclient`, enclose the name in double quotation marks.
3. In the window, issue the following command: `dsmcutil inst scheduler /name:"TSM Client Scheduler" /node:mars /password:marspswd /clientdir:"c:\program files\tivoli\tsm\baclient" /optfile:"c:\program files\tivoli\tsm\baclient\dsm.opt" /autostart:yes`. Your system is now ready to run automatic daily incremental backups. The */autostart:yes* option specifies that the scheduler service starts automatically each time the system is rebooted. You can use the */startnow:[Yes|No]* option to specify whether to start

the scheduler service after executing the command; the default is *Yes*. If you specify */startnow:No* you must start the service manually using the services control panel applet on NT, or issue the following command: `net start "TSM Client Scheduler"`.

4. The scheduler uses the backup-archive client options file to validate the node and password, and to contact the server for schedule information. This example assumes that the `dsm.opt` file is updated so that the communication parameters point to the Tivoli Storage Manager server. If you see the following message: A communications error occurred connecting to the TSM server. You should ensure that the options file contains entries that point to the correct Tivoli Storage Manager server. You should also ensure that the server is running. Use the `dsmcutil update` command to correct one of the parameters which was incorrectly specified during the `dsmcutil install`. For example, to update the client directory and options file for the specified scheduler service, enter:
`dsmcutil update scheduler /name:"TSM Central Scheduler Service"
/clientdir:c:\tsmptf3\baclient /optfile:c:\tsmptf3\baclient\dsm.opt.`
Then reissue the `net start "TSM Client Scheduler"` command.

Note:

1. If any changes that affect the scheduler service are made to the backup-archive client options file, the scheduler service has to be restarted. If you are using CAD managed scheduling this is not necessary since the scheduler is restarted by the CAD for each backup and the changes are picked up. An example of this is the Tivoli Storage Manager server address or the schedule mode. This can be done by issuing the following commands: `net stop "TSM Client Scheduler"` and then `net start "TSM Client Scheduler"`.
2. The `dsmsched.log` file contains status information for the Tivoli Storage Manager scheduler service. In this example, the file is located in this path: `c:\program files\tivoli\tsm\baclient\dsmsched.log`. You can override this file name by specifying the *schedlogname* option in the options file, `dsm.opt`.
3. Output from scheduled commands is sent to the log file. After scheduled work is performed, check the log to ensure the work completed successfully. When a scheduled command is processed the schedule log might contain the following entry: Scheduled event *eventname* completed successfully. This is merely an indication that the scheduled command associated with the *eventname* was successfully issued. No attempt is made to determine the success or failure of the command. You should assess the success or failure of the command by evaluating the return code from the scheduled command in the schedule log. The schedule log entry for the return code of the command is prefaced with the following text: Finished command. Return code is:.

Related tasks:

"Dsmcutil valid options" on page 290

Related reference:

"Query Schedule" on page 717

Examples: Configuring the CAD to manage an existing scheduler service

This section provides examples of how to configure the client service configuration utility to use scheduler services.

About this task

The example below assumes that the scheduler service name is TSM Central Scheduler Service and the CAD service name is TSM Client Acceptor, which are the default names. You can use the `dsmcutil /name` option to specify different names.

To configure the CAD to manage an existing scheduler service:

Procedure

1. Stop the scheduler service and the CAD, as follows: `dsmcutil stop /name:"tsm central scheduler service"` and then `dsmcutil stop /name:"tsm client acceptor"`.
2. Set the *managedservices* option to *schedule* in the client options file (`dsm.opt`).
3. Update the scheduler service so that it does not start automatically after a reboot: `dsmcutil update /name:"tsm central scheduler service" /autostart:no`.
4. Associate the scheduler service with the CAD: `dsmcutil update cad /name:"tsm client acceptor" /cadschedname:"tsm central scheduler service" /autostart:yes`. If this command is successful, the `dsmwebcl.log` file includes this message: Command will be executed in 1 minute. After one minute, the CAD starts the scheduler and you see information regarding the next scheduled event in the `dsmwebcl.log` file.

Related concepts:

"Dsmcutil commands: Required options and examples" on page 283

Related tasks:

"Dsmcutil valid options" on page 290

Creating a new scheduler and associating a CAD to manage the scheduler

Use step-by-step instructions to create a new scheduler and associate a CAD to manage the scheduler.

Procedure

Follow these steps to create a new scheduler and associate a CAD:

1. Set the *managedservices* option to *schedule* in the client options file (`dsm.opt`).
2. Create the scheduler service: `dsmcutil install scheduler /name:"NEW_SCHEDULE_NAME" /node:yournode /password:xxxxx /startnow:no`. Do not use the */autostart:yes* option when you install a scheduler that is managed by the CAD.
3. Create the CAD service: `dsmcutil install cad /node:yournode /password:xxxxx /autostart:yes /startnow:no`.
4. Associate the scheduler with the CAD: `dsmcutil update cad /name:"tsm client acceptor" /cadschedname:"NEW_SCHEDULE_NAME"`.
5. Start the CAD: `dsmcutil start cad /name:"tsm client acceptor"`.

Results

The CAD and scheduler starts as described. Since the CAD is controlling the scheduler, you do not see the scheduler running as a service, either through the Services applet or the NET START command. To stop the scheduler, you must stop the CAD.


dsmcutil command

The Tivoli Storage Manager Client Service Configuration Utility, **dsmcutil**, can be used to install backup-archive client services on local and remote Windows workstations.

You can use the **dsmcutil** command to install the following client services:

- Backup-Archive Scheduler Service
- Client Acceptor Service (CAD)
- Remote Client Agent Service
- Journal Engine Service
- VMware Backup Tools Service

The Client Service Configuration Utility must be run from an account that belongs to the Administrator/Domain Administrator group. The syntax for the command is as shown in the following text:



Note: Options that you specify with **dsmcutil** commands override option that you specify in your options file (dsm.opt).

The account that runs the utility must have the appropriate user rights for installing services and updating the Windows Registry on the target workstation.

If a remote workstation is specified, the account must be authorized to connect to the Windows Registry of the specified workstation.

Note: For the commands and options that are documented here, the minimum abbreviation that you can type is shown in uppercase letters.

Related concepts:

Chapter 2, “Configure the Tivoli Storage Manager client,” on page 27

Dsmcutil commands: Required options and examples

The **INSTall** command installs and configures backup-archive client services.

INSTall Scheduler

Installs and configures the Tivoli Storage Manager Scheduler Service.

These are the required **INSTall** command options:

- **/name:***service_name*
- **/password:***password*
- **/clusternode:**Yes | No (required if running the Microsoft Cluster Server (MSCS) or Veritas Cluster Server (VCS)).
- **/clustername:***cluster_name* (required if running the MSCS or VCS).

Restriction: Do not specify a clustername of more than 64 characters. If you specify more than 64 characters and you are using Veritas Storage Foundation with High Availability or a Microsoft Cluster Server configuration, you might not be able to install or start the scheduler service.

The `/clientdir:client_dir` option can also be used, the default is the current directory.

The following files must exist in the directory specified by `client_dir`:

- dsmcsvc.exe
- dscenu.txt
- dsm.opt
- dsmntapi.dll
- tsmutil1.dll

Note: If the service is being installed on a remote workstation, the fully qualified client directory path should be relative to the target workstation. UNC names are not allowed for the local system account. Multiple services can be installed on the same workstation.

Task Install a scheduler service that is named TSM Central Scheduler Service on the local workstation. Start the service automatically at system boot time. All required files must reside in the current directory and the client options file must point to the Tivoli Storage Manager server where node ALPHA1 is defined with password nodepw. The server is contacted to verify that the specified node and password are valid. When the password is validated it is generated (encrypted) into the Windows Registry:

Command:

```
dsmcutil install scheduler /name:"TSM Central Scheduler Service"  
/node:ALPHA1 /password:nodepw /autostart:yes
```

Task Install a scheduler service named TSM Central Scheduler Service on remote workstation PDC. Start the service automatically at system boot time. The required scheduler service files and the specified options file must reside on the remote workstation in the `c:\tsm\baclient` directory. The password is encrypted into the Windows Registry. The Tivoli Storage Manager server is not contacted to validate the password.

Command:

```
dsmcutil install scheduler /name:"TSM Central Scheduler Service"  
/machine:PDC /clientdir:c:\tsm\baclient /optfile:c:\tsm\baclient\dsm.opt  
/node:PDC /validate:no /autostart:yes  
/password:nodepassword
```

Task Install a scheduler service named TSM Central Scheduler Service on remote workstation PDC. Start the service automatically at system boot time. The required scheduler service files and the specified options file must reside on the remote workstation in the `c:\tsm\baclient` directory. The password is encrypted into the Windows Registry. The Tivoli Storage Manager server residing at the specified TCP/IP host and port is contacted to validate the password.

Command:

```
dsmcutil install scheduler /name:"TSM Central Scheduler Service"  
/machine:PDC /clientdir:c:\tsm\baclient /optfile:c:\tsm\baclient\dsm.opt  
/node:PDC /autostart:yes /password:nodepassword  
/commmethod:tcpip /commserver:alpha1.example.com  
/commport:1521
```


Task Install the TSM Central Scheduler Service on one node of a MSCS (or VCS) cluster. For *group-a* from workstation *node-1*, ensure that *node-1* currently owns *group-a* and then issue the following command.

Command:

```
dsmcutil install scheduler /name:"TSM Central Scheduler Service:
group-a" /clientdir:c:\tsm\baclient /optfile:q:\tsm\baclient\
dsm.opt /node:mscs-cluster-group-a /password:n
/validate:no /autostart:yes /startnow:yes
/clusternode:yes /clustername:mscs-cluster
```

INSTall CAD

Installs and configures the Client Acceptor Service (CAD). Required options are:

- **/name:***service_name*
- **/node:***node_name*
- **/password:***password*

Other valid options are:

- **/optfile:***options_file*
- **/httpport:***http_port*
- **/webports:***web_ports*

Task Install a Client Acceptor Service called TSM CAD. The client acceptor uses a node called *test* to connect to the Tivoli Storage Manager server. Use the options file *c:\tsm\baclient\dsm.opt* to connect to the server.

Command:

```
dsmcutil install cad /name:"TSM CAD" /node:test
/password:test /optfile:c:\tsm\baclient\dsm.opt
```

INSTall Journal

Installs a journaling engine service on all Windows clients. A journal database is created that stores information the client uses to determine which files are eligible for backup before an operation starts.

If necessary, you can use the *nojournal* option with the **incremental** command to specify that you want to perform a traditional full incremental backup.

The journaling engine service is named TSM Journal Service and uses the configuration file *tsmjbbd.ini* from the backup-archive client installation directory.

Note: The Journal Service is supported in a Microsoft Cluster Server environment. Multiple journal services can be installed by specifying unique pipe names using the *JournalPipe* journal config setting and client options.

There are no valid options for this command.

Task Install the journaling engine service (TSM Journal Service).

Command:

```
dsmcutil install journal
```

INSTall REMOTEAgent

Installs and configures a Remote Client Agent Service. Required options are:

- **/name:***service_name*

- **/node:***node_name*
- **/password:***password*
- **/partnername:***partner_service_name*

Other valid options are:

- **/optfile:***options_file*

Task Install a Remote Client Agent Service called TSM AGENT. The remote client agent uses a node called *test* to connect to the Tivoli Storage Manager server. The options file *c:\tsm\baclient\dsm.opt* is used to connect to. The partner client acceptor service is TSM CAD.

Command:

```
dsmcutil install remoteagent /name:"TSM AGENT" /node:test
/password:test /optfile:c:\tsm\baclient\dsm.opt /partnername:
"TSM CAD"
```

Note: Both the Remote Client Agent Service and the Client Acceptor Service must be installed to run the web client. The Client Acceptor Service must be installed before the Remote Client Agent Service. Use the **/partnername:** option to specify the name of the partner Client Acceptor Service.

INSTall vmwaretools

Installs and configures the VMware backup tools service. There are no required options.

Command:

```
dsmcutil install vmwaretools
```

REMove

Remove an installed Client Service. The required option is **/name:***service_name*.

Task Remove the specified scheduler service from the local workstation.

Command:

```
dsmcutil remove /name:"TSM Central Scheduler Service"
```

Task Remove the journaling engine service (TSM Journal Service) from the local workstation.

Command:

```
dsmcutil remove /name:"TSM Journal Service"
```

REMove vmwaretools

Remove the VMware backup tools service. There are no required options.

Command:

```
dsmcutil remove vmwaretools
```

UPDate

Updates Scheduler Service registry values. The required option for this command is **/name:***service_name*, and the registry values to update. Other valid options are:

- **/clientdir:***client_dir*
- **/optfile:***options_file*

- **/eventlogging:**Yes | No
- **/node:***node_name*
- **/autostart:**Yes | No
- **/clusternode:**Yes | No (required if running the MSCS or VCS).
- **/clustername:***cluster_name* (required if running the MSCS or VCS).

Task Update the client directory and options file for the specified scheduler service. All required client service files must reside in the specified directory.

Note: The communication options specified with the **dsmcutil** command here take precedence over those specified in the client options file.

Command:

```
dsmcutil update /name:"TSM Central Scheduler Service"
/clientdir:c:\tsmptf3\baclient /optfile:c:\tsmptf3\baclient\dsm.opt
```

Task Update the specified scheduler service to use the TCP/IP protocol to connect to a Tivoli Storage Manager server at the specified host name on the specified port.

Command:

```
dsmcutil update /name:"TSM Central Scheduler Service"
/commserver:ntl.example.com /commport:1521 /commmethod:
tcpip
```

UPDate CAD

Updates Client Acceptor Service registry values. The required option for this command is **/name:service_name**, and the registry values to update. Other valid options are:

- **/node:***node_name*
- **/password:***password*
- **/optfile:***options_file*
- **/httpport:***http_port*
- **/webports:***web_ports*
- **/cadschedname:***scheduler_name*

Task Update the Client Acceptor Service to use the specified client password and options file. All required client service files must reside in the specified directory.

Command:

```
dsmcutil update cad /name:"TSM CAD" /password:test
/optfile:c:\tsm\baclient\dsm.opt
```

UPDate REMOTEAgent

Updates Remote Client Agent Service registry values. The required option for this command is **/name:service_name**, and the registry values to update. Other valid options are:

- **/node:***node_name*
- **/password:***password*
- **/optfile:***options_file*
- **/partnername:***partner_service_name*

Task Update a Remote Client Agent Service called TSM AGENT. The remote client agent service uses a node called *test* to connect to the Tivoli Storage Manager server. The options file *c:\tsm\baclient\dsm.opt* is used to connect to the server. The partner client acceptor service is TSM CAD.

Command:

```
dsmcutil update remoteagent /name:"TSM AGENT" /node:test  
/password:test /optfile:c:\tsm\baclient\dsm.opt /partnername:  
"TSM CAD"
```

Query Scheduler

Query Scheduler Service registry values. Required options are: **/name:service_name**. Other valid options are:

- **/machine:***machine_name*
- **/clientdir**
- **/optfile**
- **/eventlogging**
- **/node**
- **/commmethod**
- **/commport**
- **/commserver**
- **/errorlog**
- **/schedlog**

Note: Do not specify a value for the non-required options. The client returns option registry values for the scheduler service you specify.

Task Query registry settings for the scheduler service you specify.

Command:

```
dsmcutil query /name:"TSM Central Scheduler Service"
```

Task Query the client directory registry setting for the scheduler service you specify.

Command:

```
dsmcutil query /name:"TSM Central Scheduler Service"
```

Query CAD

Queries Client Acceptor Service registry values. The required option for this command is **/name:service_name**. Other valid options are:

- **/machine:***machine_name*
- **/node**
- **/optfile**
- **/httpport**
- **/webports**
- **/clientdir**
- **/partnername**

Note: Do not specify a value for these options.

Task Query registry settings for the Client Acceptor Service you specify.

Command:

```
dsmcutil query cad /name:"TSM CAD"
```

Query Journal

Query the journaling engine service, TSM Journal Service on a Windows system. There are no valid options for this command.

Task Query the journaling engine service, **TSM Journal Service**.

Command:

```
dsmcutil query journal
```

Query REMOTEAgent

Queries Remote Client Agent Service registry values. The required option for this command is **/name:service_name**. Other valid options are:

- **/machine:machine_name**
- **/node**
- **/optfile**
- **/partnername**
- **/clientdir**

Note: Do not specify a value for these options.

Task Query registry settings for the specified Remote Client Agent Service.

Command:

```
dsmcutil query remoteagent /name:"TSM AGENT"
```

List

Lists installed Client Services. There are no required options.

Task Locate and list the installed backup-archive client services on the local workstation.

Command:

```
dsmcutil list
```

Task List the installed backup-archive client services on remote workstation PDC.

Command:

```
dsmcutil list /MACHINE:PDC
```

START

Use the **Start** command to start a client service. The **Start** command requires the **/name:service_name** option.

Task Start the journaling engine service, TSM Journal Service.

Command:

```
dsmcutil start /name:"TSM Journal Service"
```

STOP

Use the **Stop** command to stop a client service. The **Stop** command requires the **/name:service_name** option.

Task Stop the journaling engine service, TSM Journal Service.

Command:

```
dsmcutil stop /name:"TSM Journal Service"
```

UPDATEPW

Generate an encrypted Tivoli Storage Manager registry password. The **UPDATEPW** command requires the **/node:node_name** and the **/password:password** options. Optionally, you can use the following options:

- **/validate:Yes** | **No**
- **/clusternode:Yes** | **No** (required if running the MSCS or VCS).
- **/clustername:cluster_name** (required if running the MSCS or VCS).
- **/force:Yes** | **No**

The password is validated with the Tivoli Storage Manager server if **/validate:Yes** is specified. The password is updated on the server if you specify **/updateonserver:Yes**. If you specify this option, you must specify the current password with the **/oldpassword:** option.

Task Update the encrypted registry password for the specified node. Validate and update the password on the specified Tivoli Storage Manager server which resides on the specified TCP/IP hostname and port:

Command:

```
dsmcutil updatepw /node:alpha1 /commMethod:tcpip  
/commServer:alpha1.example.com /commPort:1500  
/password:newpw /oldpassword:oldpw /updateonserver:yes
```

Related concepts:

"Journal-based backup" on page 138

Related tasks:

"Dsmcutil valid options"

Related reference:

"Incremental" on page 683

Dsmcutil valid options

This section lists the valid **dsmcutil** options that you can specify to use the scheduler service.

About this task

/autostart:[Yes|No]

Specifies whether the Scheduler Service starts automatically at system boot time. The default is *No*.

/cadschedname:schedulename

Specifies the name of the scheduler service to manage with the CAD. Use this option when the *managedservices* option is set to *schedule* in the client options file dsm.opt. You can specify this option only with the CAD service.

/clientdir:clientdir

The fully qualified directory path where the Client Service files reside. This

directory should be relative to the target workstation where the service is installed. UNC names are not allowed if the local system account is set to logon. The default is the current directory.

/clustername:clustername

This option replaces the */group* option.

The */clustername* option specifies the cluster name to which the system belongs. You can determine the cluster name in any of the following ways:

- On MSCS, run the MSCS command, CLUSTER /LIST, from the command line or use the Cluster Administrator utility. When the Cluster Administrator utility starts, it displays a tree-like structure with the cluster name at the top.
- On VCS, use the VCS Cluster Manager - Java Console or open the main.cf file in the %VCS_HOME%\config directory.
- On VCS, use the following command:

```
haclus -display
```

Restriction: Do not specify a clustername of more than 64 characters. If you specify more than 64 characters and you are using Veritas Storage Foundation with High Availability or a Microsoft Cluster Server configuration, you might not be able to install or start the Tivoli Storage Manager scheduler service.

This option must be used with the */clusternode:Yes* option. This option must be specified when using the INSTALL command in a cluster environment. It must also be specified when using the UPDATE command to modify the cluster settings (*/clusternode* and */clustername*).

This option can also be specified when using the UPDATEPW command in a cluster environment. Normally this is not required. However, if more than one scheduler service with different cluster settings are defined for a particular node, the utility cannot determine which settings are correct. In this case, correct the discrepancies between the services.

Alternatively, you can specify this option with */clusternode:Yes* and */force:Yes*, to force the utility to show or update the password with the specified cluster settings.

This option is not required if */clusternode:No* is specified.

/clusternode:Yes | No

Specifies whether to enable support for cluster resources. The default value is *No*. You must be running the MSCS or VCS to specify */clusternode:Yes*. This option must be specified when using the INSTALL command in a cluster environment. This option must also be specified when using the UPDATE command to modify the cluster settings (*/clusternode*, */clustername*).

This option can also be specified when using the UPDATEPW command in a cluster environment. Normally this is not required. However, if more than one scheduler service with different cluster settings are defined for a particular node, the utility cannot determine which settings are correct. In this case, correct the discrepancies between the services.

Alternatively, you can specify this option with */clustername* and */force:Yes*, to force the utility to show or update the password with the specified cluster settings. If */clusternode:No* is specified, */clustername* is not required.

/commmethod:protocol

Specifies client communications protocol to communicate with a Tivoli Storage Manager server. Valid protocols are: TCP/IP and Named Pipes. If you do not specify a value, the value is obtained from the client options file or set to the default client value. You can also use this option with the UPDATEPW command to specify a communication protocol to connect to a server when updating passwords.

/commport:serverport

Specifies the protocol specific Tivoli Storage Manager server port. For TCP/IP, this is the port on the specified hostname. If this option is not specified, the value is obtained from the client options file or set to the default client value. You can also use this option with the UPDATEPW command to specify a protocol specific server port to connect to for updating passwords.

/commserver:servername

Specifies the protocol specific Tivoli Storage Manager server name. Depending on the protocol used, this can be a TCP/IP hostname or a Named Pipes name. If not specified, the value is obtained from the client options file or set to the default client value.

This option can also be used with the UPDATEPW command to specify a protocol specific server name to connect to for updating passwords.

/copyfiles

Specifies that the service installation is copied to another location prior to installing the service. Use the */srcdir* option to specify the fully qualified source path.

/errorlog:errorlog

Specifies the fully qualified name of the client error log.

/eventlogging:[Yes | No]

Turns detailed event logging on or off for the specified scheduler service. The default is *Yes*.

/force:[Yes | No]

This option can also be specified when using the UPDATEPW command in a cluster environment. Normally this is not required. However, if more than one scheduler service with different cluster settings is defined for a particular node, the utility cannot determine which settings are correct. In this case, correct the discrepancies between the services.

Alternatively, you can specify this option with */clusternode* and */clusternode* (if */clusternode:Yes* is specified), to force the utility to show or update the password with the specified cluster settings.

/httpport:httpport

Specifies a TCP/IP port address for the web client.

/machine:machinename

Specifies the name of a remote workstation to connect to.

/name:servicename

Specifies the name of the Client service. The name must be quote delimited if it contains embedded spaces.

/node:nodename

Specifies the Tivoli Storage Manager node name the Client Service uses

when connecting to the Tivoli Storage Manager server. Also used when displaying or updating the Tivoli Storage Manager registry password. The default is the workstation name.

/ntaccount:*ntaccount*

Specifies the Windows account which the service logs in as.

/ntdomain:*ntdomain*

Specifies the Windows domain which the service logs in as.

/ntpassword:*ntpassword*

Specifies the Windows password for the account under which the service logs in.

/oldpassword:*oldpw*

Current Tivoli Storage Manager server password. Used in conjunction with the **/updateonserver** option when updating a password on the server.

/optfile:*optionsfile*

The fully qualified path of the client options file. This is the options file the specified Client Service uses to connect to the Tivoli Storage Manager server. The utility also uses the file to connect to the Tivoli Storage Manager server to validate and update passwords. Note that although this option overrides the default option file in the current directory (*dsm.opt*), the Tivoli Storage Manager API requires that a default option file exists in the current directory. UNC names are not allowed if the local system account is set to logon. The default is the *dsm.opt* file in the */clientdir* directory.

/partnername:*partner service name*

This option is used when installing a Remote Client Agent Service to specify the partner Client Acceptor Service.

/password:*password*

The Tivoli Storage Manager password which is generated (encrypted) into the Windows Registry.

/schedlog:*schedlog*

Specifies the fully qualified name of the client schedule log.

/srcdir:*pathname*

Use this option in conjunction with the **/copyfiles** option to specify the fully qualified source path to copy the service installation to another location prior to installing the service.

/startnow:[*Yes* | *No*]

Specifies whether *dsmcutil* starts the specified service after executing the command; the default is *Yes*. If you specify *No*, you must start the service manually using the services control panel applet, or the NET START **name of the service**.

/updateonserver:[*Yes* | *No*]

Specifies whether the specified password is updated on the Tivoli Storage Manager server. Requires using the **/oldpassword** option.

/validate:[*Yes* | *No*]

Specifies whether to perform validation when displaying or updating the encrypted registry password. The default is *Yes*.

/webports: *webports*

Specifies the TCP/IP port number used by the Client Acceptor service and the web client agent service for communications with the web GUI.

Chapter 11. 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 314

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 options file (`dsm.opt`) or on the command line.

You can set the following types of options:

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

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.

Note: Some of the processing options that are used by the Tivoli Storage Manager central scheduler are defined in the Windows registry when the schedule services are configured. These options can also be specified in the client options file. When the scheduler runs as a service, processing options that are specified in the registry override the same options that are specified in the client options file.

Related concepts:

“Entering options with a command” on page 314

Related tasks:

“Creating and modifying the client options file” on page 29

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 all Windows clients, use one of the following protocols:

- TCP/IP
- Named pipes
- Shared memory

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 347

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 37. TCP/IP options

Option	Description
<code>httpport</code> “ <code>Httpport</code> ” on page 423	Specifies a TCP/IP port address for the web client.
<code>lanfreetcpport</code> “ <code>Lanfreetcpport</code> ” on page 453	Specifies the TCP/IP port number where the Tivoli Storage Manager storage agent is listening.
<code>lanfreetcpserveraddress</code> “ <code>Lanfreetcpserveraddress</code> ” on page 454	Specifies the TCP/IP address for the Tivoli Storage Manager storage agent.
<code>tcpbuffsize</code> “ <code>Tcpbuffsize</code> ” on page 555	Specifies the size, in kilobytes, of the internal TCP/IP communication buffer.
<code>tcpnodelay</code> “ <code>Tcpnodelay</code> ” on page 558	Specifies whether the server or client disables the delay of sending successive small packets on the network.
<code>tcpadminport</code> “ <code>Tcpadminport</code> ” on page 554	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 556	Specifies a TCP/IP address for <code>dsmcad</code> .
<code>tcpport</code> “ <code>Tcpport</code> ” on page 559	Specifies the TCP/IP port address for a Tivoli Storage Manager server.
<code>tcpserveraddress</code> “ <code>Tcpserveraddress</code> ” on page 560	Specifies the TCP/IP address for a Tivoli Storage Manager server.
<code>tcpwindowsize</code> “ <code>Tcpwindowsize</code> ” on page 560	Specifies the size, in kilobytes, of the TCP/IP sliding window for your client node.

Table 37. TCP/IP options (continued)

Option	Description
webports “Webports” on page 622	Enables the use of the web client outside a firewall by specifying the TCP/IP port number used by the client acceptor service and the web client agent service for communications with the web GUI.

Named Pipes option

This topic provides information about the namedpipename communication option.

Table 38. Named Pipes communication option

Option	Description
namedpipename “Namedpipename” on page 471	Specifies the name of a named pipe to use for communications between a Tivoli Storage Manager client and server on the same Windows server domain.

Shared memory options

This topic provides information on the shared memory options that you can use.

Table 39. Shared memory communication options

Option	Description
lanfreeshmport “Lanfreeshmport” on page 452	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 525	Specifies the unique number that is used by the client and the server to identify shared memory area used for communications.

Backup and archive processing options

You can specify client options to control some aspects of backup and archive processing.

Table 40. Backup and archive processing options

Option	Description
archmc “Archmc” on page 323	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.
asnodename “Asnodename” on page 324	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.
autofsrename “Autofsrename” on page 333	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.

Table 40. Backup and archive processing options (continued)

Option	Description
backmc "Backmc" on page 335	Specifies the management class to apply to the backup fastback subcommand for retention purposes.
changingretries "Changingretries" on page 340	Specifies the number of times the client attempts to back up or archive a file that is in use.
class "Class" on page 341	Specifies whether to list the NAS or client Application Server objects during a query backup , query filespace , or delete filespace operation.
compressalways "Compressalways" on page 350	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 351	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 354	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 360	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 359	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 359	Determines the maximum size of the data deduplication cache file.
enablededupcache "Enablededupcache" on page 389	Specifies whether you want to enable client-side data deduplication cache, so that the backup-archive client gets the changed data from the cache.
deletefiles "Deletefiles" on page 361	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.

Table 40. Backup and archive processing options (continued)

Option	Description
description "Description" on page 362	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 363	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 364	Use the diffsnapshot option to determine whether the client creates a differential snapshot.
dirmc "Dirmc" on page 367	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 368	Backs up, restores, archives, retrieves, or queries directories only.
diskcachelocation "Diskcachelocation" on page 370	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 371	Specifies the drives to include in your default client domain for an incremental backup.
domain.image "Domain.image" on page 374	Specifies the file systems and raw logical volumes that you want to include in your client domain for an image backup. This option is only valid for Windows 32-bit clients.
domain.nas "Domain.nas" on page 375	Specifies the volumes to include in your default domain for NAS image backups.
domain.vmfile "Domain.vmfile" on page 378	Specifies the virtual machines to include in file-level backups of VMware virtual machines.
domain.vmfull "Domain.vmfull" on page 381	Specifies the virtual machines to include in full image backups of VMware virtual machines.
enablearchiveretentionprotection "Enablearchiveretentionprotection" on page 388	Allows the client to connect to a data retention server.
enablelanfree "Enablelanfree" on page 392	Specifies whether to enable an available LAN-free path to a storage area network (SAN) attached storage device.
exclude exclude.backup exclude.file exclude.file.backup	Use these options to exclude a file or group of files from backup services.

Table 40. Backup and archive processing options (continued)

Option	Description
encryptiontype "Encryptiontype" on page 393	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 394	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 400	Excludes a file or a group of files that match the pattern from archive services only.
exclude.compression "Exclude options" on page 400	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 400	Excludes a directory, its files, and all its subdirectories and their files from backup processing.
exclude.encrypt "Exclude options" on page 400	Excludes specified files from encryption processing.
exclude.fs.nas "Exclude options" on page 400	Excludes file systems on the NAS file server from an image backup when used with the backup nas command.
exclude.image "Exclude options" on page 400	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 <code>exclude.image</code> .
exclude.subfile "Exclude options" on page 400	Excludes files from adaptive subfile backup processing. This option does not apply to archive processing.
fbbranch "Fbbranch" on page 406	Specifies the branch ID of the remote FastBack server to back up or archive.
fbclientname "Fbclientname" on page 407	Specifies the name of one or more FastBack clients to back up from the backup proxy.
fbpolicyname "Fbpolicyname" on page 409	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 410	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.

Table 40. Backup and archive processing options (continued)

Option	Description
fbserver "Fbserver" on page 411	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 412	Specifies the name of one or more Tivoli Storage Manager FastBack volumes to back up from the backup proxy.
filelist "Filelist" on page 414	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 418	Backs up, restores, retrieves, or queries files only.
groupname "Groupname" on page 422	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 422	Specifies whether the client GUI is returned to the Backup, Restore, Archive, or Retrieve window after an operation finishes.
ieobjtype "Ieobjtype" on page 425	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 427	Specifies the minimum size of empty regions on a volume that you want to skip during backup. This option is valid for Windows 32-bit clients.
incl excl "Incl excl" on page 429	Specifies the path and file name of an include-exclude options file.
"Include options" on page 430	Use these options to include files or assign management classes for backup processing.
include include.backup include.file	
include.archive "Include options" on page 430	Includes files or assigns management classes for archive processing.
include.compression "Include options" on page 430	Includes files for compression processing if you set the compression option to <i>yes</i> . This option applies to backups and archives.
include.encrypt "Include options" on page 430	Includes the specified files for encryption processing. By default, the client does not perform encryption processing.

Table 40. Backup and archive processing options (continued)

Option	Description
include.fs "Include options" on page 430	Use the include.fs option to specify processing options for a file system. Use the include.fs option to specify which drives use open file support and to control how full file space incremental backups are processed.
include.fs.nas "Include options" on page 430	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, by using the toc option with the include.fs.nas option in your client options file (dsm.opt). For more information, see "Toc" on page 563.
include.image "Include options" on page 430	Specifies a file system or logical volume to be included for image backup processing. This option also provides a way to specify an explicit management class assignment for a specified file system or logical volume. The backup image command ignores all other include options. Use the include.fs option to specify which drives use open file support and to control how full file space incremental backups are processed.
include.subfile "Include options" on page 430	Includes files for adaptive subfile backup processing. This option does not apply to archive processing.
include.systemstate "Include options" on page 430	Assigns management classes for backup of the Windows system state. The default is to bind the system object to the default management class.
incrbydate "Incrbydate" on page 445	Use with the incremental command to request an incremental backup by date.
incremental "Incremental" on page 446	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.
incrthreshold "Incrthreshold" on page 446	The incrthreshold option specifies the threshold value for the number of directories in any journaled file space that might have active objects on the server, but no equivalent object on the workstation.
memoryefficientbackup "Memoryefficientbackup" on page 462	Specifies a memory-saving backup algorithm for incremental backups when used with the incremental command.

Table 40. Backup and archive processing options (continued)

Option	Description
mode "Mode" on page 463	<p>Use the mode option with these commands, as follows:</p> <p>backup image To specify whether to perform a selective or incremental image backup of client file systems.</p> <p>backup nas To specify whether to perform a full or differential image backup of NAS file systems.</p> <p>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.</p> <p>backup vm To specify whether to perform a full or incremental backup of a VMware virtual machine when <code>vmbackuptype=fullvm</code>, and when you have installed IBM Tivoli Storage Manager for Virtual Environments.</p>
monitor "Monitor" on page 468	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 474	Suppresses the confirmation prompt that is presented by the delete group , delete archive , expire , restore image , and set event commands.
nojournal "Nojournal" on page 474	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 478	Specifies the client options file you want to use when you start a backup-archive client session.
postsnapshotcmd "Postsnapshotcmd" on page 486	During an online image backup or open file support operation, this option allows you to manually open an application after the snapshot provider starts a snapshot. This option is only valid if the OFS or online image support is enabled.

Table 40. Backup and archive processing options (continued)

Option	Description
preservelastaccessdate "Preservelastaccessdate" on page 489	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 492	During an online image backup or open file support operation, this option allows you to manually quiesce an application before the snapshot provider starts a snapshot. This option is only valid if the OFS or online image support is enabled.
resetarchiveattribute "Resetarchiveattribute" on page 507	Specifies whether the client resets the Windows archive attribute on files that are successfully backed up to a Tivoli Storage Manager server. This option is valid for all Windows clients.
skipntpermissions "Skipntpermissions" on page 528	Specifies whether to back up, archive, retrieve, or restore Windows security information.
skipntsecuritycrc "Skipntsecuritycrc" on page 528	Specifies whether to compute the security CRC for permission comparison during subsequent backups. Use this option on all Windows clients.
snapdiff "Snapdiff" on page 529	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.
snapshotproviderfs "Snapshotproviderfs" on page 535	Use the snapshotproviderfs option to enable snapshot-based file backup and archive operations, and to specify a snapshot provider.
snapshotproviderimage "Snapshotproviderimage" on page 536	Use the snapshotproviderimage option to enable snapshot-based online image backup, and to specify a snapshot provider.
snapshotroot "Snapshotroot" on page 537	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.
subdir "Subdir" on page 548	Specifies whether to include subdirectories of a named directory.

Table 40. Backup and archive processing options (continued)

Option	Description
subfilebackup "Subfilebackup" on page 549	Specifies whether the client uses adaptive subfile backup.
subfilecachepath "Subfilecachepath" on page 550	Specifies the path where the client cache is located for adaptive subfile backup processing.
subfilecachesize "Subfilecachesize" on page 551	Specifies the client cache size for adaptive subfile backup.
tapeprompt "Tapeprompt" on page 553	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 563	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 567	Use the type option with the query node command to specify the type of node to query.
v2archive "V2archive" on page 570	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 575 (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.
vmchost "Vmchost" on page 583	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 584	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 586	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.

Table 40. Backup and archive processing options (continued)

Option	Description
vmlist "Vmlist" on page 598	Used with the backup VM , restore VM , or query VM commands to specify the host name or list of host names of the virtual machine to back up.
vmmxvirtualdisks "Vmmxvirtualdisks" on page 599	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 610	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.

The following options are backup-archive client options that apply only to Tivoli Storage Manager HSM for Windows migrated files.

- Restorecheckstubaccess
- Restoremigstate
- Skipmigrated

Related concepts:

➡ Options for backing up migrated files: skipmigrated, checkreparsecontent, stagingdirectory

➡ Options for restoring migrated files: restorecheckstubaccess, restoremigstate

Restore and retrieve processing options

You can use client options to control some aspects of restore and retrieve processing.

Table 41 lists the restore and retrieve processing options that are available.

Table 41. Restore and retrieve processing options

Option	Description
asrmode "Asrmode" on page 327	Use this option with the restore , and restore systemstate commands to specify whether to perform a restore operation in system ASR recovery mode. This option is used in the context of restore commands that are generated in the asr.sif file by the backup asr command only. Do not use this option outside the context of ASR recovery mode.
backupsetname "Backupsetname" on page 336	The backupsetname option specifies either the name of the backup set, or the name of the file or tape device that contains the backup set. This option is used with the location option.
dirsonly "Dirsonly" on page 368	Qualifies the operation (backup, archive, restore, retrieve) to process directories alone.

Table 41. Restore and retrieve processing options (continued)

Option	Description
disablenqr “Disablenqr” on page 368	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 414	Specifies a file that contains a list of files to be processed by the specified command.
filesonly “Filesonly” on page 418	Qualifies the operation (backup, archive, restore, retrieve) to process files alone.
fromdate “Fromdate” on page 420	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 420	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.
fromtime “Fromtime” on page 421	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 422	Specifies whether the client GUI is returned to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.
ifnewer “Ifnewer” on page 426	Replaces an existing file with the latest backup version only if the backup version is newer than the existing file.
imagetofile “Imagetofile” on page 428	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.
inactive “Inactive” on page 428	Displays a list of active and inactive files when used with the pick option.
latest “Latest” on page 456	Restores the most recent backup version of a file whether it is active or inactive.
localbackupset “Localbackupset” on page 456	Specifies whether the backup-archive client GUI bypasses initial logon with the server to restore a local backup set on a stand-alone workstation.
monitor “Monitor” on page 468	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 474	suppresses the confirmation prompt that is presented by the delete group , delete archive , expire , restore image , and set event commands.
optfile “Optfile” on page 478	Specifies the client options file you want to use when you start a backup-archive client session.

Table 41. Restore and retrieve processing options (continued)

Option	Description
pick "Pick" on page 481	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 482	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 483	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 490	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 499	Specifies whether to overwrite an existing file, or to prompt you for your selection when you restore or retrieve files.
showmembers "Showmembers" on page 526 (does not apply to Mac OS X)	Displays all members of a group.
subdir "Subdir" on page 548	Specifies whether you want to include subdirectories of a named directory.
tapeprompt "Tapeprompt" on page 553	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 564	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 565	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 567	Use the type option with the query node command to specify the type of node to query.
verifyimage "Verifyimage" on page 575	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 http://www.ibm.com/support/knowledgecenter/SSATMW_7.1.6/hsmwin/welcome.html.

- 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 42 lists the scheduling options that are available.

Table 42. Scheduling options

Option	Description
cadlistenonport “Cadlistenonport” on page 338	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 457	Specifies whether the client acceptor manages the web client, the scheduler, or both.
maxcmdretries “Maxcmdretries” on page 459	Specifies the maximum number of times the client scheduler attempts to process a scheduled command that fails.
postschedulecmd/postnschedulecmd “Postschedulecmd/Postnschedulecmd” on page 484	Specifies a command to process after running a schedule.
preschedulecmd/prenschedulecmd “Preschedulecmd/Prenschedulecmd” on page 487	Specifies a command to process before running a schedule.
querschedperiod “Querschedperiod” on page 494	Specifies the number of hours the client scheduler waits between attempts to contact the server for scheduled work.
retryperiod “Retryperiod” on page 511	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.
runasservice “Runasservice” on page 512	Forces the client command process to continue running, even if the account that started the client logs off. Use this option on all Windows clients.
schedcmddisabled “Schedcmddisabled” on page 513	Specifies whether to disable the scheduling of generic commands specified by your Tivoli Storage Manager administrator.
schedlogmax “Schedlogmax” on page 515	Specifies the maximum size of the scheduler log and web client log, in megabytes.
schedlogname “Schedlogname” on page 517	Specifies the path and file name where you want to store schedule log information.

Table 42. Scheduling options (continued)

Option	Description
schedlogretention “Schedlogretention” on page 518	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 519	Specifies which schedule mode to use, <i>polling</i> or <i>prompted</i> .
schedrestretrdisabled “Schedrestretrdisabled” on page 521	Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing restore or retrieve schedule operations.
sessioninitiation “Sessioninitiation” on page 523	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 540	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 541	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 557	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 519) for details.
tcpclientport “Tcpclientport” on page 557	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 519) for details.

Format and language options

Format and language options allow you to select different formats for date, time and numbers for different languages.

Table 43. Format and language options

Option	Description
dateformat “Dateformat” on page 356	Specifies the format for displaying dates.
language “Language” on page 455	Specifies the language used for messages.
numberformat “Numberformat” on page 476	Specifies the format for displaying numbers.

Table 43. Format and language options (continued)

Option	Description
timeformat "Timeformat" on page 562	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 44. Command processing options

Option	Description
quiet "Quiet" on page 496	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 521	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 522	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 574	Specifies that processing information should be displayed on your screen. The alternative is quiet. This option can be overridden by the server.

Authorization options

Authorization options control access to a Tivoli Storage Manager server.

Table 45 lists the authorization options that are available.

Table 45. Authorization options

Option	Description
autodeploy "Autodeploy" on page 332	Specifies whether you want to enable or disable an automatic deployment of the client if a restart is required.
password "Password" on page 478	Specifies a Tivoli Storage Manager password.
passwordaccess "Passwordaccess" on page 480	Specifies whether you want to use a generated password or be prompted for a password each time you start the client.
revokeremoteaccess "Revokeremoteaccess" on page 512	Restricts an administrator with client access privileges from accessing your workstation through the web client.

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 46 lists the error processing options that are available.

Table 46. Error processing options

Option	Description
errorlogmax “Errorlogmax” on page 396	Specifies the maximum size of the error log, in megabytes.
errorlogname “Errorlogname” on page 397	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 398	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 47 lists the transaction processing options that are available.

Table 47. Transaction processing options

Option	Description
collocatebyfilespec “Collocatebyfilespec” on page 346	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 348	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 349	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 369	Specifies the maximum disk I/O buffer size (in kilobytes) that the client can use when reading files.
largecommbuffers “Diskbuffsize” on page 369	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.

Table 47. Transaction processing options (continued)

Option	Description
resourceutilization "Resourceutilization" on page 508	Use the resourceutilization option in your client options file dsm.opt to regulate the level of resources the Tivoli Storage Manager server and client can use during processing.
txnbytelimit "Txnbytelimit" on page 566	Specifies the number of kilobytes the client program buffers before it sends a transaction to the server.
usedirectory "Usedirectory" on page 568	Provides a convenient way to simplify client communication configuration by overriding commethod parameters set in the client options file and instead querying the Active Directory for the communication method and server with which to connect.

Web client options

Several backup-archive client options are used to configure the Tivoli Storage Manager web client.

Table 48 lists the web client options that are available.

Table 48. Web client options

Option	Description
httpport "Httpport" on page 423	Specifies a TCP/IP port address for the web client.
managedservices "Managedservices" on page 457	Specifies whether the client acceptor service manages the web client, the scheduler, or both.
revokeremoteaccess "Revokeremoteaccess" on page 512	Restricts administrator access on a client workstation through the web client.
webports "Webports" on page 622	Enables the use of the web client outside a firewall by specifying the TCP/IP port number used by the client acceptor service and the web Client Agent service for communications with the web client.

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 49 lists the diagnostics options that are available.

Table 49. Diagnostics options

Option	Description
console "Console" on page 352	Use the console option with the query systeminfo command to output system information to the console.
filename "Filename" on page 417	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 718

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.

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 50 on page 315.

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="Project A" c:\devel\proj1\*
```

- For options that do not include parameters, enter a command, a dash (-), and the option name. For example,

```
dsmc incremental -quiet
```

Note: Use a leading dash (-) to indicate that the following text is the name of an option. If an object name begins with a dash, you must surround it in either single 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 c:\devel\proj1\*  
dsmc selective c:\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" c:\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.
- The maximum number of bytes for a file name and file path is 6255 combined. However, the file name itself cannot exceed 255 bytes and the path that leads to the file cannot exceed 6000 bytes. Furthermore, directory names (including the directory delimiter) within a path are limited to 255 bytes. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.

Table 50 lists client command options that you can enter only on the command line with specific commands.

Table 50. Client command options

Command option	Description	Commands
archmc "Archmc" on page 323	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 341	Specifies whether to display a list of NAS objects or client objects when you use the following commands.	query backup delete filespace query filespace
restore systemstate restore backupset		
console "Console" on page 352	Use the console option with the query systeminfo command to output system information to the console.	query systeminfo
deletefiles "Deletefiles" on page 361	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 362	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 363	Displays management class, file space, backup, and archive information, depending on the command with which it is used.	delete filespace query archive query backup query filespace query mgmtclass

Table 50. Client command options (continued)

Command option	Description	Commands
dirsonly "Dirsonly" on page 368	Backs up, restores, archives, retrieves, or queries directories only.	archive incremental query archive query backup restore restore backupset retrieve selective
filelist "Filelist" on page 414	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 417	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 418	Backs up, restores, retrieves, or queries files only.	archive incremental query archive query backup restore restore backupset retrieve selective
fromdate "Fromdate" on page 420	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 420	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

Table 50. Client command options (continued)

Command option	Description	Commands
fromtime "Fromtime" on page 421	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 422	Specifies the fully qualified name for a group.	backup group
ifnewer "Ifnewer" on page 426	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
imagnetofile "Imagnetofile" on page 428	Use the imagnetofile 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.	restore image
inactive "Inactive" on page 428	Displays a list of active and inactive files when used with the pick option.	delete group query backup query group query image query nas query systemstate restore restore group restore image restore nas restore systemstate
incrbydate "Incrbydate" on page 445	Requests an incremental backup by date.	incremental
incremental "Incremental" on page 446	Applies changes to the base image by using information from incremental backups that are made after the original image backup.	restore image
latest "Latest" on page 456	Restores the most recent backup version of a file whether it is active or inactive.	restore restore group
mode "Mode" on page 463	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 50. Client command options (continued)

Command option	Description	Commands
monitor "Monitor" on page 468	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.	backup nas restore nas
nojournal "Nojournal" on page 474	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 474	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 478	Specifies the client options file you want to use when you start a backup-archive client session.	dsmc.exe
pick "Pick" on page 481	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 asr restore group restore image restore nas retrieve
pitdate "Pitdate" on page 482	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 query systemstate restore restore group restore image restore nas restore systemstate All query and restore system object commands
pittime "Pittime" on page 483	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 query systemstate restore restore image restore nas restore systemstate All query and restore system object commands

Table 50. Client command options (continued)

Command option	Description	Commands
preservepath “Preservepath” on page 490	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
runasservice “Runasservice” on page 512	Forces the client command process to continue running, even if the account that started the client logs off. Use this option on all Windows clients.	schedule
showmembers “Showmembers” on page 526	Displays all members of a group.	query group query systemstate restore group
todate “Todate” on page 564	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 565	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 567	Use the type option with the query node command to specify the type of node to query.	query node
v2archive “V2archive” on page 570	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 575	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 only valid for Windows 32-bit platforms.	restore image
virtualfsname “Virtualfsname” on page 575	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 **commmethod** and **optfile** options. Options in this category are not valid in interactive, macro, or scheduler modes. They generate an error and cause processing to stop.

Table 51 on page 320 lists the options that are valid only on the initial command line.

Table 51. Options that are valid on the initial command line only

Options valid on the initial command line

asrmode	preschedulecmd/prenschedulecmd (can be included in the schedule definition)
backupregistry	presnapshotcmd
commmethod	querschedperiod
computername	resourceutilization
deduplication	retryperiod
diskbuffsize	runasservice
editor	schedlogmax
enablededupcache	schedlogname
enablelanfree	schedlogretention
errorlogmax	schedmode
errorlogname	sessioninitiation
errorlogretention	subfilebackup
incrthreshold	subfilecachepath
lanfreecommmethod	subfilecachesize
lanfreeshmport	tcpbuffsize
lanfreetcport	tcpaddress
maxcmdretries	tcpclientaddress
namedpipename	tcpclientport
nodename	tcpport
optfile	tcpserveraddress
password	tcpwindowsize
postschedulecmd/postnschedulecmd (can be included in the schedule definition)	txnbytelimit
postsnapshotcmd	usedirectory
	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 52 on page 321 lists the options that can be set by the server.

Table 52. Options that can be set by the Tivoli Storage Manager server


Options that can be set by the Tivoli Storage Manager server

- “Casesensitiveaware” on page 339
- “Changingretries” on page 340
- “Collocatebyfilespec” on page 346
- “Compressalways” on page 350
- “Compression” on page 351
- “Deduplication” on page 360
- “Dirmc” on page 367
- “Disablenqr” on page 368
- “Diskcachelocation” on page 370
- “Domain” on page 371
- “Domain.image” on page 374
- “Domain.nas” on page 375
- “Encryptiontype” on page 393
- “Encryptkey” on page 394
- “Exclude options” on page 400
- “Incl excl” on page 429
- “Include options” on page 430
- maxcandprocsmaxcandprocs
- maxmigratorsmaxmigrators
- “Memoryefficientbackup” on page 462
- “Postschedulecmd/Postnschedulecmd” on page 484
- “Postsnapshotcmd” on page 486
- “Preschedulecmd/Prenschedulecmd” on page 487
- “Preserve lastaccessdate” on page 489
- “Presnapshotcmd” on page 492
- “Querschedperiod” on page 494
- “Quiet” on page 496
- “Resetarchiveattribute” on page 507
- “Resourceutilization” on page 508
- “Retryperiod” on page 511
- “Schedmode” on page 519
- “Scrolllines” on page 521
- “Scrollprompt” on page 522
- “Snapshotproviderfs” on page 535
- “Snapshotproviderimage” on page 536
- “Stagingdirectory” on page 547
- “Subdir” on page 548
- “Subfilebackup” on page 549
- “Subfilecachepath” on page 550
- “Subfilecachesize” on page 551
- “Tapeprompt” on page 553
- “Txnbytelimit” on page 566
- “Verbose” on page 574
- “Vmchost” on page 583
- “Vmcuser” on page 586
- “Vmlist” on page 598
- “Vmprocessvmwithindependent” on page 604
- “Vmprocessvmwithprdm” on page 606

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.

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.

During a domain incremental backup, where `systemstate` is specified as part of the domain, the `absolute` option does not force a full backup of system state objects. To force a domain incremental backup operation to create a full backup of system state objects, you must add `systemstatebackupmethod full` to the client options file.

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 c:\foo\*.c
```

Adlocation

You can use the adlocation option with the **query adobjects** or **restore adobjects** commands to indicate whether the Active Directory objects are to be queried or restored from the local Active Directory Deleted Objects container or from a system state backup on the Tivoli Storage Manager server.

Supported Clients

This option is valid for supported Windows Server clients. The Tivoli Storage Manager API does not support this option.

Syntax

►—ADLOCation—local
server—◄◄

Parameters

server

Specifies that the Active Directory objects are to be queried or restored from a system state backup on the Tivoli Storage Manager server. Valid for all supported Windows server clients.

local

Specifies that the Active Directory objects are to be queried or restored from the local Active Directory Deleted Objects container. This is the default.

Example

Command line:

```
query adobjects "cn=Jim Smith" -adlocation=server
```

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 Windows clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—ARCHMc =—*managementclass*————►►

Parameters

managementclass
Specifies an available management class in the active policy set of your policy domain. This management class overrides the default management class and any include statements for the files and directories you are archiving.

Examples

Command line:
dsmc archive -archmc=ret2yrs c:\plan\proj1\budget.jan*

Asnodename

Use the `asnodename` option to allow an agent node to back up, archive, restore, retrieve, and query data on behalf of a target node.

An *agent node* is a client node that the Tivoli Storage Manager administrator grants the authority to perform client operations on behalf of a *target node*. The target node is the client node that the agent node performs the actions for. The administrator uses the **grant proxynode** command on the Tivoli Storage Manager server to grant this authority.

Agent nodes can be used to distribute the workload of backing up a computer’s volumes, across multiple client systems. Each system that is involved in the backup uses its own agent node name, but the backup data is stored in a common file space that is owned by the target node.

For example, assume that you plan to back up four volumes that belong to a node that is named SCORPIO, but the backup operation takes too long to run. You can distribute part of the workload to three other machines: TAURUS, ARIES, and LEO. SCORPIO and the three other machines each back up one of SCORPIO’s volumes. Each node that is involved in the backup connects to the server by using its own agent node name, and each node specifies a unique value for the `asnodename` option. Do not use a computer name or cluster name for the `asnodename` value. The following table illustrates an example configuration.

Table 53. Setting the value of the `asnodename` option to distribute backups.

Host name	NODENAME option value	ASNODENAME option value	Volume backed up	Server file space name
SCORPIO	SCORPIO	TARGET_SCORPIO	\\scorpio\r\$	\\target_scorpio\r\$

Table 53. Setting the value of the `asnodename` option to distribute backups. (continued)

Host name	NODENAME option value	ASNODENAME option value	Volume backed up	Server file space name
TAURUS	TAURUS	TARGET_SCORPIO	\\scorpio\s\$	\\target_scorprio\s\$
ARIES	ARIES	TARGET_SCORPIO	\\scorpio\t\$	\\target_scorprio\t\$
LEO	LEO	TARGET_SCORPIO	\\scorpio\u\$	\\target_scorprio\u\$

To create the relationships between the target node and the proxy nodes, the Tivoli Storage Manager server administrator needs to take the following actions:

1. Register nodes SCORPIO, TAURUS, ARIES, LEO, and TARGET_SCORPIO.
2. Grant nodes SCORPIO, TAURUS, ARIES, and LEO proxy authority to node TARGET_SCORPIO

When you back up or archive data without the `asnodename` option, the backed up data is stored in a file space on the server that matches the UNC name of the drive on which the original data exists.

When you use the `asnodename` option to back up data on behalf of a target node, the data is stored in a file space that is owned by the target node. However, instead of using the host name in the file space name, the target node name is used in the file space name. For example, if node TAURUS backs up data on SCORPIO's S drive and sets the `asnodename` option value to `-asnodename=target_scorprio`, the backup data is stored in a file space named `\\target_scorprio\s$`. The file space is owned by the TARGET_SCORPIO node.

When you restore or retrieve data, the default behavior is to restore or retrieve the data to a location that matches the file space name.

Continuing with the preceding example, if node SCORPIO uses `-asnodename=target_scorprio` to restore data from `\\target_scorprio\s$`, the client attempts to restore the data to the S drive on a computer named TARGET_SCORPIO. This operation does not produce the expected result because, in this sample configuration, there is no computer that is named TARGET_SCORPIO.

In the following example, the **restore** command is entered on the SCORPIO node. The command restores all files and subdirectories from the `Users\andy\education` directory in the `\\target_scorprio\s$` file space to the S drive on the computer that is named SCORPIO:

```
dsmc restore \\target_scorprio\s$\users\andy\education\* s:\
-subdir=yes -asnodename=target_scorprio
```

The following considerations apply 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.
- You cannot use `asnodename` with the **backup nas** command.
- You cannot use `asnodename` with the `fromnode` option.
- If you use `asnodename` to backup and restore volumes that are in a cluster configuration, do not use `clusternode` yes.

- You cannot use `asnodename` to back up or restore system state.
- If an agent node restores data from a backup set, the system state object in the backup set is not restored.
- You can use `asnodename` with the **backup image** command, but you must specify the volume by UNC name. You cannot use the drive letter.
- If you use the same `asnodename` value to back up files from different machines, you need to keep track which files or volumes are backed up from each system so that you can restore them to the correct location.
- All agent nodes in a multiple node environment should be of the same platform type.
- Do not use target nodes as traditional nodes, especially if you encrypt your files before backing them up to the server.

Supported Clients

This option is valid for all Windows clients.

Options File

Place this option in the `dsm.opt` file. You can set this option on the **General** tab of the Preferences editor.

Syntax

►►—ASNODENAME— *—targetnode—*—————►►

Parameters

targetnode

Specifies the node name on the Tivoli Storage Manager server under which you want to back up or restore data.

Examples

Options file:

```
asnodename target_scorpio
```

Command line:

This command backs up the entire F: drive to a server file space named `\\target_scorpio\f$`.

```
dsmc incremental f: -asnodename=target_scorpio
```

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.

Asrmode


Use the **asrmode** option with the **restore** and **restore systemstate** commands to specify whether to perform a restore operation in system ASR recovery mode.

This option is used in the context of **restore** commands generated in the **asr.sif** file by the **backup asr** command only.

Supported Clients

This option is valid for supported Windows clients that are running in a Windows Preinstallation Environment; both BIOS and UEFI boot architectures are supported.

Syntax



Parameters

No Specifies that the client does not perform the restore operation in system ASR recovery mode.

Yes

Specifies that the client performs the restore operation in ASR recovery mode. This is the default for Windows clients during ASR recovery. These clients are running in Windows Preinstallation Environment (WinPE) during ASR recovery.

Examples

Command line:

```
restore systemstate -asrmode=yes  
restore systemstate -asrmode=yes -inactive -pick
```

This option is valid for an interactive session, but cannot be changed by entering the option while running an interactive session.

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 c:\test\file.txt  
04/21/07 15:25:05 ANS1651I Backed Up:  
    \\spike\c$\test\file.txt  
04/21/07 15:25:05 ANS1652I Archived:  
    \\spike\c$\test\file.txt  
04/21/07 15:25:05 ANS1653I Updated:  
    \\spike\c$\test\file.txt  
04/21/07 15:25:05 ANS1654E Failed:  
    \\spike\c$\test\file.txt  
04/21/07 15:25:05 ANS1655I Restored:  
    \\spike\c$\test\file.txt  
04/21/07 15:25:05 ANS1656I Retrieved:  
    \\spike\c$\test\file.txt  
04/21/07 15:25:05 ANS1657I Expired:  
    \\spike\c$\test\file.txt  
04/21/07 15:25:05 ANS1658I Deleted:  
    \\spike\c$\test\file.txt  
04/21/07 15:25:05 ANS1659I Skipped:  
    \\spike\c$\test\file.txt
```

The following is an example of the messages that are issued when the audit log is configured to capture the full level of information (in addition to all messages issued for the basic level of audit logging):

```
04/21/07 15:25:05 ANS1660I Excluded:
  \\spike\c$\test\file.txt
04/21/07 15:25:05 ANS1661I Unchanged:
  \\spike\c$\test\file.txt
```

The audit log is not a substitute or a replacement for the standard error log (`dsmerror.log`) or for the schedule log (`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.

When using the Windows backup-archive client, all object names are written in the UNC format. The Windows backup-archive client creates the audit log as a Unicode file.

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

The **auditlogging** command is not supported with backup commands which interact with image-level objects such as **backup image** or **restore image**. The **auditlogging** command is supported with backup commands that interact with file-level objects such as **backup groups**, and **backup systemstate**.

If you have enabled audit logging for an operation and there is a failure trying to write to the audit log (for example, the disk on which the audit log resides is out of space), the audit logging is disabled for the rest of the operation and the return code for the operation is set to 12, regardless of the outcome of the operation.

Supported Clients

This option is valid for all clients.

Options File

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

Syntax



Parameters

off
Specifies that the audit logging facility is not engaged. This is the default.

basic
Specifies that the audit log captures a basic level of information.

full

Specifies that the audit log captures a more extensive level of information.

Examples

Run an incremental backup with audit logging enabled.

Command line:

```
dsmc i -auditlogging=basic
```

Back up a list of files using the maximum level of auditing, which enables a separate application, such as a Perl script, to verify the results.

```
dsmc i -filelist=file.lst -auditlogging=full  
-auditlogname="c:\program files\tivoli\tsm\baclient\  
temp_audit001.log"
```

Auditlogname

The `auditlogname` option specifies the path and file name where you want to store audit log information. This option applies when audit logging is enabled.

Supported Clients

This option is valid for all clients.

Options File

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

Syntax

►►—AUDITLOGName—*filespec*—————►►

Parameters

filespec

Specifies the path and file name where you want 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.

In Uniform Naming Convention (UNC) format, the path must contain a drive letter. In the following example, the path contains the drive letter D\$: `\\computer7\D$\logs\tsmaudit.log`.

Examples

Run an incremental backup with audit logging enabled.

Options file:

Store the audit log in a non-default path.

```
auditlogname c:\mypath\myaudit.log
```

Command line:

Back up a list of files using the maximum level of auditing, which would enable a separate application, such as a Perl script, to verify the results:

```
dsmc i -filelist=file.lst -auditlogging=full
-auditlogname="c:\program files\tivoli\tsm\baclient\
temp_audit001.log"
```

Sample output

The following is a sample execution and output file:

```
C:\Program Files\Tivoli\TSM\baclient>dsmc i
c:\test\* -sub=yes -auditlogging=full
IBM Tivoli Storage Manager
Command Line Backup/Archive Client Interface
Client Version 5, Release 5, Level 0.0
Client date/time: 04/21/2007 15:52:25
(c) Copyright by IBM Corporation and other(s)
1990, 2007. All Rights Reserved.

Node Name: PATMOS
Session established with server PATMOS_5331: Windows
Server Version 5, Release 3, Level 3.1
Server date/time: 04/21/2007 15:52:26
Last access: 04/21/2007 15:52:06
```

```
Incremental backup of volume 'c:\test\*'
Normal File-->          1,048,576 \\patmos\c$\test
\dir1\file1 [Sent]
Normal File-->          1,048,576 \\patmos\c$\test
\dir1\file2 [Sent]
Normal File-->           1,024 \\patmos\c$\test
\dir1\file3 [Sent]
Normal File-->          1,048,576 \\patmos\c$\test
\dir2\file1 [Sent]
Normal File-->          1,048,576 \\patmos\c$\test
\dir2\file2 [Sent]
Normal File-->           1,024 \\patmos\c$\test
\dir2\file3 [Sent]
Successful incremental backup of '\\patmos\c$\test\*'
```

```
Total number of objects inspected:      12
Total number of objects backed up:       6
Total number of objects updated:         0
Total number of objects rebound:        0
Total number of objects deleted:         0
Total number of objects expired:         0
Total number of objects failed:          0
Total number of subfile objects:         0
Total number of bytes transferred:      400.85 KB
Data transfer time:                      0.00 sec
Network data transfer rate:              0.00 KB/sec
Aggregate data transfer rate:            382.85 KB/sec
Objects compressed by:                   91%
Subfile objects reduced by:              0%
Elapsed processing time:                  00:00:01
ANS1900I Return code is 0.
ANS1901I Highest return code was 0.
```

The following are the audit log contents:

```
04/21/2007 15:52:25 ANS1650I Command:
i c:\test\*
04/21/2007 15:52:26 ANS1661I Unchanged:
\\patmos\c$\test
04/21/2007 15:52:26 ANS1661I Unchanged:
\\patmos\c$\test\dir1
04/21/2007 15:52:26 ANS1661I Unchanged:
\\patmos\c$\test\dir2
04/21/2007 15:52:26 ANS1661I Unchanged:
\\patmos\c$\test\file1
04/21/2007 15:52:26 ANS1661I Unchanged:
```

```

\\patmos\c$\test\file2
04/21/2007 15:52:26 ANS1661I Unchanged:
\\patmos\c$\test\file3
04/21/2007 15:52:26 ANS1651I Backed Up:
\\patmos\c$\test\dir1\file1
04/21/2007 15:52:26 ANS1651I Backed Up:
\\patmos\c$\test\dir1\file2
04/21/2007 15:52:26 ANS1651I Backed Up:
\\patmos\c$\test\dir1\file3
04/21/2007 15:52:26 ANS1651I Backed Up:
\\patmos\c$\test\dir2\file1
04/21/2007 15:52:26 ANS1651I Backed Up:
\\patmos\c$\test\dir2\file2
04/21/2007 15:52:26 ANS1651I Backed Up:
\\patmos\c$\test\dir2\file3

```

Related information

For more information about the audit logging facility refer to “Auditlogging” on page 328.

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 Windows 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.

Important:

- When you set autodeploy to yes, if a restart of the client workstation is required to complete the deployment, you cannot disable the restart. The client workstation will be restarted. If it is important that the workstation is not automatically restarted, set autodeploy to noreboot. The deployment will be canceled if a restart is required. The current client is not affected.
- If a restart is required, the deployment manager initiates a restart for the client computer and exits. However, it is possible that you cancel or interrupt the restart. Since the deployment manager is already terminated, a

message is not sent to the server to indicate the failure of the restart. The deployment result is still successful. You must restart the computer so that the new client deployment completes.

No Specifies that the client is not automatically deployed from the server.

NOREboot

Specifies that the deployment manager never automatically restarts the client computer, even if a restart is required. If a restart is required, allowing automatic deployment to many machines with the *NOREboot* parameter can result in only a partial update of, potentially, many clients.

To alleviate this problem, the deployment manager tries to detect if a restart is required. If a restart is required, the deployment manager cancels the deployment before the new client installation. This guarantees that the client computer still has a working backup-archive client, and the new client deployment can be rescheduled.

There are rare cases where the deployment manager cannot detect the restart; for example, if client processes are started from a script. In these cases, the new client installation will continue, but a manual restart of the client computer is required.

Examples

Options file:

autodeploy no

Command line:

Does not apply.

Options file:

autodeploy noreboot

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 3

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 `\\your-node-name\h$` to `\\your-node-name\h$_OLD`. If the new file space name is too long, the suffix replaces the last characters of the file space name, as follows:

`\\your-node-name_OLD`

If the new file space name already exists on the server, the server renames the new file space `\\your-node-name_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, to archive files from your H: disk named \\your-node\h\$, issue the following archive command:

```
arc h:\logs\*.log
```

Before the archive takes place, the server renames the file space to \\your-node\h\$_OLD. The archive places the data specified in the current operation into the Unicode-enabled file space named \\your-node\h\$. The new Unicode-enabled file space now contains only the \logs directory and the *.log files specified in the operation. 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.

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


Supported Clients

This option is valid for all Windows clients. The server can define the `autofsrename` option and override the `autofsrename` setting on the client. The Tivoli Storage Manager API does not support this option.

Options File

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

Syntax



Parameters

Yes

Specifies that the Tivoli Storage Manager server 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
```

Related concepts:

“Restore from file spaces that are not Unicode-enabled” on page 737

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 all Windows 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 filespace**
- **query image**
- **query systemstate**
- **restore image**
- **restore systemstate**

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 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 image -backupsetname=WEEKLY_BSET.21435678
dsmc query backup c:\* -subdir=yes
    -backupsetname=weekly_accounting_data.32145678
dsmc restore image e:
    -backupsetname=weekly_backup_data.12345678
dsmc restore systemstate
    -backupsetname=daily_backup_data.43215678
```

Related information

“Restore data from a backup set” on page 206

Basesnapshotname

The `basesnapshotname` option specifies the snapshot to use as the base snapshot, when you perform a snapshot differential (`snappdiff`) backup of a NetApp filer volume. If you specify this option, you must also use the `snappdiff` 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 Windows 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_121103231100_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 options file (dsm.opt).

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 457

“Schedmode” on page 519

Casesensitiveaware

The `casesensitiveaware` option specifies whether the Windows backup-archive client attempts to filter out file and directory objects that have name conflicts that are caused by different capitalization of the object names.

NTFS and ReFS volumes are case-sensitive and allow case-sensitive file names to be stored. Although the Windows operating system is not case-sensitive, applications such as Windows Services for UNIX (SFU) uses POSIX conventions and allow case-sensitive file names. SFU is typically included with Windows operating systems such as Windows Powered OS and Windows Storage Server. These operating systems are typically deployed on hardware (for example, NAS hardware) which is acting as a dedicated file server in a heterogeneous environment.

If there are UNIX clients that store files on NTFS or ReFS volumes in these Windows file server environments, use the `casesensitiveaware` option. If this option is not used in these environments, unpredictable results occur during backup and archive operations if case-sensitive file name conflicts are encountered. For homogeneous Windows file server environments, the `casesensitiveaware` option is not necessary.

For example, if there is a set of objects that are called 'MyWork.xls', 'MYWORK.xls', and 'mywork.xls', because the Windows operating system is not case-sensitive, applications cannot distinguish between two objects named 'mywork.xls' and 'MyWork.xls'

For this reason, the Windows backup-archive client cannot guarantee the restore integrity of such objects. When a name casing conflict arises, the backup-archive client can guarantee only the restore integrity of the first file in an alphabetical sort. On an ASCII-based operating system such as Windows, this means that capital letters come first, alphabetically, before their lowercase counterparts, so 'MySwor.xls' would alphabetically precede 'mywork.xls'.

In this example, if the `casesensitiveaware` option is used, only 'MyWork.xls' is processed. An error message is issued for 'mywork.xls' and it is skipped. If

'mywork.xls' is a directory, then the directory subtree 'mywork.xls' would be skipped. In all cases, messages are written to both the local error log and to the Tivoli Storage Manager server console to indicate the exact file names of the objects that are skipped.


Supported Clients

This option is valid for all Windows clients. The server can also define this option.

Options File

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

Syntax



Parameters

yes

Specifies that the client will attempt to identify object names which differ in casing only and filter out objects which have casing conflicts and cannot be guaranteed to be restored properly.

no

Specifies that the client will not attempt to identify object names which differ in casing only. This is the default.

Changingretries

The `changingretries` option specifies how many additional times you want the client to attempt to back up or archive a file that is in use. Use this option with the **archive**, **incremental**, and **selective** commands.

This option is applied only when 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. Open file support can be used to back up files that are locked or in use.

Supported Clients

This option is valid for all Windows clients. The server can also define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Backup** tab, **Number of retries if file is in use** field of the Preferences editor.

Syntax

►►—CHAngingretries— *numberretries* —►►

Parameters

numberretries

Specifies the number of times a backup or archive operation is attempted if the file is in use. The range of values is zero through 4; the default is 4.

Examples

Options file:

changingretries 3

Command line:

-cha=3

Class

The class option specifies whether to display a list of NAS or client objects when using the **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 for all Windows 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

►►—CLASS =

<i>client</i>
<i>nas</i>

 —►►

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

Clientview

The `clientview` option is available to users who have upgraded from the IBM Tivoli Storage Manager Express backup client to the IBM Tivoli Storage Manager enterprise client.

You must be connected to the Tivoli Storage Manager Version 5.4 or higher server to use this option. The `clientview` option allows you to choose either the express view or the standard view of the client graphical user interface (GUI).

Supported Clients

This option is valid for all Windows clients.

Options File

Place this option in the `dsm.opt` file. To switch to the Express view:

1. In the backup-archive client GUI, select **Edit > Preference** from the menu bar.
2. From the **General** tab of the Preferences editor, in the **Client View** field, click **Express**.
3. Click **OK** to save your change.

To switch to the Standard view:

1. In the backup-archive client GUI, click **Modify Settings**.
2. From the **General** tab of the Preferences Editor, in the **Client View** field, click **Standard**.
3. Click **OK** to save your change.

Syntax



Parameters

standard

Specifies that the standard, or enterprise, view of the backup-archive client GUI should be used. The standard view contains the advanced features of the backup-archive client GUI. This is the default.

express

Specifies that the express view of the backup-archive client GUI should be used. The express view contains the same features as the Express backup client GUI.

Clusterdiskonly

The `clusterdiskonly` option specifies whether the backup-archive client allows the backup of only clustered disks in specific environments.

The backup-archive client allows for the backup of only clustered disks when the client is running in the following environments:

- In a Microsoft Cluster Server (MSCS)
- When failover clustering is employed on a supported Windows Server client

- In a VERITAS Cluster Server (VCS) environment, when you set `clusternode yes`

The backup-archive client previously allowed only backups and restores of data on clustered drives that were mounted as a drive letter.

It is common to find clustered drives that are mounted as volume mount points. Windows Server operating systems allow users to surpass the 26-drive-letter limitation by allowing volume mount points to be defined on a clustered server. The client can protect data on cluster disks that are mounted as drive letters on Windows Server OS computers. The client can also protect data on cluster disks that are mounted as volume mount points. The backup-archive client can automatically determine whether a volume that is using a volume mount point is a cluster volume.

When you set `clusterdisksonly yes`, the backup-archive client continues to segregate local drives from cluster drives when it evaluates the ALL-LOCAL domain option. When `clusterdisksonly no` is specified, you must explicitly define the backup domains. When `clusterdisksonly no` is specified, the backup-archive client also bypasses enumeration of cluster resources to determine which resources represent cluster drives.


Supported Clients

This option is valid for all supported Windows Server clients.

Options File

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

Syntax



Parameters

Yes

Specifies that the client allows only the processing of cluster drives. Yes is the default.

No

Specifies that the client allows the processing of any disk when `clusternode yes` is set.

Examples

Scenario 1: Back up a node that manages the local (non-clustered) drives and the system state information

This is the node that is dedicated to the restoration of the physical system if a hardware failure occurs. There are no clustered drives that are mounted as volume mount points.

Options file:

```

CLUSTERNODE NO (default)
CLUSTERDISKSONLY YES (default)
DOMAIN ALL-LOCAL (default)
EXCLUDE c:\...\file.txt
  
```

Scenario 1b: Back up a node that manages the local (non-clustered) drives and the system state information and bypass enumeration of cluster resources

This is a scenario similar to scenario 1, which can be deployed if the backup-archive client takes an inappropriate amount of time during startup processing. During initialization of the backup-archive client, all of the cluster resources are enumerated to determine which resources represent cluster disk devices. This processing can be skipped by setting clusterdiskonly no.

Options file:

```
CLUSTERNODE NO (default)
CLUSTERDISKONLY NO
DOMAIN C: D: (local drives must be explicitly enumerated)
EXCLUDE c:\...\file.txt
```

Scenario 2: Back up a node that manages the clustered drives within a cluster resource group and bypass enumeration of cluster resources

This is a scenario that can be deployed if the backup-archive client takes an inappropriate amount of time during startup processing. During initialization of the backup-archive client, all of the cluster resources are enumerated to determine which resources represent cluster disk devices. This processing can be skipped by setting clusterdiskonly no.

Options file:

```
CLUSTERNODE YES
CLUSTERDISKONLY NO
DOMAIN f: g:
EXCLUDE f:\...\file.txt
```

Scenario 3: Back up a node that manages the clustered drives within a cluster resource group, by using volume mount points as cluster resources

In this scenario, it is assumed that the node is responsible for backing up a cluster resource group that has two drives, f: and f:\mnt. There are clustered drives that are mounted as volume mount points (Windows Server operating systems). Ensure that you define the incremental processing domain as only the volumes within a cluster resource group. If you have multiple cluster resource groups, assign a unique client node to manage each cluster resource group.

Options file

```
CLUSTERNODE YES
CLUSTERDISKONLY YES
DOMAIN f: f:\mnt
EXCLUDE f:\mnt\...\file.txt
```

Table 54 lists the clusternode and clusterdiskonly combinations.

Table 54. Clusternode and clusterdiskonly combinations

Clusternode	Clusterdiskonly	When to use
no	yes	This is the default behavior if nothing is specified; since the clusterdiskonly option is set to clusterdiskonly yes, the cluster disk map is built. This combination is used for backing up local drives.

Table 54. *Clusternode and clusterdiskonly combinations (continued)*

Clusternode	Clusterdiskonly	When to use
yes	yes	This is the default way to run in a cluster node to back up cluster disks, including disks that are exposed as mount points; the cluster disk map is built.
yes	no	For clients that run on Windows Server operating systems, you must specify clusterdiskonly no only if you want to bypass cluster volume enumeration for performance reasons.

Clusternode

The `clusternode` option specifies how the backup-archive client manages cluster drives.

The backup-archive client manages clustered drives in the following environments:

- A Microsoft Cluster Server (MSCS)
- Failover Clustering on Windows Server systems
- VERITAS Cluster Server (VCS)

When the `clusternode yes` is set, only shared cluster drives are available for backup and archive processing. When you set `clusternode yes`, the node name defaults to the cluster name.

To back up local drives or Windows Server system state, you must set `clusternode no`.

Note: You must set the `clusternode yes` for all Tivoli Storage Manager-managed cluster operations. Inconsistent use of the `clusternode` option for a given Tivoli Storage Manager cluster node name can cause the cluster node name encrypted password to be invalidated, and prompt the user to reenter the password during the next Tivoli Storage Manager program invocation.

Use the `optfile` option to properly call the correct (cluster) `dsm.opt` for all Tivoli Storage Manager programs to ensure proper functionality for cluster related operations. See the `optfile` option description for more information.


Supported Clients

This option is valid for Windows Server operating system clients.

Options File

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

Syntax



Parameters

Yes

Specifies that you want the client to manage cluster drives in the following environments:

- A MSCS
- Failover Clustering on Windows Server systems
- VCS

No Specifies that you want to back up local disks. This is the default.

Examples

Options file:

```
cluster no
```

Command line:

```
-cluster=yes
```

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

Related information

“Optfile” on page 478

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 Windows clients. The server can also define this option.

Options File

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

Syntax



Parameters

Yes

Specifies that you want the 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 options file (dsm.opt). 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.

NAMedpipes

The interprocess communication method that permits message data streams to pass between a client and a server. Use this communication method with a Tivoli Storage Manager server running on the same workstation as the client.

SHAREdmem

Use the shared memory communication method when the client and server are running on the same system. This provides better performance than the TCP/IP protocol.

Note: Use of this communication method requires that both client and server run under the same Windows account.

Examples

Options file:

```
commmethod tcpip
```

Use only TCP/IP V4

```
commmethod 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 `commmethod V6Tcpip` are specified.

Command line:

```
-commm=tcpip
```

```
-commm=V6Tcpip
```

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

Commrestartduration

The `commrestartduration` option specifies the maximum number of minutes you want the client to try to reconnect to a Tivoli Storage Manager server after a communication error occurs.

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 options file (dsm.opt). You can set this option on the **Communication** tab, **Common Options** section of the Preferences editor.

Syntax

►►—COMMRESTARTDuration— *minutes* —————►►

Parameters

minutes

The maximum number of minutes you want the client to attempt to reconnect with a server after a communication failure occurs. The range of values is zero through 9999; the default is 60.

Examples

Options file:

commrestartduration 90

Command line:

Does not apply.

Commrestartinterval

The commrestartinterval option specifies the number of seconds you want the client to wait between attempts to reconnect to a Tivoli Storage Manager server after a communication error occurs.

Note: Use this option only when commrestartduration is a value greater than zero.

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

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Communication** tab, **Common Options** section of the Preferences editor.

Syntax

►►—COMMRESTARTInterval— *seconds* —————►►

Parameters

seconds

The number of seconds you want the client to wait between attempts to

reconnect with a server after a communication failure occurs. The range of values is zero through 65535; the default is 15.

Examples

Options file:

```
commrestartinterval 30
```

Command line:

Does not apply.

Compressalways

The `compressalways` option specifies whether to continue compressing an object if it grows during compression.

Use this option with the `compression` option, 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 options file (`dsm.opt`). You can set this option on the **Backup** tab, **Continue Compressing if Object Grows** check box of the Preferences editor.

Syntax



Parameters

Yes

File compression continues even if the file grows as a result of compression. This is the default.

No Backup-archive client objects are resent uncompressed if they grow during compression. API behavior depends on the application. Application backups might fail.

Examples

Options file:

```
compressalways yes
```

Command line:

```
-compressa=no
```

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

Compression

The compression option compresses files before you send them to the server.

Compressing your files reduces data storage for backup versions and archive copies of your files. It can, however, affect Tivoli Storage Manager throughput. A fast processor on a slow network connection benefits from compression, but a slow processor on a fast network connection does not.

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

The **backup image** command uses the compression option value specified in the dsm.opt 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 options file (dsm.opt) to exclude specific files or groups of files from compression processing.
- Use the include.compression option in your client options file (dsm.opt) to include files within a broad group of excluded files for compression processing.

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 client options file (dsm.opt). You can set this option on the Backup tab, **Compress objects** check box of the Preferences editor.

Syntax



Parameters

No Files are not compressed before they are sent to the server. This is the default.

Yes

Files are compressed before they are sent to the server.

Examples

Options file:

```
compression yes
```

Command line:

```
-compressi=no
```

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

Related reference:

“Deduplication” on page 360

“Exclude options” on page 400

“Include options” on page 430

Console

Use the console option with the **query systeminfo** command to output information to the console.

- DSMOPTFILE - The contents of the dsm.opt file.
- ENV - Environment variables.

- ERRORLOG - The Tivoli Storage Manager error log file.
- FILE - Attributes for the file name that you specify.
- FILESNOTTOBACKUP - Enumeration of Windows Registry key:

```
HKEY_LOCAL_MACHINE\
  SYSTEM\
    CurrentControlSet\
      BackupRestore\
        FilesNotToBackup
```

This key specifies those files that backup products should not back up. The **query inclexcl** command indicates that these files are excluded per the operating system.

- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- KEYSNOTTORESTORE - Enumeration of Windows Registry key:

```
HKEY_LOCAL_MACHINE\
  SYSTEM\
    ControlSet001\
      BackupRestore\
        KeysNotToRestore
```

This key specifies those Windows Registry keys that backup products should not restore.

- MSINFO - Windows system information (output from MSINFO32.EXE).
- OPTIONS - Compiled options.
- OSINFO - Name and version of the client operating system
- POLICY - Policy set dump.
- REGISTRY - Windows Tivoli Storage Manager-related Windows Registry entries.
- SCHEDLOG - The contents of the Tivoli Storage Manager schedule log (usually dsmsched.log).
- SFP - The list of files protected by Windows System File Protection, and for each file, indicates whether that file exists. These files are backed up as part of the SYSFILES system object.
- SFP=*filename* - Indicates whether the specified file (*filename*) is protected by Windows System File Protection. For example:
SFP=C:\WINNT\SYSTEM32\MSVCRT.DLL
- SYSTEMSTATE - Windows system state information.
- CLUSTER - Windows cluster information.

Note: The **query systeminfo** command is intended primarily as an aid for IBM support to assist in diagnosing problems, although users who are familiar with the concepts addressed by this information might also find it useful. If you use the console option, no special formatting of the output is performed to accommodate screen height or width. Therefore, the console output might be difficult to read due to length and line-wrapping. In this case, 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 417

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 529 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:

- All Windows clients except IA 64 bit

Enter the createnewbase option on the command line. Specify this option with the snapdiff option.

Syntax

►►—Createnewbase—

No
Yes
IGNore
MIGRate

—►►

Parameters

No Specifies that a snapshot difference incremental is run. If 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 83

Related reference:

“Snapdiff” on page 529

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.

Use this option if you want to change the default date format for the language of the message repository you are using.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time you start the client. Consult the documentation on your local system for details about setting up your locale definition.

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 asr**
- **query backup**
- **query filespace**
- **query image**
- **query systemstate**
- **restore**
- **restore image**
- **restore nas**
- **retrieve**
- **restore registry**
- **restore systemstate**
- **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 options file (`dsm.opt`). You can set this option on the **Regional Settings** tab, **Date Format** drop-down list of the Preferences editor.

Syntax

►► `DATEformat` — *format_number* —————►►

Parameters

format_number

Displays the date using one of the following formats. Select the number that corresponds to the date format you want to use:

- 1 MM/DD/YYYY

This is the default for the following available translations:

- US English
 - Chinese (Traditional)
 - Korean
- 2** DD-MM-YYYY
- This is the default for the following available translations:
- Brazilian Portuguese
 - Italian
- 3** YYYY-MM-DD
- This is the default for the following available translations:
- Japanese
 - Chinese (Simplified)
 - Polish
- 4** DD.MM.YYYY
- This is the default for the following available translations:
- German
 - French
 - Spanish
 - Czech
 - Russian
- 5** YYYY.MM.DD
- This is the default for the following available translations:
- Hungarian
- 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`.

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 client options file (dsm.opt). You can set this option on the **Deduplication > Deduplication Cache Location** text box of the Preferences editor. The option can be set in the client option set on 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.

In Uniform Naming Convention (UNC) format, the path must contain a drive letter. In the following UNC format example, the path contains the drive letter D\$: \\computer7\D\$\stgmgr\dedupecache.

Examples

Options file:

dedupcachepath c:\logs\dedup\

Command line:

Does not apply.

Related reference:

“Enablededupcache” on page 389

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 client options file (dsm.opt). 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 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 client options file (dsm.opt). 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 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 430

“Exclude options” on page 400

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



Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc archive c:\foo\*.c -deletefiles  
dsmc rest image c: -incre -deletefiles
```

Description

The description option assigns or specifies a description for files when performing archive, delete archive, retrieve, query archive, or query backupset.

For example, if you want to archive a file named budget.jan and assign to it the description "2002 Budget for Proj 1", you would enter:

```
dsmc archive -des="2003 Budget for Proj 1" c:\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" c:\allproj\*.x
```

You can then use the description to retrieve all of the files.

Examples

Command line:

```
dsmc archive -des="2003 Budget for Proj 1" c:\foo\ *.prj
```

Detail

Use the **detail** option to display management class, file space, backup, 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** option with the **query adobjects** command to display detailed information about Active Directory objects, including all of their attributes.

Use the **detail** option with the **query adobjects** command to display detailed information about Active Directory objects, including all of their attributes.

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 filesystem.
- 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 filesystem**
- **incremental**
- **query adobjects**
- **query archive**
- **query backup**
- **query filesystem**
- **query inclexcl**
- **query mgmtclass**
- **query systemstate**
- **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 filesystem -detail
dsmc query backup file1 -detail
dsmc query systemstate -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. Whatever method is used to create named snapshots, snapshot names that differ only by case will not work properly with the `snappdiff` option. Snapshots that are created by the client will not have the casing problem. Snapshots that are created by methods outside of Tivoli Storage Manager are never deleted by the client.

Supported Clients

This option is valid for all Windows clients except IA 64 bit.

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 `snappdiff` incremental backup from a snapshot that was taken of a network share `//homestore.example.com/vol/vol1` mounted on drive `H:`, where `homestore.example.com` is a file server.

```
incremental -snappdiff H:
```

Perform a `snappdiff` incremental backup from a snapshot that was taken of a network share `//homestore.example.com/vol/vol1` mounted on drive `H:`, where `homestore.example.com` is a file server. The `-diffsnapshot` option value of `LATEST` means the operation uses the latest snapshot (the active snapshot) for volume `H:`.

```
incremental -snappdiff H: -diffsnapshot=latest
```

Related concepts:

“Snapshot differential backup with an HTTPS connection” on page 142

Related tasks:

“Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 83

Related reference:

"Snapdiff" on page 529
"Snapdiffhttps" on page 534
"Createnewbase" on page 354

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 Windows 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 client options file (`dsm.opt`). You can set this option on the **Backup** tab, **Directory Management Class** section in the Preferences editor.

Syntax

►►—`DIRMc`— *—mgmtclassname—*—————►►

Parameters

mgmtclassname

Specifies the name of the management class 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:

```
dirmc 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 273 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 c:*
```

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



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 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 options file (`dsm.opt`).

Syntax



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.

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. Estimate 2 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 400 MB. 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.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

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

Syntax

►►—DISKCACHELocation— —path—◄◄

Parameters

path

Specifies the location where the disk cache database is created if `memoryefficientbackup=diskcachemethod`. The default location is to create the disk cache file in the root of the file space being processed.

In Uniform Naming Convention (UNC) format, the path must contain a drive letter. In the following UNC format example, the path contains the drive letter D\$: `\\computer7\D$\temp\diskcache`.

Examples

Options file:

```
diskcachelocation c:\temp
diskcachelocation c:\tivoli\data
```

Command line:

Does not apply.

See “Include options” on page 430 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 drives to process during an incremental backup:

- 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 the system state:

```
domain all-local -systemstate
```

Example 2: This example uses multiple domain statements to back up all local file systems except for the system state:

```
domain all-local domain -systemstate
```

Example 3: This example excludes the system state from a backup operation. If no other domain statement is used, the result is an empty domain. Nothing is backed up.

```
domain -systemstate
```

If you start the incremental command with a file specification, the client ignores any domain definitions and backs up only the file specification.


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`. You can set this option on the **Backup** tab, **Domain for Backup** section of the Preferences editor.

Syntax



Parameters

all-local

Back up all local volumes on the system, and the Windows system state. This is the default setting. Local volumes are defined as volumes that are formatted with a supported file system (ReFS, NTFS, FAT32, FAT, or GPFS) on a direct-attached storage device, including SAN and iSCSI attached storage. Directories that are mapped to drive letters by using the Windows **subst** command are included in a backup if the mapped directory is on a local disk.

The following types of volumes are not included when `all-local` is specified:

- Network attached volumes, including CIFS shares that are mapped to drive letters.

- Removable volumes, including CD/DVD drives, USB thumb drives, and floppy diskette drives. Some USB-attached hard disks are included in the all-local domain if Windows does not classify them as a removable storage device.

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.

systemstate

Back up the Windows system state. The systemstate domain is included in the all-local domain.

-systemstate

Exclude system state from backup processing.

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 c: d: e:
domain c: systemstate
domain ALL-LOCAL -systemstate
domain ALL-LOCAL -c:
domain ALL-LOCAL -\\florence\e$
```

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="c: d:"
-domain="ALL-LOCAL -c: -systemstate"
```

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 drive. 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 55. 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

Domain.image

The **domain.image** option specifies what you want to include in your client domain for an image backup.

Raw logical volumes must be named explicitly.

If you do not specify a file system with the **backup image** command, the file systems you specify with the **domain.image** option are backed up.

When you specify a file system with the **backup image** command, the **domain.image** option is ignored.

If you do not use the **domain.image** option to specify file systems in your client options file, and you do not specify a file system with the **backup image** command, a message is issued and no backup occurs.

Supported Clients

This option is valid for 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 options file (dsm.opt). 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 d: e: f: domain.image f:\mnt\raw\rawmnt1  
f:\mnt\fs\fsmnt1
```

Command line:

Does not apply.

Domain.nas

The `domain.nas` option specifies the volumes to include in your NAS image backups.

You can specify `all-nas` to include all the mounted file systems on the NAS file server, except those you exclude with the `exclude.fs.nas` option.

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.opt), 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 volumes that you specify on the command line to the volumes defined in your dsm.opt file. For example, if you enter `domain.nas nas1/vol/vol0 nas1/vol/vol1` in your dsm.opt 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.opt file, `all-nas` takes precedence.

Supported Clients

This option is valid for all Windows clients. The server can also define this option.

Options File

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

Syntax



Parameters

domain

Defines the volumes you want to process. You cannot exclude volumes by specifying the dash (-) operator.

all-nas

Processes all mounted volumes on the NAS file server, except those you exclude with the `exclude.fs.nas` option. This is the default. If there is no `domain.nas` statement in the `dsm.opt` file and no volumes specified on the command line, 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 652.

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 755 or “**Query VAPP**” on page 721.

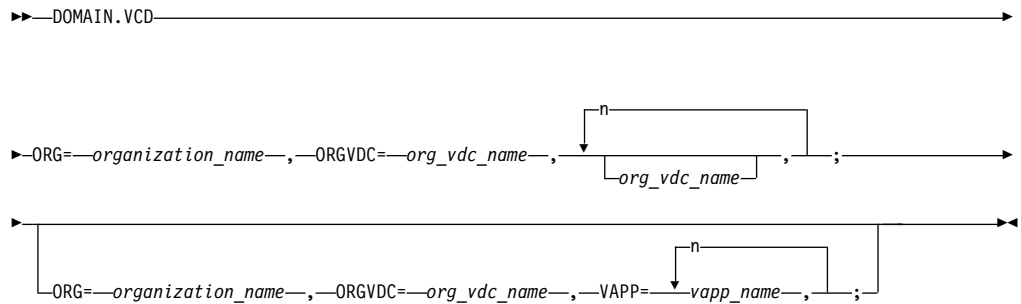
Supported Clients

This option is valid for Windows 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 652, “**Query VAPP**” on page 721, and “**Restore VAPP**” on page 755

Domain.vmfile

The `domain.vmfile` option specifies the virtual machines to include when you perform off-host file-level backups of VMware virtual machines.

VMware file-level backups are only available for Windows guest backups that are performed from a Windows backup proxy. Template virtual machines are not included in file-level backups.

The virtual machines that are specified with the `domain.vmfile` option are only processed when the **backup vm** command is issued without a virtual machine, a list of virtual machines, or the `-vm list` option specified.

You can specify the `all-windows` domain to include all Windows virtual machines that are defined to the Virtual Center or to the ESX Server that is specified on the `vmchost` option.

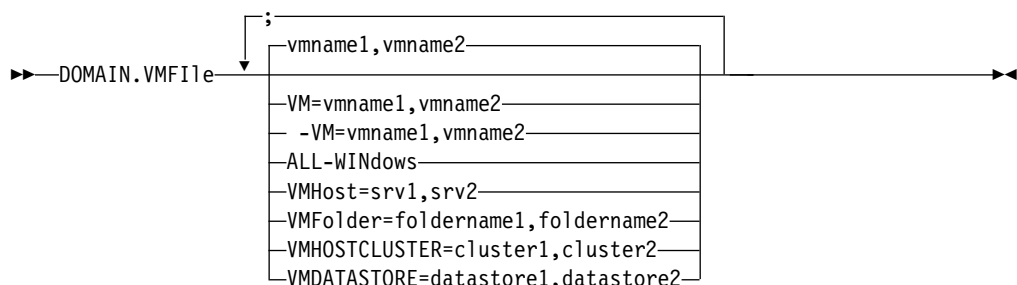
Supported Clients

This option is valid for Windows clients. The server can also define this option. This option is not valid for Hyper-V backups.

Options File

Place this option in the client options file (`dsm.opt`), or on the command line. You can also set this option on the **VM Backup** tab of the Preferences editor.

Syntax



Important: Multiple keywords must be separated by a semicolon. There cannot be any spaces following the semicolons. Multiple machine or domain names must be separated by commas, with no space characters. See `vm=vmname` for examples.

You can use these keywords to easily add groups of virtual machines:

```
all-windows
vmhost
vmfolder
vmhostcluster
vmdatastore
```

When you include virtual machines by using one of these domain levels, the backup-archive client detects any new virtual machines added to the domain after the initial VMware environment was set up.

Parameters

vmname

Defines the virtual machine host name that you want to process or a list of virtual machine host names that are separated by commas.

vm=vmname

The VM keyword specifies that the next set of values are virtual machine host names. VM is the default and it is not required.

-vm=vmname

You can exclude individual virtual machines from the domain by specifying the exclude operator (-) before this keyword. For example, `-VM=vm1,vm2` excludes the vm1 and vm2 virtual machines from a backup operation. You cannot use any domain exclusion operator (-) in front of any domain keywords.

all-windows

Process all virtual machines that are defined to the Virtual Center or the ESX server that is specified on the `vmhost` option that have a guest operating system type of Windows. This parameter also allows for automatic inclusion of newly added virtual machines.

vmhost=hostname

Process 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 newly 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.

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*** keyword applies only if the VM host is the ESXi you connect to. If it is not, a warning level message is issued to the console and is recorded in the client `dsmerror.log` file; it is also recorded as a server event message.

vmfolder=foldername

Process all virtual machines that are defined to the Virtual Center or to the ESX server that is specified on the VMHOST that are defined as residing in the VMware folder that is specified by the *foldername* value. The *foldername* can include multiple VMware folders that are separated by commas. Allows for automatic inclusion of newly added virtual machines.

vmhostcluster=hostclustername

Process 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.

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 issued to the console and is recorded in the client dsmerror.log file; it is also recorded as a server event message.

vmdatastore=datastorename

Process 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 have their configured datastore location match the datastore name that is specified by *datastorename*. The datastore name can include multiple datastores that are separated by commas: VMDATASTORE=datastore1,datastore2

It is possible for virtual machines to 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 located in a different datastore than the one specified on the domain.

Examples

Options file:

Run a file-level backup of all virtual machines that have Windows operating systems.

```
domain.vmfile all-windows
```

Run a file-level VM backup of all virtual machines that are defined to VM folder TSMSvt and TSMdev.

```
domain.vmfile "vmfolder=TSMSvt,TSMdev"
```

File-level VM backup of all virtual machines that are defined to VM folders TSMSvt and TSMdev and virtual machines testvm1 and testvm2.

```
domain.vmfile "vmfolder=TSMSvt;VM=testvm1,testvm2"
```

Process all Windows virtual machines for file-level **backup vm** command, except for the virtual machines named machine1 and machine 2.

```
domain.vmfile "all-windows;-vm=machine1,machine2"
```

Process the virtual machines vm1 and vm2 during a file-level backup.

```
-domain.vmfile="vm1,vm2"
```

Command line:


```
dsmc backup vm -domain.vinfile="vm1,vm2"
```

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 `vmhost` option. The `vmhost` 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 `vmhost`.

For Hyper-V virtual machine backups, you can use this option to specify which Hyper-V virtual machines are processed when you run a **backup vm -vmbackuptype=hypervfull** command, without specifying any Hyper-V virtual machine names. The `domain.vmfull` option applies only to Hyper-V backups that use either `mode=ifincremental` or `mode=iffull`. Those modes also require a Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V license.

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.

Attention: For Microsoft Hyper-V operations, the only valid domain-level parameter for the `domain.vmfull` option is **all-vm**. You can also include one or more virtual machines by using the `VM=` keyword, or exclude virtual machines by using the `-VM=` syntax. As an alternative to **domain.vmfull**, you can also use the `vmlist` option.

Supported Clients

This option can be used with supported Windows clients.

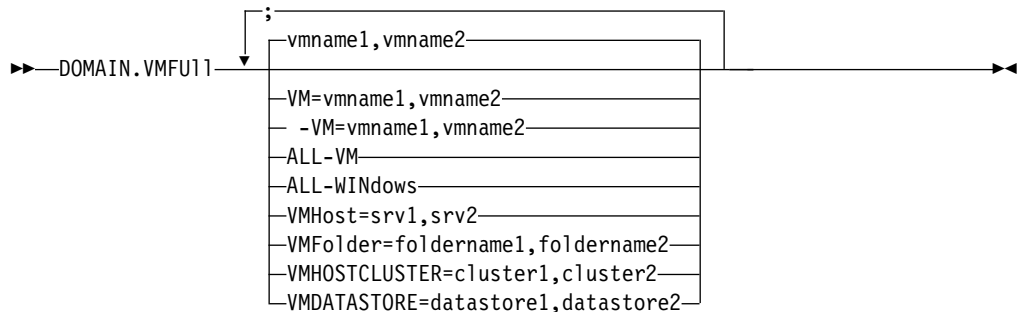
The server can also define this option.

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*.

Syntax for Hyper-V virtual machines



Syntax rules: Multiple keywords must be separated by a semicolon. There cannot be any spaces after the semicolons. Multiple machine or domain names must be separated by commas, with no space characters. For examples, see *vm=vmname*.

Parameters

vmname

Defines the virtual machine name that you want to process. You can supply a list of virtual machine host names, by separating the names with commas (*vm1, vm2, vm5*).

For Hyper-V virtual machines, the names are case-sensitive and must match the capitalization that is shown on the Hyper-V host in the **Hyper-V Manager > Virtual Machines** view.

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.

For Hyper-V virtual machines, this option specifies that a **backup vm** operation processes all Hyper-V virtual machines that are known to the Hyper-V host.

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 405

"Include.vmdisk" on page 440

Enable8dot3namesupport

The enable8dot3namesupport option specifies whether the client backs up and restores short 8.3 names for files that have long names on NTFS file systems.

Supported Clients

This option is valid for all Windows clients.

A file with a long file name might not have a short 8.3 name if short name generation is disabled on the Windows system. This option is effective only for NTFS file systems.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the General tab of the Preferences editor.

Syntax



Parameters

No Short 8.3 names for files with long file names are not backed up or restored. This is the default.

Yes

Short 8.3 names for files with long file names are backed up and restored.

Each short name uses up to 14 additional bytes in the server database.

Although this is a small number, if there are many files with short 8.3 names on many Windows systems, this can increase the size of the Tivoli Storage Manager server database.

Important: Consult with your Tivoli Storage Manager server administrator before you use this option.

The first backup that runs with this option causes all files that have short 8.3 names to be updated on the Tivoli Storage Manager server, even if the files have not otherwise changed. This is because the client is adding the short 8.3 names to the active backup versions.

If this option is enabled for restore, the client attempts to set the short 8.3 name for restored files, even if short name generation is disabled on the Windows system. The client must run under a Windows account that possesses the SE_RESTORE_NAME privilege in order for this option to be effective. See your system administrator if you have questions about account privileges.

During restore, the short 8.3 name of a file is not restored if another object in the same directory already has the same short 8.3 name. In this case, the file is restored and an informational message is logged indicating that the short name could not be set. If the file must be restored with its original short name, you must resolve the conflict with the existing file, and then try the restore again.

Important: This parameter can cause unexpected results in some cases. For example, if the short name of a file changes between the last time the file was backed up and the time it is restored, and there is a link or registry entry that refers to the newer short name, then restoring the file with the older short name invalidates the references to the newer short name.

Examples

Options file:

```
enable8dot3namesupport yes
```

Command line:

```
-enable8dot3namesupport=yes
```

Enablearchiveretentionprotection

The enablearchiveretentionprotection option allows the client to connect to a Tivoli Storage Manager data retention server. 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 client options file (dsm.opt) and is not valid in a client option set from the server. It is not valid on any command line.

Syntax



Parameters

No The data retention server connection is refused. This is the default.

Yes

The client connects to a data retention server.

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 client options file (`dsm.opt`). 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:

-enablededupcache=no

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

Related reference:

"Deduplication" on page 360

"Dedupcachepath" on page 359

"Dedupcachesize" on page 359

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 options file (`dsm.opt`). The option can be set in the client option set on Tivoli Storage Manager server.

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 **dsmsc** 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 c:\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 448

“Instrlogname” on page 449

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 739 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 all Windows clients.

Options File

Place this option in the client options file (`dsm.opt`). 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 450.

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 options file (dsm.opt). You can also set this option on the **Authorization** tab of the Preferences editor. The server can override this.

Syntax



Parameters

AES128

AES 128-bit data encryption. AES 128-bit 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 63 characters in length and include the following characters:

A-Z	Any letter, A through Z, uppercase or lowercase. You cannot specify national language characters.
0-9	Any number, 0 through 9
+	Plus
.	Period
_	Underscore
-	Hyphen
&	Ampersand

Note:

1. The 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 Windows Registry. 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 options file (`dsm.opt`). You can set this option on the **Authorization** tab, **Encryption Key Password** section of the Preferences editor.

Syntax



Parameters

save

The encryption key password is saved in the 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 password can be up to 63 bytes in length. The key is retrieved from the password file on restore and retrieve operations.

When the `save` option is specified for an API application, 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 63 bytes in length.

When the `prompt` option is specified for an API application, the key password must be provided by the application using the API in the `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 options file (dsm.opt).

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.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **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.

Examples

Options file:

```
errorlogname c:\temp\dsmerror.log
```

Command line:

```
-errorlogname=c:\temp\dsmerror.log
```

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

The location of the log file specified using the Client Service Configuration Utility or the client configuration wizard overrides the location specified in the client options file (dsm.opt).

Errorlogretention

The `errorlogretention` option specifies how many days to maintain error log entries before pruning, and whether to save the pruned entries 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 options file (dsm.opt).

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 dsmerrorlog.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 dsmerrorlog.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 dsmerrorlog.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.

You can exclude remotely accessed files by specifying Universal Naming Convention (UNC) names in your exclude statement.

Note:

1. When you exclude a file that was previously included, existing backup versions become inactive during the next incremental backup.
2. The exclude statements are not case sensitive.
3. The server can define exclude options with the `incl excl` option.
4. As with other include-exclude statements, you can use the `incl excl` option to specify a file that can be in Unicode format, containing exclude statements with file names in Unicode.

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 c:\any\test
```

To exclude subdirectories that begin with `test` under the `any` directory, enter the following:

```
exclude.dir c:\any\test*
```

Note: Defining an exclude statement without using a drive letter, such as `exclude.dir code`, excludes the `code` directory on any drive from processing.

Supported Clients

This option is valid for all clients.

Options File

Place these options in the client options file (`dsm.opt`). You can set these options on the **Include-Exclude** tab, **Define Include-Exclude Options** section of the Preferences editor.

Syntax

► *options* — *pattern* —►

exclude, exclude.backup, exclude.file, exclude.file.backup

Use these options to exclude a file or group of files from backup services.

exclude.archive

Excludes a file or a group of files that match the pattern from archive services *only*.

exclude.compression

Excludes files from compression processing if the `compression` option is set to `yes`. This option applies to backups and archives.

exclude.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
- SYSTEMState
- Asr

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 c:\test\dan\data1` excludes the `c:\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.

Note: Defining an exclude statement without using a drive letter, such as `exclude.dir code`, excludes the `code` directory on any drive from processing.

exclude.encrypt

Excludes the specified files from encryption processing. This option does not affect whether files are excluded from backup or archive processing, only whether they are excluded from encryption processing.

exclude.fs.nas

Excludes file systems on the NAS file server from an image backup when used with the **backup nas** command. The NAS node name must be prefixed to the file system name, for example: `netappsj1/vol/vol1`. To apply 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.dir` statements. This option is valid for all Windows clients.

exclude.subfile

Excludes files from adaptive subfile backup processing. This option does not apply to archive processing.

Table 56. System services components and corresponding keywords

Component	Keyword
Background Intelligent Transfer Service	BITS
Event log	EVENTLOG
Removable Storage Management	RSM
Cluster Database	CLUSTERDB
Remote Storage Service	RSS
Terminal Server Licensing	TLS
Windows Management Instrumentation	WMI
Internet Information Services (IIS) metabase	IIS
DHCP database	DHCP
Wins database	WINSDB

Parameters

pattern

Specifies the file or group of files that you want to exclude.

Note: For NAS file systems: You must prefix the NAS node name to the file specification to specify the file server to which the exclude statement applies. If you do not specify a NAS node name, the file system identified refers to the NAS node name specified in the client options file (dsm.opt) or on the command line.

If the pattern begins with a single or double quote or contains any embedded blanks or equal signs, you must surround the value in either single (') or double (") quotation marks. The opening and closing quotation marks must be the same type of quotation marks.

- For the exclude.image option, the pattern is the name of a file system or raw logical volume.

Examples

Options file:

```
exclude ?:\...\swapper.dat
exclude "?:\ea data. sf"
exclude ?:\io.sys
exclude ?:\...\spart.par
exclude c:\*\budget.fin
exclude c:\devel\*
exclude.dir c:\home\jodda
exclude.archive c:\home\*.obj
exclude.encrypt c:\system32\mydocs\*
exclude.compression c:\test\file.txt
exclude.subfile c:\test\file.txt
exclude.fs.nas netappsj\vol\vol0
exclude.dedup c:\Users\Administrator\Documents\Important\...\*
exclude.dedup e:\*\* ieobjtype=image
exclude.dedup ALL ieobjtype=systemstate
exclude.dedup ALL ieobjtype=ASR
```

Command line:

Does not apply.

Related information

See “Exclude files with UNC names” on page 97 for examples of statements using UNC file names.

See “System files to exclude” on page 96 for a list of files that you should always exclude.

“Incl excl” on page 429

See “Include and exclude groups of files with wildcard characters” on page 98 for a list of wildcard characters that you can use. Then, if necessary, use the `include` option to make exceptions.

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.dir` and other include-exclude statements first. The client then considers any `exclude.compression` statements. For example, consider the following include-exclude list:

```
exclude c:\test\*.*
exclude.compression c:\test\file.txt
include c:\test\file.txt
```

The client examines the statements (reading from bottom to top) and determines that the `c:\test\file.txt` file is a candidate for backup, but is not a candidate for compression processing.

- Include-exclude compression processing is valid 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 351

Adaptive subfile backup processing

This topic lists some items to consider if you want to exclude files from adaptive subfile backup processing using the `exclude.subfile` option.

- You must set the `subfilebackup` option to *yes* to enable include-exclude adaptive subfile backup processing. If you do not specify the `subfilebackup` option or you set the `subfilebackup` option to *no*, the backup-archive client does not perform adaptive subfile backup processing.

If you set the `subfilebackup` option to *yes* and no `exclude.subfile` statements exist, the client considers all files for adaptive subfile backup processing.

- The client processes `exclude.dir` and other include-exclude statements first. The client then considers any `exclude.subfile` statements. For example, consider the following include-exclude list:

```
exclude c:\test\*.*
exclude.subfile c:\test\file.txt
include c:\test\file.txt
```

The client examines the statements (reading from bottom to top) and determines that the `c:\test\file.txt` file is a candidate for backup, but is not a candidate for adaptive subfile backup.

- Include-exclude adaptive subfile backup processing is valid for backup and restore processing only. The `exclude.subfile` option does not affect whether files are excluded from backup or restore processing, only whether they are excluded from adaptive subfile backup processing.
- As with other include-exclude statements, you can specify `exclude.subfile` statements in Unicode using the `incl excl` option.

The `subfilebackup` option does not work correctly for migrated files. If you use a combination of `subfilebackup` and `non-subfilebackup` for migrated files, your data might be corrupted on the server.

Related reference:

“`Incl excl`” on page 429

“`Subfilebackup`” on page 549

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`” on page 405

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 Windows 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. 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 658 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 658

“**Restore VM**” on page 758

“`Domain.vmfull`” on page 381

“`Include.vmdisk`” on page 440

“`INCLUDE.VMTSMVSS`” on page 443

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 all Windows 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 all Windows clients.

Options File

None. You can specify this option only on the command line. The server can also define or override this option.

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=\\myFbServer.company.com\REP
```

Backs up all volumes for FastBack clients fbclient1 and fbclient2 that are found in policy Policy1.

Command line:

```
dsmc backup fastback -fbpolicyname=Policy1  
-fbclientname=fbclient1  
-fbvolume=c:,f: -fbserver=myFbServer  
-fbreposlocation=\\myFbServer.company.com\REP
```

Backs up volumes C:\ and F:\ for FastBack client fbclient1 found in policy Policy1.

Command line:

```
dsmc backup fastback -fbpolicyname=Policy1  
-fbclientname=fbWindowsClient,fbLinuxClient  
-fbserver=myFbServer  
-fbreposlocation=\\myFbServer.company.com\REP
```

Backs up all volumes for FastBack client fbWindowsClient found in policy Policy1.

The volumes for Linux FastBack client `fbLinuxClient` will not be backed up from the Windows backup-archive client. To back up or archive volumes from a Linux FastBack client, use the Linux backup-archive client.

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 Windows volumes will be backed up or archived to the Tivoli Storage Manager server by the Windows 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 all Windows clients.

Options File

None. You can specify this option only on the command line. The server can also define or override this option.

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.

On Windows systems, you do not need to specify the fbreposlocation option when the backup-archive client is installed on a DR Hub server or a FastBack server workstation. When the backup-archive client is installed on a dedicated client proxy, the repository location fbreposlocation option is required.

If you specify the fbreposlocation option for a FastBack DR Hub, specify only the base directory of the DR Hub repository with this option. Then use the fbbranch option to indicate the Branch ID of the server to back up. If you specify the fbreposlocation option for the FastBack server, use the format \\<fbserver>\REP. In this case, do not use the fbbranch option.

Supported Clients

This option is valid for all Windows 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:

The `fbreposlocation` option is only required on a dedicated proxy machine. If the `fbreposlocation` option is specified on a machine where a FastBack server or FastBack DR Hub is installed, it is ignored.

Use this command when the Tivoli Storage Manager dedicated proxy client is connecting to a remote Tivoli Storage Manager FastBack server repository:

```
dsmc backup fastback -fbpolicyname="Policy 1" -fbserver=myFbServer  
-fbreposlocation=\\myFbServer.company.com\REP
```

A repository location is required.

`myFbServer` is the short host name of the machine where the FastBack server is installed.

Command line:

Use this command when the Tivoli Storage Manager dedicated proxy client is connecting to a remote repository on a Tivoli Storage Manager FastBack disaster recovery hub:

```
dsmc backup fastback -fbpolicyname="Policy 1" -fbserver=myFbServer  
-fbreposlocation=\\myfbdrhub.company.com\REP  
-fbbranch=aFbServerBranch
```

A repository location is required.

The `myFbServer` parameter specifies the short host name of the Tivoli Storage Manager FastBack Server whose FastBack branch is specified using the `FBBranch` option.

The `fbbranch` option specifies the branch ID of a FastBack server on the disaster recovery hub.

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 all Windows clients.

Options File

None. You can specify this option only on the command line. The server can also define or override this option.

Syntax

► — -FBServer — *server_name* —►

Parameters

server_name

Specifies the short hostname of the machine on which a FastBack server is installed.

Examples

Command line:

The Tivoli Storage Manager backup-archive client is running on a FastBack server machine whose short name is myFbServer:

```
dsmc archive fastback -fbpolicyname=Policy1 -fbserver=myFbServer
```

Command line:

The Tivoli Storage Manager backup-archive client is running on a Tivoli Storage Manager FastBack DR Hub machine and is connecting to a FastBack Server branch repository branch1. The short host name of the FastBack server is myFbServer:

```
dsmc archive fastback -fbpolicyname=Policy1 -fbserver=myFbServer  
-fbbranch=branch1
```

Command line:

The backup-archive client is running on a dedicated proxy machine and is connecting to a remote FastBack server repository. The FastBack server is installed on a machine whose short name is myFbServerMachine:

```
dsmc archive fastback -fbpolicyname=Policy1 -fbserver=myFbServerMachine  
-fbreposlocation=\\myFbServerMachine.company.com\Rep
```

Command line:

The backup-archive client is running on a dedicated proxy machine and is connecting to a remote Tivoli Storage Manager FastBack repository on a FastBack DR Hub. The FastBack Server with branch ID branch1 is installed on a machine whose short name is myFbServer.

```
dsmc backup fastback -fbpolicyname=Policy1 -fbserver=myFbServer  
-fbreposlocation=\\myDrHubMachine.company.com\Rep  
-fbbranch=branch1
```

Fbvolumename

Use the fbvolumename option with the **backup fastback** or **archive fastback** commands.

Command line:

```
dsmc archive fastback -fbpolicyname=Policy1 -fbclientname=client1  
-fbvolumename=c:,f: -fbserver=myFbServer  
-fbreposlocation=\\myFbServer.company.com\REP
```

Archives volumes C: and F: 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 623.
- 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 `quotessareliteral` 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 497. If `quotessareliteral` 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:

```
c:\myfiles\directory\file1  
c:\tivoli\mydir\yourfile.doc  
..\notes\avi\dir1  
..\fs1\dir2\file3  
"d:\fs2\Ha Ha Ha\file.txt"  
"d:\fs3\file.txt"
```

To override standard processing of quotation marks and wildcard characters, see “Quotesareliteral” on page 497 and “Wildcardsareliteral” on page 623.

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, `d:\dir\` represents the restore destination for all entries in the filelist.

```
restore -filelist=c:\filelist.txt d:\dir\
```

However, in the following **selective** command, the file specification `d:\dir\` is ignored.

```
selective -filelist=c:\filelist.txt d:\dir\
```

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 `c:\dir0\myfile` and `c:\dir0`. If the `\dir0` directory does not exist on the server, the `c:\dir0` directory is sent to the server a second time.

Supported Clients

This option is valid for all 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=c:\avi\filelist.txt
```

Related information

“Quotesareliteral” on page 497

“Wildcardsareliteral” on page 623

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.
- ENV - Environment variables.
- ERRORLOG - The Tivoli Storage Manager error log file.
- FILE - Attributes for the file name that you specify.
- FILESNOTTOBACKUP - Enumeration of Windows Registry key:

```
HKEY_LOCAL_MACHINE\  
  SYSTEM\  
    CurrentControlSet\  
      BackupRestore\  
        FilesNotToBackup
```

This key specifies those files that back up products should not back up. The **query inclexcl** command indicates that these files are excluded per the operating system.

- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- KEYSNOTTORESTORE - Enumeration of Windows Registry key:

```
HKEY_LOCAL_MACHINE\  
  SYSTEM\  
    ControlSet001\  
      BackupRestore\  
        KeysNotToRestore
```

This key specifies those Windows Registry keys that back up products should not restore.

- MSINFO - Windows system information (output from `MSINFO32.EXE`).
- OPTIONS - Compiled options.
- OSINFO - Name and version of the client operating system.
- POLICY - Policy set dump.
- REGISTRY - Tivoli Storage Manager-related Windows Registry entries.
- SCHEDLOG - The contents of the schedule log (usually `dsmsched.log`).
- SFP - The list of files protected by Windows System File Protection, and for each file, indicates whether that file exists. These files are backed up as part of the `SYSFILES` system object.
- SFP=*filename* - Indicates whether the specified file (*filename*) is protected by Windows System File Protection. For example:
`SFP=C:\WINNT\SYSTEM32\MSVCRT.DLL`
- SYSTEMSTATE - Windows system state information.
- CLUSTER - Windows cluster information.

Note: The **query systeminfo** command is intended primarily as an aid for IBM support to assist in diagnosing problems, although users who are familiar with the concepts addressed by this information might also find it useful. If you use the console option, no special formatting of the output is performed to accommodate screen height or width. Therefore, the console output might be difficult to read due to length and line-wrapping. In this case, 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 352

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
```

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 options file (`dsm.opt`).

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 60

Related tasks:

“Configuring the client for automated failover” on page 63

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 c:\Windows\Program  
Files\*.exe
```

Fromnode

The `fromnode` option permits one node to perform commands for another node. A user on another node must use the **set access** command to permit you to query, restore, or retrieve files for the other node.

Use the `fromnode` option with the following commands:

- **query archive**
- **query backup**
- **query filespace**
- **query group**
- **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 d:\
```

Note: The backup-archive client can use file space information when restoring files. The file space information can contain the name of the computer from which the files were backed up. If you restore from another backup-archive client node and do not specify a destination for the restored files, the client uses the file space information to restore the files. In such a case, the client attempts to restore the files to the file system on the original computer. If the restoring computer has access to the file system of the original computer, you can restore files to the original file system. If the restoring computer can not access the file system of the original computer, the client can return a network error message. If you want to restore the original directory structure but on a different computer, specify only the target file system when you restore. This is true when restoring files from another node and when retrieving files from another node.

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 c:\*
```

Groupname

Use the `groupname` option with the **backup group** command to specify the name for a group. You can only perform operations on new groups or the current active version of the group.

Supported Clients

This option is valid for all Windows clients.

Syntax

►►GROUPName =— —*name*◄◄

Parameters

name

Specifies the name of the group which 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=c:\dir1\filelist1 -groupname=group1  
-virtualfsname=\virtfs -mode=full
```

Gui treeview after backup

The `gui treeview after backup` 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). You can set this option on the **General** tab, **Return to tree window after function completed** check box of the Preferences editor.

Syntax



Parameters

No Returns you to the Tivoli Storage Manager main window after a successful operation completes. This is the default.

Yes

Returns you to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.

Examples

Options file:

```
guitreeviewafterbackup yes
```

Command line:

Does not apply.

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.opt). 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:

-httpport=1502

Hsmreparsetag

The hsmreparsetag option specifies a unique reparse tag that is created by an HSM product installed on your system.

Many HSM products use reparse points to retrieve or recall migrated files. After a file is migrated, a small stub file, with the same name as the original file, is left on the file system. The stub file is a reparse point that triggers a recall of the original file when a user or application accesses the stub file. The reparse point includes a unique identifier called a *reparse tag* to identify which HSM product migrated the file.

If the Tivoli Storage Manager backup-archive client does not recognize the reparse tag in a stub file, the Backup-Archive Client causes the HSM product to recall the original file. You can prevent files from being recalled if you specify the reparse tag with the hsmreparsetag option.

The backup-archive client recognizes the reparse tag of HSM products from the following companies:

- International Business Machines Corp.
- Wisdata System Co. Ltd.
- BridgeHead Software Ltd.
- CommVault Systems, Inc.
- Data Storage Group, Inc.
- Enigma Data Solutions, Ltd.
- Enterprise Data Solutions, Inc.
- Global 360
- GRAU DATA AG
- Hermes Software GmbH
- Hewlett Packard Company
- International Communication Products Engineering GmbH
- KOM Networks
- Memory-Tech Corporation
- Moonwalk Universal
- Pointsoft Australia Pty. Ltd.

- Symantec Corporation

If the HSM product you use is not in the preceding list, use the `hsmreparsetag` option to specify the reparse tag. Ask your HSM product vendor for the reparse tag used by the product.

Supported clients

This option is valid for all Windows clients.

Option file

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

Syntax

►—HSMREPARSETAG—*reparse_tag_value*—►

Parameters

reparse_tag_value

A decimal (base 10) or hexadecimal (base 16) value that specifies the reparse tag.

Examples

Options file:

Specify an HSM reparse tag in decimal format:

```
hsmreparsetag 22
```

Specify an HSM reparse tag in hexadecimal format:

```
hsmreparsetag 0x16
```

Command line:

Does not apply.

Ieobjtype

Use the `ieobjtype` option to specify an object type for a client-side data deduplication operation within include-exclude statements.

The `ieobjtype` option is an additional parameter to the `include.dedup` or `exclude.dedup` options.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API also supports this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Include/Exclude** tab of the Preferences editor. The option can be set in the client option set on Tivoli Storage Manager server.

Syntax



Parameters

File

Specifies that you want to include files for, or exclude files from, client-side data deduplication processing. File is the default.

Image

Specifies that you want to include images for, or exclude images from, client-side data deduplication processing.

System State

Specifies that you want to include system state for, or exclude system state from, client-side data deduplication processing.

Asr

Specifies that you want to include automatic system recovery objects for, or exclude ASR objects from, client-side data deduplication processing.

Examples

Options file:

```
exclude.dedup e:\*\* ieobjtype=image
```

Command line:

Does not apply.

Related reference:

"Exclude options" on page 400

"Include options" on page 430

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 -ifnewer d:\logs\*.log
```

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.

Place the `include.image` statement containing the `imagegapsize` value in your `dsm.opt` file.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager API does not support this option.

Options File

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

Syntax

►►—IMAGEGapsize— *size* —►►

Parameters

size

Specifies the minimum size of empty regions in a formatted logical volume

that should be skipped during an image backup. You can specify k (kilobytes) m (megabytes) or g (gigabytes) qualifiers with the value. Without a qualifier, the value is interpreted in KB. 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.

Note: Because of operating system limitations, use this option for NTFS file systems only. If you specify an `imagegapsize` that is greater than 0 for a file system other than NTFS, you get a warning message.

Examples

Options file:

```
imagegapsize 1m
```

Include-exclude list example: `include.image h: MYMC imagegapsize=1m`

Command line:

```
-imagegapsize=64k
```

Imagetofile

Use the `imagetofile` option with the **restore image** command to specify that you want to restore the source image to a file.

You might need to restore the image to a file if bad sectors are present on the target volume, or if you want to manipulate the image data. Later, you can use a data copy utility of your choice to transfer the image from the file to a disk volume.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager API does not support this option.

Syntax

►—IMAGETOfile—————►

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore image d: e:\diskD.img -imagetofile
```

Inactive

Use the `inactive` option to display both active and inactive objects.

You can use the `inactive` option with the following commands:

- **delete group**
- **query asr**
- **query backup**

- **query image**
- **query nas**
- **query systemstate**
- **query vm** (vmbackuptype=fullvm and vmbackuptype=hypervfull)
- **restore**
- **restore group**
- **restore image**
- **restore nas**
- **restore systemstate**
- **restore vm** (vmbackuptype=fullvm and vmbackuptype=hypervfull)

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 -inactive c:\id\projecta\ -pick
```

Incl excl

The `incl excl` option specifies the path and file name of an include-exclude options file.

Multiple `incl excl` statements are permitted. However, you must specify this option for each include-exclude file.

Ensure that you store your include-exclude options file in a directory to which all users have read access.

When processing occurs, the include-exclude statements within the include-exclude file are placed in the list position occupied by the `incl excl` option, in the same order, and processed accordingly.

Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Include-Exclude** tab of the Preferences editor.

Syntax

►—INCLExcl— *filespec* —►

Parameters

filespec

Specifies the path and file name of *one* include-exclude options file.

Examples

Options file:

```
incl excl c:\dsm\backup.excl
```

Command line:

Does not apply.

Related information

For more information about creating an include-exclude options file, see “Creating an include-exclude list” on page 92.

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.

Windows users: Create an include-exclude file in Unicode format by performing the following steps:

1. Open Notepad.
2. Enter your include and exclude statements. You might need to copy file names with characters from other code pages using Microsoft Windows Explorer.
3. Click **File** and then click **Save As**.
4. Select the **Save as Unicode** check box, specify the file and target directory, and then save the file.
5. Place an `incl excl` option specifying the include-exclude file you just created in your client options file (dsm.opt).
6. Restart the 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, archive, and image services.
- Files that are included for backup or archive processing that you want to include for encryption processing.
- Files that are included for backup or archive processing that you also want to include for compression processing.
- Files that are included for backup processing that you also want to include for adaptive subfile backup processing.
- Objects to which you want to assign a specific management class.
- A management class to assign to all objects to which you do not explicitly assign a management class.
- File spaces to which you want to assign memory-efficient backup processing
- File spaces where you want to use the `diskcache` location 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.

You can include remotely accessed files by specifying Universal Naming Convention (UNC) names in your include statement.

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.dir` statement overrides all include statements that match the pattern.
2. The include statements are not case-sensitive.
3. The server can also define these options with the `incl excl` option.

Supported Clients

This option is valid for all clients. The server can also define `include.fs.nas`.

Options File

Place these options in the client options file (`dsm.opt`). 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 c:\test\*\ archmc
include.backup c:\test\*
```

include.archive

Includes files or assigns management classes for archive processing.

include.compression

Includes files for compression processing if you set the `compression` option to `yes`. This option applies to backups and archives.

include.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
- SYSTEMState
- Asr

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

If open file support has been configured, the client performs a snapshot backup or archive of files that are locked (or in use) by other applications. The snapshot allows the backup to be taken from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to the file

system are not included in the backup. You can set the `snapshotproviderfs` parameter of the `include.fs` option to `none` to specify which drives do not use open file support.

To control how the client processes your file space for incremental backup, you can specify these additional options in your `dsm.opt` file as values of the `include.fs` option: `diskcachelocation` and `memoryefficientbackup`.

```
include.fs d: memoryefficientbackup=diskcachem
           diskcachelocation=e:\temp
include.fs e: memoryefficientbackup=diskcachem
           diskcachelocation=c:\temp
```

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 client options file (`dsm.opt`).

include.image

Includes a file space or logical volume, or assigns a management class when used with the **backup image** command. The **backup image** command ignores all other include options.

By default, the client performs an offline image backup. To enable and control an online image operation, you can specify these options in your `dsm.opt` file as values of the `include.image` option: `snapshotproviderimage`, `presnapshotcmd`, `postsnapshotcmd`.

include.subfile

Includes files for adaptive subfile backup processing. This option does not apply to archive processing. This option is valid for all Windows clients.

include.systemstate

This option binds system state backups to the specified management class. If you specify this option, specify all as the pattern. If you do not specify this option system state backups are bound to the default management class.

Parameters

pattern

Specifies the objects to include for backup or archive processing or to assign a specific management class.

Note: For NAS file systems: You must prefix the NAS node name to the file specification to specify the file server to which the include statement applies. If you do not specify a NAS node name, the file system identified refers to the NAS node name specified in the client options file (`dsm.opt`) or on the command line.

If the pattern begins with a single or double 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 file system or raw logical volume.

Note: When you specify `include.systemstate`, the only valid pattern is **all**.

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 57. Other optional parameters

optional_parameter	Use with option
ieobjtype "Ieobjtype" on page 425	include.dedup
memoryefficientbackup "Memoryefficientbackup" on page 462	include.fs
diskcachelocation "Diskcachelocation" on page 370	include.fs
postsnapshotcmd "Postsnapshotcmd" on page 486	include.image
presnapshotcmd "Presnapshotcmd" on page 492	include.image
snapshotproviderfs "Snapshotproviderfs" on page 535	include.image
snapshotproviderimage "Snapshotproviderimage" on page 536	include.image

Examples

Options file:

Windows only:

```
include c:\proj\text\devel.*
include c:\proj\text\* textfiles
include ?:\* managall
include WAS_ND_NDNODE mgmtclass
```

```

include WAS_APPNODE mgmtclass
include.backup c:\win98\system\* mybackupclass
include.archive c:\win98\system\* myarchiveclass
include.encrypt c:\win98\proj\gordon\*
include.compress c:\test\file.txt
include.subfile c:\test\file.txt
include.image h: MGMTCLASSNAME
    snapshotproviderimage=vss

include.image x:
    snapshotproviderimage=none
include.image y:
    snapshotproviderimage=vss
include.image z: MGMTCLASSNAME
    snapshotproviderimage=none
include.fs c:
    snapshotproviderfs=vss

include.systemstate ALL mgmtc3
include.dedup c:\Users\Administrator\Documents\Important\...\*
include.dedup e:\*\* ieobjtype=image
include.dedup ALL ieobjtype=systemstate
include.dedup ALL ieobjtype=ASR

```

To encrypt all files on all drives:

```
include.encrypt ?:\...\*
```

Command line:

Does not apply.

Related concepts:

“Exclude files with UNC names” on page 97

Related tasks:

“Configuring Open File Support” on page 83

Compression, encryption, and adaptive subfile backup processing

This topic lists some items to consider if you want to include specific files or groups of files for compression, encryption, and adaptive subfile backup processing during a backup or archive operation.

- 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.
- You must set the subfilebackup option to *yes* to enable include-exclude adaptive subfile backup processing. If you do not specify the subfilebackup option or you set the subfilebackup option to *no*, the backup-archive client does not perform adaptive subfile backup processing.
- The client processes `exclude.dir` and other include-exclude statements first. The client then considers any `include.compression`, `include.encrypt` and `include.subfile` statements. For example, consider the following include-exclude list:

```

exclude c:\test\file.txt
include.compression c:\test\file.txt
include.encrypt c:\test\file.txt
include.subfile c:\test\file.txt

```

The client examines the `exclude c:\test\file.txt` statement first and determines that `c:\test\file.txt` is excluded from backup processing and is, therefore, not a candidate for compression, encryption, or adaptive subfile backup processing.

- Include-exclude compression and encryption processing is valid for backup and archive processing *only*.
- Include-exclude adaptive subfile backup processing is valid for backup and restore processing *only*.
- As with other include-exclude statements, you can use the `incl excl` option to specify a file that is in Unicode format, which contains `include.compression`, `include.encrypt` and `include.subfile` statements specifying Unicode files. See “`Incl excl`” on page 429 for more information.

Related reference:

“Compression” on page 351

“Subfilebackup” on page 549

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: `/vol/vol0`.
- NAS file system designations that are specified on the command line require brace delimiters ({ and }) around the file system names, such as: `{/vol/vol0}`. Do not use brace delimiters if you specify this option in the option file.

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 563.

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 440

“INCLUDE.VMTSMVSS” on page 443

“INCLUDE.VMSNAPSHOTATTEMPTS” on page 441

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 Windows 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 vAapp. 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 571

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 Windows clients that are configured to back up VMware or Hyper-V 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. 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 601

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 Windows 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. 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 658 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 658

“**Restore VM**” on page 758

“Domain.vmfull” on page 381

“Exclude.vmdisk” on page 405

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 Windows 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

This is a required positional parameter. It specifies the name of the virtual machine to attempt the total number of snapshots for, if a backup attempt fails due to snapshot failure.

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 Windows 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. 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 621

"`Exclude.vmdisk`" on page 405

"`Include.vmdisk`" on page 440

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.

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 NTFS security information, 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.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—INCRemental—◄◄

Examples

Command line:

```
res i d: -incremental
```

Incrthreshold

The `incrthreshold` option specifies the threshold value for the number of directories in any journaled file space that might have active objects on the server, but no equivalent object on the workstation.

When a Windows client deletes a file or directory with a long name, it sometimes reports this using a compressed name. After the object is deleted, the compressed name might be reused and the deletion notice can no longer identify a unique object. During a journaled incremental backup of a file space, this can result in the *no active version* response from the server resulting in an unsuccessful expire for an object.

The `incrthreshold` option allows you to specify what to do when this condition arises:

- If you set the `incrthreshold` option to 0 (the default), no action is taken. The primary consequence is that, during a restore of such a directory, these objects might be inadvertently restored. When the next non-journaled incremental backup is run on this directory, the Tivoli Storage Manager server expires all objects in the directory that exist on the server but not on the workstation.
- If you specify a value greater than zero, the client saves the directory name of an object in the journal during journaled backups. During a full file space journaled incremental backup, if the number of directories in the file space is greater than or equal to this value, a full incremental backup of each directory occurs. This takes place automatically after completion of the journaled backup and does not require entry of another command.
- If you set the `incrthreshold` option to 1, the client performs a full incremental backup of these directories whenever a *no active version* response is received during a full file space journaled incremental backup.

Supported Clients

This option is for all Windows clients.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Backup > Threshold for non-journal incremental backup** field of the Preferences editor.

Syntax

►►—`INCRThreshold—numberdirectories`—————►►

Parameters

numberdirectories

Specifies the threshold value for the number of directories in any journaled file space that might contain active files that should be expired. When this threshold is reached during a full file space journaled incremental, the client initiates an incremental backup on each such directory at the completion of the journaled backup. The range of values is 0 through 2,000,000,000; the default is 0.

Examples

Options file:

```
incrthreshold 1
```

Command line:

```
-increthreshold=1
```

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

Related information

See “Incremental” on page 683 for more information about journaled backups.

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.

If you do not specify the `instrlogmax` option, the instrumentation log can grow without any limit on its size. You must manually manage the log contents to prevent the log from depleting disk resources.

Supported Clients

This option is valid for all clients and the Tivoli Storage Manager API.

Options File

Place this option in the client options file (`dsm.opt`). The option can be set in the client option set on the Tivoli Storage Manager server.

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 c:\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 390

“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 options file (`dsm.opt`). The option can be set in the client option set on the Tivoli Storage Manager server.

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.

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).

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 Windows clients:

```
instrlogname c:\mydir\mydsminstr.log
```

Command line:

For Windows clients:

```
dsmc sel c:\mydir\* -subdir=yes -instrlogname=c:\temp\mydsminstr.log
```

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

Related reference:

“Enableinstrumentation” on page 390

“Instrlogmax” on page 448

Journalpipe

The `journalpipe` option specifies the pipe name of a journal daemon session manager to which the backup clients attach.

Supported Clients

This option is for all Windows clients.

Options File

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

```
JournalPipe \\.\pipe\jnlSessionMgr1
```

Syntax

►►—JOURNALPipe—pipename—————◄◄

Parameters

pipename

Specify the name of the pipe the client attaches to when performing a journal-based backup. The default pipe name is `\\.\pipe\jnlSessionMgr`.

Examples

Options file:

```
JOURNALPipe \\.\pipe\jnlSessionMgr
```

Command line:

This option cannot be set on the command line.

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 TCPip` in the client options file (`dsm.opt`).

For Windows, use the `lanfreeshmport` option to uniquely identify the storage agent to which the client is trying to connect.

Supported Clients

This option is valid for all Windows clients.

Options File

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

Syntax

►►—LANFREECommmethod— *commmethod* —►►

Parameters

commmethod

Specifies the supported protocol for 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.

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.

NAMedpipes

The interprocess communication method that permits message data streams to pass between a client and a server. This is the default. Do not specify the `lanfreetcpport` option if you want to use the `NAMedpipes` communication method for LAN-free communication.

SHAREdmem

Use the shared memory communication method when the client and Storage Agent are running on the same system. Shared memory provides better performance than the TCP/IP protocol. The backup-archive client must have local administrator permissions.

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 453

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 all Windows clients.

Options File

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

Syntax

►—LANFREEShmport— *port_address*—►

Parameters

port_address

Specifies the number that is used to connect to the storage agent. The range of values is 1 through 32767.

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 450

Lanfreetcppport

The lanfreetcppport 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 lanfreetcppport option if you want to use the NAMEDpipes communication method for LAN-free communication.

Supported Clients

This option is valid for all Windows clients.

Options File

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

Syntax

►►—LANFREETCPPort— *port_address*————►►

Parameters

port_address

Specifies the TCP/IP port number where the Storage Agent is listening. The range of values is 1 through 32767; the default is 1500.

Note: The client lanfreetcppport 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 450

Lanfreessl

Use the lanfreessl option to enable Secure Sockets Layer (SSL), to provide secure client and Storage Agent communications.

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 `lanfreecommethod=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 for all supported Windows clients.

Options File

Place this option in the client system-options file.

Syntax



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.

Language

The language option specifies the national language in which to present client messages.

You can use US English (ENU) with all clients.

The language that is displayed by the backup-archive client Java GUI is defined by the Windows display locale and not the Windows system locale. For example, if the Windows system and input locale is French, but the display locale is Russian, the language that is displayed by the Java GUI is Russian by default, if the language option is not used. If you want to display the Java GUI in US English or another language, you can override the default display language by specifying the language option.

Tip: The language option does not affect the web client. The web client displays in the language associated with the locale of the browser. If the browser is running in a locale that client does not support, the web client is displayed in US English.

Supported Clients

This option is valid for all Windows clients.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Regional Settings** tab, **Language** drop-down list of the Preferences editor.

Syntax

►►—LANGUAGE— *language*—►►

Parameters

language

Specifies the language that you want to use. The available languages include:

- ENU (English, United States).
- PTB (Brazilian Portuguese)
- CHS (Chinese, Simplified)

- CHT (Chinese, Traditional)
- FRA (Standard French)
- DEU (Standard German)
- ITA (Standard Italian)
- JPN (Japanese)
- KOR (Korean)
- ESP (Standard Spanish)
- CSY (Czech)
- HUN (Hungarian)
- PLK (Polish)
- RUS (Russian)

Examples

Options file:

language enu

Command line:

Does not apply.

Latest

Use the `latest` option 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 c:\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 `dsm.opt` file.

Syntax



Parameters

No Specifies that the GUI attempts initial logon with the server and enables all functions. This is the default.

Yes Specifies that the GUI does not attempt initial logon with the server and enables only the restore functionality.

Examples

Options file:

```
localbackupset yes
```

This option is not valid with the **dsmc** command-line client. (Does not apply to Mac OS X)

Related information

“Restore Backupset” on page 739

Manageservices

The `manageservices` 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. Set the `passwordaccess` option to `generate` in your client options file (`dsm.opt`) and generate a password, so Tivoli Storage Manager can manage your password automatically.

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.

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.opt`). 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.

schedule

Specifies that the client acceptor daemon manages the scheduler. Both `webclient` and `schedule` are the defaults for Mac OS X.

Examples

Options file:

The following are examples of how you might specify the `managedservices` option in your client options file (`dsm.opt`).

Task Specify that the client acceptor daemon manages 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.

Command line:

Does not apply.

Related information

“Passwordaccess” on page 480

See “Configuring the scheduler” on page 36 for instructions to set up the client acceptor daemon to manage the scheduler.

“Sessioninitiation” on page 523

“Cadlistenonport” on page 338

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 options file (`dsm.opt`). 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.

This option is valid for clients that act as data mover nodes that protect VMware virtual machines or on clients that protect Microsoft Hyper-V virtual machines. 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.

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.

This option is valid for clients that act as data mover nodes that protect VMware virtual machines or on clients that protect Microsoft Hyper-V virtual machines. 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.

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 disk cache file created by the initial disk cache incremental backup can require up to 5 GB of disk space for each million files or directories being backed up.

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. Estimate 2 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 400 MB. 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.


Supported Clients

This option is valid for all clients. The server can also define this option.

Options File

Place this option in the client user-options file (`dsm.opt`), or on the initial command line. You can also set this option on the **Performance Tuning** tab 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 430

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.

For Microsoft Hyper-V virtual machines, this parameter specifies whether to perform a full image backup, an incremental-forever full backup, or an incremental-forever-incremental backup of Hyper-V 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



For backing up Microsoft Hyper-V 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.

Microsoft Hyper-V virtual machine parameters

full

Specifies that you want to perform an image backup of all objects on a Hyper-V virtual machine's disks. This is the default mode, and only mode that you can use for Hyper-V virtual machine backups if you do not have a license for Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V.

IFIncremental

Specifies that you want to perform an incremental-forever-incremental backup of a Hyper-V 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 Microsoft Hyper-V license. With that license installed, this is the default backup mode for Hyper-V backups.

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 Hyper-V virtual machines, see Incremental-forever backup strategy.

IFFull

Specifies that you want to perform an incremental-forever-full backup of a Hyper-V virtual machine. An incremental-forever-full backup backs up all used blocks on a virtual machine's disks. This mode requires a Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V license. By default, when the license for Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V is detected, the first backup of a Hyper-V 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.

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 a full VM backup of a Hyper-V VM named msvm1

```
dsmc backup vm msvm1 -mode=full  
-vmbackuptype=hypervfull
```

Task Perform an incremental-forever-full VM backup of a Hyper-V VM named msvm1

```
dsmc backup vm msvm1 -mode=iffull  
-vmbackuptype=hypervfull
```

Task Perform an incremental-forever-incremental backup of a Hyper-V VM named msvm1

```
dsmc backup vm msvm1 -mode=ifincremental  
-vmbackuptype=hypervfull
```

Task Perform the NAS image backup of the entire file system.

```
dsmc backup nas -mode=differential -nasnodename=nas1  
{/vol/vol0} {/vol/vol1}
```

Task Back up the c: drive using an image incremental backup that backs up only new and changed files after the last full image backup.

```
dsmc backup image c: -mode=full
```

Task Perform a full backup of all the files in filelist c:\dir1\filelist1 to the virtual file space name \virtfs containing the group leader c:\group1 file.

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

Related reference:

“Backup VM” on page 658

“Backup Group” on page 642

“Backup Image” on page 644

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 all Windows clients.

Syntax



Parameters

Yes

Specifies that you want to monitor the current NAS image backup or restore operation and display processing information on your screen. This is the default.

No Specifies that you do not want to monitor the current NAS image backup or restore operation.

Examples

Command line:

```
backup nas -mode=full -nasnodename=nas1 -monitor=yes  
{/vol/vol0} {/vol/vol1}
```

Myprimaryserver

The `myprimaryserver` option specifies the primary server name that the client uses to log on to the secondary server in failover mode.

During the normal (non-failover) logon process, the `myprimaryserver` option is sent to the client and is saved in the `dsm.opt` file. Do not edit this option during normal operations.

Important: If you change the value for the `myprimaryserver` option, authentication information such as the Tivoli Storage Manager password and encryption key will no longer work with the new primary server. You will be prompted for the password and encryption key for operations that require authentication. Therefore, do not change this value even if you change the secondary server connection information.

Supported Clients

This option is valid only for Windows clients.

Options File

This option is placed in the client options file (dsm.opt).

Syntax

►—MYPRIMARYServer—*primary_servername*—►

Parameters

primary_servername

Specifies the name of the primary server to be used for authentication during a failover. The primary server is the Tivoli Storage Manager server that a client uses for normal production.

Examples

Options file:

```
*** These options should not be changed manually
REPLSERVERNAME          TARGET
REPLTCPSERVERADDRESS    192.0.2.9
REPLTCPSPORT            1501
REPLSERVERGUID          60.4a.c3.e1.85.ba.11.e2.af.ce.00.0c.29.2f.07.d3

MYREPLICATIONServer    TARGET
MYPRIMARYSERVERNAME    SERVER1
*** end of automatically updated options
```

Command line:

Does not apply.

Related concepts:

“Automated client failover configuration and use” on page 60

Related tasks:

“Configuring the client for automated failover” on page 63

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 in the client options file (dsm.opt).

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 the secondary server in the dsm.opt file, and how to reference the secondary server. The connection information for the secondary server is located within the **REPLSERVERName** stanza. The **MYREPLICATIONServer** option points to the secondary server name that is specified by the **REPLSERVERName** stanza.

```
REPLSERVERNAME      TargetReplicationServer1
  REPLTCPSEVERADDRESS TargetReplicationServer1
  REPLTCPPOrt        1505
  REPLSSLPOrt        1506
  REPLSERVERGUID     91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00

COMMMethod          TCPip
TCPPOrt              1500
TCPSeveraddress      server_hostname1.example.com
PASSWORDAccess      prompt
MYREPLICATIONServer TargetReplicationServer1
MYPRIMARYSERVER      Server1
```

Related concepts:

“Automated client failover configuration and use” on page 60

Related tasks:

“Configuring the client for automated failover” on page 63

Namedpipename

The `namedpipename` option specifies the name of a named pipe to use for communications between a client and a server on the same Windows server domain.

Supported Clients

This option is valid for all Windows clients.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Communication** tab of the Preferences editor.

Syntax

►►—`NAMEDpipename`— *name*—————►►

Parameters

name

The name of a named pipe. The default is `\\.\pipe\Server1`.

Examples

Options file:

```
namedpipename \\.\pipe\dsmser1
```

Command line:

```
-namedpipename=\\.\pipe\dsmser1
```

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

Nasnodename

The `nasnodename` option specifies the node name for the NAS file server when processing NAS file systems. The 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 options file (`dsm.opt`).

You can override the default value in the `dsm.opt` file by entering a different value on the command line. If you do not specify the `nasnodename` option in the `dsm.opt` file, you must specify this option on the command line when processing NAS file systems.

You can use the `nasnodename` option with the following commands:

- **backup nas**
- **delete filespace**
- **query backup**
- **query filespace**
- **restore nas**

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. Place the `nasnodename` option in your client options file (`dsm.opt`). The value in the client options file is the default, but this value can be overridden on the command line. If the `nasnodename` option is not specified in the client options file, you must specify this option on the command line when processing NAS file systems.

Use the `class` option to specify the class of the file space to delete. To display a list of file spaces belonging to a NAS node so that you can choose one to delete, use the `-class=nas` option.

To delete NAS file spaces using the web client, see the topic for backing up your data.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **General** tab of the Preferences editor.

Syntax

►—NASNodename— —*nodename*—►

Parameters

nodename

Specifies the node name for the NAS file server.

Examples

Options file:

`nasnodename nas2`

Command line:

`-nasnodename=nas2`

Nodename

Use the `nodename` option in your client options file 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.

If you are working from a different workstation, you can use the `nodename` option even if the `passwordaccess` option is set to `generate`, however, the password is stored in the Windows Registry. To prevent this, use the `virtualnodename` option instead of `nodename`.

The node name is not necessarily the TCP/IP host name.

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.opt` file, or a `virtualnodename` entry in the client options file (`dsm.opt`), or a virtual node name specified on a command line, the default login ID is the name that the **hostname** command returns.
- If a `nodename` entry exists in the `dsm.opt` file, the `nodename` entry overrides the name that the **hostname** command returns.
- If a `virtualnodename` entry exists in the client options file (`dsm.opt`), or a virtual node name is specified on a command line, it cannot be the same name as the name returned by the **hostname** command. When the server accepts the virtual node name, a password is required (if authentication is on), even if the `passwordaccess` option is `generate`. When a connection to the server is established, access is permitted to any file that is backed up using this login ID.

Supported Clients

This option is valid for all clients.

Options File

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

Command line:

`-nodename=cougar`

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

“Virtualnodename” on page 576

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.
- Attribute changes to an object require a backup of the entire object.

For these reasons, you might want to use the `nojournal` option periodically to perform a traditional full incremental backup.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager API does not support this option.

Syntax

►►—NOJournal—►►

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc incr c: -nojournal
```

Related concepts:

“Journal-based backup” on page 687

Noprompt

The `noprompt` option suppresses the confirmation prompt that is presented by the **delete group**, **delete archive**, **expire**, **restore image**, and **set event** commands.

- **delete archive**
- **delete backup**
- **delete group**
- **expire**
- **restore image**

Supported Clients

This option is valid for all 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 c:\home\project\*
```

Nrtablepath

The `nrtablepath` 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 options file (`dsm.opt`).

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.

Restriction: The node replication table cannot be created in the `C:\` directory. If you choose to specify a location for the node replication table, do not specify the `C:\` directory.

Example

Options file:

nrtblepath C:\nrtbl

Command line:

Does not apply.

Related tasks:

“Determining the status of replicated client data” on page 65

“Configuring the client for automated failover” on page 63

Numberformat

The numberformat option specifies the format you want to use to display numbers.

Use this option if you want to change the default number format for the language of the message repository you are using.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time the client is called. Consult the documentation on your local system for details about setting up your locale definition.

Note: The numberformat option does not affect the web client. The web client uses the number format for the locale that the browser is running in. If the browser is not running in a supported locale, the web client uses the number format for US English.

You can use the numberformat option with the following commands:

- **delete archive**
- **delete backup**
- **expire**
- **query archive**
- **query asr**
- **query backup**
- **query image**
- **query nas**
- **query systemstate**
- **restore**
- **restore image**
- **restore nas**
- **restore registry**
- **restore systemstate**
- **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

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`.

Examples

Command line:

```
dsmc query session -optfile=myopts.opt
```

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

Password

The `password` option specifies a Tivoli Storage Manager password.

If you do not specify this option and your administrator has set authentication to `On`, you are prompted for a password when you start a 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 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 Windows systems:

Enclose the command parameters in quotation marks (").

Command line example:

```
dsmc set password "t67@#$$%^&" "pass2><w0rd"
```

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="secret>shhh"
```

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 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 client options file (dsm.opt). 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.

API applications must supply the password when a session is initiated. The application is responsible for obtaining the password.

generate

Encrypts and stores your password locally and generates a new password when the old password expires. The new password is randomly generated by the 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.

Examples

Options file:

```
passwordaccess generate
```

Command line:

Does not apply.

Related information

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 asr**
- **restore group**
- **restore image**
- **restore nas**
- **restore vm**
- **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 c:\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 asr**
- **query backup**
- **query group**
- **query image**
- **query nas**
- **query systemstate**

- **query vm** (vmbackuptype=fullvm and vmbackuptype=hypervfull)
- **restore**
- **restore group**
- **restore image**
- **restore nas**
- **restore systemstate**
- **restore vm** (vmbackuptype=fullvm and vmbackuptype=hypervfull)

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 -pitdate=08/01/2003 c:\myfiles\
```

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 asr**
- **query backup**
- **query image**
- **query nas**
- **query systemstate**
- **query vm**(vmbackuptype=fullvm and vmbackuptype=hypervfull)
- **restore**
- **restore image**
- **restore nas**
- **restore systemstate**

- **restore vm** (vmbackuptype=fullvm and vmbackuptype=hypervfull)

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 -pitt=08/01/2003 c:\myfiles\
```

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.

- ## Supported Clients

Options File

Syntax



If your administrator uses a blank or null string on the `postschedulecmd` option, you cannot run a post-schedule command.

Examples

```
posts startdb.cmd
posts 'rename c:\myapp\logfile.log logfile.new'
posts 'net start "simple service"'
posts 'rename "c:\myapp\log file.log" "log file.new"'
posts '"C:\Program Files\MyTools\runreport.bat"
    log1.in log2.in'
```

Related reference:

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.

This option can be used in conjunction with the `presnapshotcmd` option to allow you to quiesce an application while a snapshot is created, and then to restart that application after the snapshot is started. This option is only valid if OFS or online image backup has been configured.

For an online image backup, use this option with the **backup image** command, the `include.image` option, or in the `dsm.opt` file.

For open file support operations, use the `postsnapshotcmd` option in an `include.fs` statement or in your client options file (`dsm.opt`).

If the `postsnapshotcmd` fails the operation continues, but appropriate warnings are logged.

Attention: During image backup operations or snapshot differential backup operations, if the command that you include on either the `presnapshotcmd` or `postsnapshotcmd` statement starts an asynchronous process, the command might not complete before the backup operation finishes. If the command does not complete before the backup completes, temporary files might be locked, which prevents them from being deleted. A database event occurs and the following message is recorded in the `dsmerror.log` file:

```
ANS0361I DIAG: ..\..\common\db\cacheobj.cpp( 777): dbDelete():  
remove('C:\adsm.sys\SystemExcludeCache_24400820.TsmCacheDB'):  
errno 13: "Permission denied".
```

The file that is specified in the message (`cacheobj.cpp`) can be manually deleted after the command that was started by the `presnapshotcmd` or `postsnapshotcmd` option completes.

Supported Clients

This option is valid for all Windows 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 options file (`dsm.opt`). You can also set this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

►►—POSTSNAPshotcmd— —"*cmdstring*"—————◄◄

Parameters

"cmdstring"

Specifies the quiesce command to process.

Use a blank, or null, string for *"cmdstring"* if you want to prevent any commands from running that the administrator uses for `postsnapshotcmd`. If you specify a blank or null string, it prevents the administrator from using a command on this option. If your administrator uses a blank or null string on the `postsnapshotcmd` option, you cannot run a post-snapshot command.

Use the `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 "restart application"
```

The command string is a valid command for restarting your application.

Command line:

```
backup image -postsnapshotcmd="restart application"
```

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

Related information

“Include options” on page 430

“Srvprepostscheddisabled” on page 540

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 `prenschedulecmd` command is not tracked, and does not influence the return code of the scheduled event.
2. The server can also define the `preschedulecmd` option (and the `prenschedulecmd` option).


Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Scheduler** tab, in the **Schedule Command** dialog box in the Preferences editor.

Syntax



Parameters

cmdstring

Specifies the command to process. Use only one `preschedulecmd` option. You can enter a command to be executed before a schedule using this option.

Specify the command string just as you would enter it from the operating system command prompt; if the string you specify would require quotation marks to run it at a Windows prompt, include the quotation marks as needed. If the command string contains any blank spaces, enclose the command string in single quotation marks.

In this example, single quotation marks are needed because the command string contains space characters:

```
'net stop someservice'
```

In this next example, double quotation marks are needed because both the file being renamed, and the new file name, contain space characters. Because the command string does contain space characters, the entire string must be enclosed in single quotation marks.

```
presc 'rename "c:\myapp\log file.log" "log file.old"'
```

Use a blank or null string for *cmdstring* if you want to prevent any commands from running that the Tivoli Storage Manager server administrator uses for `postschedulecmd` and `preschedulecmd`. If you specify a blank or null string on either option, it prevents the administrator from using a command on both options.

If your administrator uses a blank or null string on the `preschedulecmd` option, you cannot run a pre-schedule command.

Examples

Options file:

```
presc stopdb.cmd
presc 'rename c:\myapp\logfile.log logfile.old'
presc 'net stop "simple service"'
presc 'rename "c:\myapp\log file.log" "log file.old"'
presc '"C:\Program Files\MyTools\runreport.bat"
      log1.in log2.in'
```

Command line:

```
-preschedulecmd='"quiesce database"'
```

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

Related concepts:

Chapter 8, “Client return codes,” on page 265

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 can affect applications that rely on accurate last-access dates such as a Storage Resource Management (SRM) application.
3. The last access date cannot be preserved on files that are write-protected either by the read-only attribute or by a restrictive NTFS security permission.
4. You cannot reset the last access date of read-only files. The `preserveLastAccessDate` option ignores read-only files and does not change their date.

Supported Clients


This option is valid for all Windows clients.

The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Backup** tab of the Preferences editor.

Syntax



Parameters

No A backup or archive operation can change the last access date. This value is the default.

A backup or archive operation does not change the last access date.

Options file:

Command line:

Preservepath

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.

- restore
- restore backupset
- restore group
- retrieve

Syntax



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.

Restores the entire path, starting from the root, into the specified directory. The entire path includes all the directories except the file space name.

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 the server file space contains the following backup copies:

```
c:\h1\m1\file.a
c:\h1\m1\file.b
c:\h1\m1\l1\file.x
c:\h1\m1\l1\file.y
```

This command:

```
dsmc res backupset my.backupset.file /fs/h1/m1/ /u/ann/ -su=yes
creates a local backupset file named "my.backupset.file".
```

Restores these directories and files:

```
c:\ann\h1\m1\file.a
c:\ann\h1\m1\file.b
```

This command:

```
dsmc res c:\h1\m1\ c:\ann\ -preser=nobase.
```

Restores these directories and files:

```
c:\ann\file.a
c:\ann\file.b
```

This command:

```
dsmc res c:\h1\m1\ c:\ann\ -preser=subtree.
```

Restores these directories and files:

```
c:\ann\m1\file.a
c:\ann\m1\file.b
```

This command:

```
dsmc res c:\h1\m1\ c:\ann\ -preser=none.
```

Restores these directories and files:

```
c:\ann\file.a
c:\ann\file.b
```

This command:

```
dsmc res c:\h1\m1\ c:\ann\ -su=yes -preser=
complete
```

Restores these directories and files:

```
c:\ann\h1\m1\file.a
c:\ann\h1\m1\file.b
c:\ann\h1\m1\l1\file.x
c:\ann\h1\m1\l1\file.y
```

This command:

```
dsmc res c:\h1\m1\ c:\ann\ -su=yes -preser=nobase.
```

Restores these directories and files:

```
c:\ann\file.a
c:\ann\file.b
c:\ann\l1\file.x
c:\ann\l1\file.y
```

This command:

```
dsmc res c:\hl\ml\ c:\ann\ -su=yes -preser=subtree.
```

Restores these directories and files:

```
c:\ann\ml\file.a  
c:\ann\ml\file.b  
c:\ann\ml\l1\file.x  
c:\ann\ml\l1\file.y
```

This command:

```
dsmc res c:\hl\ml\ c:\ann\ -su=yes -preser=none.
```

Restores these directories and files:

```
c:\ann\file.a  
c:\ann\file.b  
c:\ann\file.x  
c:\ann\file.y
```

This command:

```
dsmc res backupset c:\hl\ml\ c:\ann\ -su=yes  
-preser=nobase -loc=file
```

Restores these directories and files:

```
c:\ann\file.a  
c:\ann\file.b  
c:\ann\file.x  
c:\ann\file.y
```

Presnapshotcmd

The `presnapshotcmd` option allows you to run 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.

This option can be used in conjunction with the `postsnapshotcmd` option to allow you to quiesce an application while a snapshot is created, and then to restart that application after the snapshot is started. This option is only valid if OFS or online image backup has been configured.

For an online image backup, use this option with the **backup image** command, the `include.image` option, or in the `dsm.opt` file.

For open file support operations, use the `presnapshotcmd` option in an `include.fs` statement or in your client options file (`dsm.opt`).

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.

Attention: During image backup operations or snapshot differential backup operations, if the command that you include on either the `presnapshotcmd` or `postsnapshotcmd` statement starts an asynchronous process, the command might not complete before the backup operation finishes. If the command does not complete before the backup completes, temporary files might be locked, which prevents them from being deleted. A database event occurs and the following message is recorded in the `dsmerror.log` file:

```
ANS0361I DIAG: ..\..\common\db\cacheobj.cpp( 777): dbDelete():  
  remove('C:\adsm.sys\SystemExcludeCache__24400820.TsmCacheDB'):  
    errno 13: "Permission denied".
```

The file that is specified in the message (`cacheobj.cpp`) can be manually deleted after the command that was started by the `presnapshotcmd` or `postsnapshotcmd` option completes.

Supported Clients

This option is valid for all Windows 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 options file (`dsm.opt`). You can set also this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

►►—PRESNAPSHOTcmd— —"*cmdstring*"——►►

Parameters

"cmdstring"

Specifies the quiesce command to process.

Use a blank, or null, string for *"cmdstring"* if you want to prevent any commands from running that the administrator uses for `presnapshotcmd`. If you specify a blank or null string, it prevents the administrator from using a command on this option. If your administrator uses a blank or null string on the `presnapshotcmd` option, you cannot run a pre-snapshot command.

Use the `srvprepostsnapdisabled` option to prevent the Tivoli Storage Manager server administrator from 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 "<insert your application quiesce command here>  
application"
```

The command string is a valid command for quiescing your application.

Command line:

```
backup image -presnapshotcmd="<insert your application quiesce command  
here> application"
```

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

Related information

"Include options" on page 430

"Srvprepostscheddisabled" on page 540

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.

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 options file (`dsm.opt`). 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

Single objects that span multiple volumes only include one volume in the total number of volumes statistics. For example, if c:\bigfile spans two volumes, only one of the volumes is counted in the estimated number of volumes.

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 query backup k:\.* -subdir=yes -QUERYSUMMARY

IBM Tivoli Storage Manager

Command Line Backup/Archive Client Interface

Client Version 6, Release 1, Level 0.0 0324FA

Client date/time: 04/03/2008 05:46:42

(c) Copyright by IBM Corporation and other(s) 1990, 2008. All Rights Reserved.

Node Name: BARKENSTEIN

Session established with server BARKENSTEIN_SERVER1: Windows

Server Version 5, Release 4, Level 1.0

Server date/time: 04/03/2008 05:46:43 Last access: 04/03/2008 05:46:09

Size		Backup Date	Mgmt Class	A/I File
----		-----	-----	----
0	B	04/02/2008 20:21:51	STANDARD	A \\barkenstein\k\$\
0	B	04/02/2008 20:21:51	STANDARD	A \\barkenstein\k\$\jack_test
0	B	04/01/2008 12:37:07	STANDARD	A \\barkenstein\k\$\
				System Volume Information
0	B	04/01/2008 12:37:07	STANDARD	A \\barkenstein\k\$\Test1
0	B	04/02/2008 20:21:51	STANDARD	A \\barkenstein\k\$\TestTree
0	B	04/01/2008 12:37:07	STANDARD	A \\barkenstein\k\$\Tree150
0	B	04/02/2008 19:49:20	STANDARD	A \\barkenstein\k\$\Tree150.1
0	B	04/01/2008 12:37:07	STANDARD	A \\barkenstein\k\$\Tree150.2
0	B	04/02/2008 19:50:51	STANDARD	A \\barkenstein\k\$\Tree150.3
0	B	04/01/2008 12:37:07	STANDARD	A \\barkenstein\k\$\Tree1500
0	B	04/02/2008 10:41:40	STANDARD	A \\barkenstein\k\$\Tree150_2
0	B	04/02/2008 20:02:31	STANDARD	A \\barkenstein\k\$\tree18
0	B	04/02/2008 20:15:04	STANDARD	A \\barkenstein\k\$\Tree18.test
0	B	04/01/2008 12:37:07	STANDARD	A \\barkenstein\k\$\Tree30
0	B	04/01/2008 12:37:07	STANDARD	A \\barkenstein\k\$\Tree30.2
0	B	04/02/2008 19:52:30	STANDARD	A \\barkenstein\k\$\tree30.test
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file1
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file10
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file11
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file12
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file13
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file14
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file15
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file16
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file17
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file18
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file19
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file2
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file20
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file21
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file3
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file4
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file5
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file6
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file7
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file8
11,788	B	04/02/2008 19:55:32	DEFAULT	A \\barkenstein\k\$\file9
11,788	B	04/02/2008 13:31:06	DEFAULT	A \\barkenstein\k\$\file910
10,964	B	04/01/2008 12:37:07	DEFAULT	A \\barkenstein\k\$\filea
10,964	B	04/01/2008 12:37:07	DEFAULT	A \\barkenstein\k\$\fileb
10,964	B	04/01/2008 12:37:07	DEFAULT	A \\barkenstein\k\$\x

Summary Statistics

Total Files	Total Dirs	Avg. File Size	Total Data	Memory Est.
-----	-----	-----	-----	-----
25	16	11.41 KB	285.37 KB	10.58 KB

Estimated Number of Volumes: 2

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 options file (dsm.opt). You can set this option on the **Command Line** tab, **Do not display process information on screen** checkbox of the Preferences editor.

Syntax

►►—QUIET—◄◄

Parameters

There are no parameters for this option.

Examples

Options file:
quiet

Command line:
-quiet

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

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=c:\important_files.txt`, where `important_files.txt` contains the list of files to process.

`important_files.txt` contains the following list of files:

```
c:\gpfs\home\myfiles\"file"1000
c:\gpfs\home\myfiles\'file'
c:\gpfs\home\myfiles\file'ABC
c:\gpfs\home\myfiles\ABC"file"
```

Related information

For information about the `filelist` option, see "Filelist" on page 414.

For information about syntax for file specifications, see "Specifying input strings that contain blank spaces or quotation marks" on page 114.

"Wildcardsareliteral" on page 623

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 options file (`dsm.opt`). You can set this option on the **Restore** tab, **Action for files that already exist** section of the Preferences editor.

Syntax



Parameters

Prompt

For nonscheduled operations, you specify whether to overwrite existing files. For scheduled operations, existing files are not overwritten and no prompts are displayed. This is the default.

All

All existing files are overwritten, including read-only files. All locked files are replaced when the system is rebooted. If access to a file is denied, you are prompted to skip or overwrite the file. No action is taken on the file until there is a response to the prompt.

Yes

Existing files are overwritten, *except* read-only files. For nonscheduled operations, you specify whether to overwrite existing read-only files. For scheduled operations, existing read-only files are not overwritten and no prompts are displayed. If access to a file is denied, the file is skipped.

No Existing files are not overwritten. No prompts are displayed.

Note: You can choose to replace locked files when the system is rebooted. The client cannot perform an in-place restore of active files. However, it stages restored versions of active files for replacement during the next reboot, except for files containing named streams, sparse files, and directories. You can only restore these files if they are unlocked.

Examples

Options file:

```
replace all
```

Command line:

```
-replace=no
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This 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 client options file (`dsm.opt`).

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 the secondary server in the `dsm.opt` file, and how to reference the secondary server. The connection information for the secondary server is located within the **REPLSERVERName** stanza. The **MYREPLICATIONServer** option points to the secondary server name that is specified by the **REPLSERVERName** stanza.

```
REPLSERVERNAME      TargetReplicationServer1
  REPLTCPSEVERADDRESS TargetReplicationServer1
  REPLTCPPOrt        1505
  REPLSSLPOrt        1506
  REPLSERVERGUID     91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00

COMMMethod          TCPip
TCPPOrt              1500
TCPSeveraddress      server_hostname1.example.com
PASSWORDAccess       prompt
MYREPLICATIONServer TargetReplicationServer1
MYPRIMARYSERVER      Server1
```

Related concepts:

“Automated client failover configuration and use” on page 60

Related tasks:

“Configuring the client for automated failover” on page 63

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 options file (dsm.opt).

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 the secondary server in the dsm.opt file, and how to reference the secondary server. The connection information for the secondary server is located within the **REPLSERVERName** stanza. The **MYREPLICATIONServer** option points to the secondary server name that is specified by the **REPLSERVERName** stanza.

```
REPLSERVERNAME    TargetReplicationServer1
  REPLTCPSEVERADDRESS TargetReplicationServer1
  REPLTCPPEORT      1505
  REPLSSLPEORT      1506
  REPLSERVERGUID    91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00.00

COMMMethod        TCPip
TCPPEORT           1500
TCPSEVERaddress    server_hostname1.example.com
PASSWORDAccess     prompt
MYREPLICATIONServer TargetReplicationServer1
MYPRIMARYSERVER    Server1
```

Related concepts:

“Automated client failover configuration and use” on page 60

Related tasks:

“Configuring the client for automated failover” on page 63

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.

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 client options file (`dsm.opt`).

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 the secondary server in the `dsm.opt` file, and how to reference the secondary server. The connection information for the secondary server is located within the **REPLSERVERName** stanza. The **MYREPLICATIONServer** option points to the secondary server name that is specified by the **REPLSERVERName** 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

COMMMethod          TCPip
TCPPOINT            1500
TCPSEVERADDRESS      server_hostname1.example.com
PASSWORDAccess      prompt
MYREPLICATIONServer TargetReplicationServer1
MYPRIMARYSERVER      Server1
```

Related concepts:

“Automated client failover configuration and use” on page 60

Related tasks:

“Configuring the client for automated failover” on page 63

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 client options file (`dsm.opt`).

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 the secondary server in the `dsm.opt` file, and how to reference the secondary server. The connection information for the secondary server is located within the **REPLSERVERName** stanza. The **MYREPLICATIONServer** option points to the secondary server name that is specified by the **REPLSERVERName** stanza.

```
REPLSERVERNAME      TargetReplicationServer1
  REPLTCPSEVERADDRESS TargetReplicationServer1
  REPLTCPPOST        1505
  REPLSSLPORT        1506
  REPLSERVERGUID      91.0f.ef.90.5c.cc.11.e1.ae.34.08.00.00.00.00

COMMMethod          TCPip
TCPPost             1500
TCPSeveraddress      server_hostname1.example.com
PASSWORDAccess      prompt
MYREPLICATIONServer TargetReplicationServer1
MYPRIMARYSERVER      Server1
```

Related concepts:

“Automated client failover configuration and use” on page 60

Related tasks:

“Configuring the client for automated failover” on page 63

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

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

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 commethod 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 the secondary server in the dsm.opt file, and how to reference the secondary server. The connection information for the secondary server is located within the **REPLSERVERName** stanza. The **MYREPLICATIONServer** option points to the secondary server name that is specified by the **REPLSERVERName** 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

  COMMMethod         TCPip
  TCPPOrt             1500
  TCPServeraddress    server_hostname1.example.com
  PASSWORDAccess      prompt
  MYREPLICATIONServer TargetReplicationServer1
  MYPRIMARYSERVER     Server1
```

Related concepts:

“Automated client failover configuration and use” on page 60

Related tasks:

Resetarchiveattribute

Use the `resetarchiveattribute` option to specify whether the backup-archive client resets the Windows archive attribute on files that are successfully backed up to a Tivoli Storage Manager server.

The client also resets the archive attribute during incremental backups if it is determined that there is already an active object on the server. The `resetarchiveattribute` option is useful in conjunction with applications, such as the IBM Tivoli Storage Productivity Center, as a simple way to report on the backup status of files.

The Windows archive attribute is used to indicate that a file has changed since the last backup. After the client resets the archive attribute, the Windows operating system turns the attribute back to ON after the file has been modified. The client does not use the Windows archive attribute to determine if a file is a candidate for incremental backup, but only manipulates this attribute for reporting purposes. The client uses a much more sophisticated method to determine candidacy for incremental backup.

There are several applications which manipulate or examine the Windows archive attribute. Be aware of the ramifications of using the `resetarchiveattribute` option in conjunction with these products.

If you set the `resetarchiveattribute` option to yes, after a file has been successfully backed up to the Tivoli Storage Manager server, the client resets the Windows archive attribute on the local file system:

- The Windows archive attribute is reset during incremental and selective backups after the file has been successfully committed to the Tivoli Storage Manager server database. This attribute is not reset for archive, or image operations.
- The Windows archive attribute is not reset when processing system objects or system state objects.
- The Windows archive attribute is not reset for directory entries.

In addition, in order for the local file system to reflect the current active object inventory on the Tivoli Storage Manager server, the `resetarchiveattribute` option instructs the client to reset the Windows archive attribute on the local file system if it is determined during incremental backup that a valid, active backup copy of the file already exists on the server. This behavior is not displayed in the following cases:

- Incremental backup operations which do not examine the stored client attributes on the server, such as journal-based backup or incremental-by-date processing.
- Files that are not examined during an incremental backup operation because they are excluded from backup processing.

The client does not guarantee the accuracy of the current setting of the Windows archive attribute. For example, if the `resetarchiveattribute` option is set to yes and a file examined by a reporting product indicates that the Windows archive attribute is OFF for a particular file, this does not necessarily mean that a valid, active backup copy of the file exists on the Tivoli Storage Manager server. Factors that could contribute to this type of situation include:

- An independent software vendor product is manipulating the Windows archive attribute

- A file space was deleted from the server
- A backup tape was lost or destroyed

There should be no significant performance degradation when using the `resetarchiveattribute` option. The `resetarchiveattribute` option does not affect restore processing.


Supported Clients

This option is valid for all Windows clients. The server can also define this option.

Options File

This option is valid in the client options file (dsm.opt) or server client options set. You can set this option on the **Backup** tab of the Preferences editor.

Syntax



Parameters

Yes

Specifies that you want to reset the Windows archive attribute for files during a backup operation.

No Specifies that you do not want to reset the Windows archive attribute for files during a backup operation. This is the default.

Examples

Options file:

```
resetarchiveattribute yes
```

Related information

“Full and partial incremental backup” on page 136

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 client options file (dsm.opt). 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 346

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 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 options file (*dsm.opt*). 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:

retryp 10

Command line:

-retryperiod=10

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

Revokeremoteaccess

The `revokeremoteaccess` option restricts an administrator with client access privilege from accessing a client workstation that is running the web client.

This option does not restrict administrators with client owner, system, or policy privilege from accessing your workstation through the web client.


Supported Clients

This option is valid for all clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Web Client** tab of the Preferences editor.

Syntax



Parameters

None

Does not revoke access to administrators who have client access authority for the client. This is the default.

Access

Revokes access to administrators who have client access authority for the client.

Examples

Options file:

```
revokeremoteaccess none
```

Command line:

Does not apply.

Runasservice

The `runasservice` option forces the client command process to continue running, even if the account that started the client logs off.

Use this option with the **AT** command and the **dsmc sched** command when you schedule client command batch jobs. The `runasservice` option is *not* valid in any options file (`dsm.opt` or `tmasr.opt`).

Important: Use the scheduler service when running Tivoli Storage Manager services unattended. Set `runasservice=yes` only to schedule client commands using the Windows **AT** command. Setting `runasservice=yes` might interfere with other interactive uses of the backup-arcserve client.

Options File

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

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 717

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 client options file (dsm.opt).

Syntax

►—SCHEDCMException—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  
  
schedcmexception "start dir c: /s"  
  
schedcmexception "start echo hello, world!"
```

Related information

“Schedcmddisabled” on page 513

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 `dsmweblog.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 options file (`dsm.opt`).

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 options file (`dsm.opt`). 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.

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`.

Examples

Options file:

```
schedlogname c:\mydir\schedlog.jan
```

Command line:

```
-schedlogn=c:\mydir\schedlog.jan
```

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

Related information

See “Errorlogname” on page 397 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 (dsmsched.pru and dsmwebcl.pru), and the existing logs (dsmsched.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 dsmsched.log file is copied to a file called dsmsched.pru, the content of dsmwebcl.log is copied to dsmwebcl.pru, and new log entries are created in both dsmsched.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 options file (dsm.opt).

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 `dsmsched.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.opt 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.opt file or on the schedule command; the client can then be contacted at either an address or a port of your choice (useful for client systems with multiple network interface cards).

Note:

1. When changing the setting of this option in the client options file (dsm.opt) you must stop and restart the scheduler service for the setting to take effect.
2. The server can also define this option.


Supported Clients

This option is valid for all clients.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Scheduler** tab, in the **Schedule Mode** section in the Preferences editor.

Syntax



Parameters

P0lling

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 querieschedperiod option.

PRompted

The client scheduler waits for the server to contact your client node when scheduled work needs to be done.

Note:

1. Use schedmode prompted in conjunction with the autodeploy option, to enable the scheduler to process the client deployment schedule immediately.
2. If you use the **dsmsc 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:

"Autodeploy" on page 332

"Cadlistenonport" on page 338

"Tcpclientaddress" on page 557

"Tcpclientport" on page 557

Schedrestretrdisabled

The schedrestretrdisabled option specifies whether to disable the execution of restore or retrieve schedule operations.


Supported Clients

This option is valid for all clients. The server cannot define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client options file (dsm.opt) for the scheduler. You can set this option on the **Scheduler** tab 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.

Place this option in the client 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.

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 attempt to initiate sessions.

Attention: You cannot use the **dsmcad** for scheduling when you set the `sessioninitiation` option to `serveronly`

Note: If you set the `sessioninitiation` option to `serveronly`, the client setup wizard and scheduler service are unable to authenticate to the Tivoli Storage Manager server. In this case, you can execute scheduler from the command line (`dsmc schedule`) and enter the password for your node when prompted or use the following **dsmcutil** command to write the password into the Windows Registry:

```
dsmcutil updatepw /node:nnn /password:ppp /validate:no
```

To avoid this problem, when configuring the client scheduler using the setup wizard, ensure that the **Contact the TSM Server to validate password** check box on the TSM Authentication page is unchecked.

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 `HAddress` 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 `LAddress` 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 client options file (dsm.opt). You can set this option on the **Scheduler** tab, **Session Initiation** field of the Preferences editor.

Syntax



Parameters

Client

Specifies that the client initiates sessions with the server by communicating on the TCP/IP port defined with the server option TCP_{PORT}. 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 36

"Tcpclientport" on page 557

Shmport

The `shmport` option specifies the TCP/IP port address of a server when using shared memory. All shared memory communications start with a TCP/IP connection.

Note: The value specified for the `shmport` option in the client options file (`dsm.opt`) must match the value specified for `shmport` in the server options file.

Supported Clients

This option is valid for all Windows clients.

Options File

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

Syntax

►►—SHMPort— *port_number*—————►►

Parameters

port_number

Specifies the port number. You can specify a value from 1 to 32767. The default value is 1510.

Examples

Options file:

shmport 1580

Command line:

Does not apply.

Showmembers

Use the showmembers option to display all members of a group.

You can use the showmembers option with the **query group**, **query systemstate**, 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 Windows clients.

Syntax

►►—SHOWMembers—————►►

Parameters

There are no parameters for this option.

Examples

Command line:

restore group {virtfs}* -pick -showmembers

Skipmissingsyswfiles

Use the Skipmissingsyswfiles option to specify whether the Tivoli Storage Manager client skips certain missing VSS writer files and continues the system state backup.

Setting the `skipmissingsyswfile` option to `yes` causes certain VSS writer files that are not found during a system state backup to be skipped. This option is effective only for missing files from the following VSS writers:

- System Writer
- Windows Deployment Service Writer
- Event Log writer

Consider the following items before you use the `skipmissingsyswfile` option:

- Setting the `skipmissingsyswfile` option to `yes` enables backups that might have failed to complete with previous versions of the backup-archive client.
- There is a small risk of an inconsistent backup because a file is skipped.
- This risk is minimized by these factors:
 - The backup can be done only when the system is running.
 - Critical system files are protected from deletion by Microsoft Windows.

Supported Clients

This option is valid for Windows clients.

Options File

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

Syntax



Parameters

Yes

Specifies that you want the backup-archive client to skip certain files that are not found during system state backup. The files that are not found are logged to both the error log and the server activity log. The final return code is set to 8. This is the default.

No Specifies that you want the backup-archive client to stop the backup when files are not found during system state backup. The files that are not found are logged to the error log and to the server activity log. The final return code is 12.

Examples

Options file:

```
SKIPMISSingsyswfiles yes
```

Command line:

```
-SKIPMISSingsyswfiles=yes
```

Related reference:

“Backup Systemstate” on page 650

Skipntpermissions

The skipntpermissions option bypasses processing of Windows file system security information.

You can use this option for incremental backups, selective backups, restore operations, and for archive and retrieve operations.

Supported Clients

This option is valid for all Windows clients.

Options File

Place this option in the client options file (dsm.opt). It applies to **incremental**, **selective**, **restore**, **archive**, and **retrieve** commands. You can also set this option on the **General** tab of the Preferences editor.

Syntax



Parameters

No If you specify *No*, Windows file system security information is backed up, restored, archived, or retrieved. This is the default setting.

Yes

If you specify *Yes*, Windows file system security information is not backed up, restored, archived, or retrieved.

Examples

Options file:

skipntp yes

Command line:

-skipntp=yes

Skipntsecuritycrc

The skipntsecuritycrc option controls the computation of the security cyclic redundancy check (CRC) for a comparison of Windows NTFS or ReFS security information during an incremental or selective backup, archive, restore, or retrieve operation.

If you set the skipntsecuritycrc option to no (the default), performance might be slower because the program must retrieve all the security descriptors.

Use this option with the following commands:

- **archive**
- **incremental**
- **restore**
- **retrieve**
- **selective**

Supported Clients

This option is valid for all Windows clients.

Options File

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

Syntax



Parameters

No If you specify *No*, the security CRC is generated during a backup. This is the default setting.

Yes

If you specify *Yes*, the security CRC is not generated during a backup. All the permissions are backed up, but the program cannot determine if the permissions are changed during the next incremental backup. When the `skipntpermissions` option is set to *yes*, the `skipntsecuritycrc` option does not apply.

Examples

Options file:

```
skipnts no
```

Command line:

```
-skipnts=no
```

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.

The `snapdiff` (snapshot difference) option is for backing up NAS/N-Series file server volumes that are NFS or CIFS attached.

Restriction: None of the NetApp predefined shares, including C\$, works with the Tivoli Storage Manager snapshot difference option because the backup-archive client cannot determine their mount points programmatically.

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 83.

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.

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 *snappdiff* option for any file system that is not mounted to the root of the volume. After you backed up the data with the *snappdiff* option, the snapshot that was used as the base snapshot is deleted from the snapshot directory.

On Windows 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 `snappdiffhttps` 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 `snappdiffhttps` 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 `snappdiff` 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 `snappdiff` option, an error message is generated.

No effect

The option can be used, but it is ignored.

Table 58. Incremental command: Related options

Option	Where specified	With <code>snappdiff</code>
<code>asnodename</code> “Asnodename” on page 324	Client options file (<code>dsm.opt</code>) or command line.	Valid
<code>autofsrename</code> “Autofsrename” on page 333	Client options file (<code>dsm.opt</code>) only.	No effect
<code>basesnapshotname</code> “Basesnapshotname” on page 337	Client options file (<code>dsm.opt</code>) or command line.	Valid
<code>changingretries</code> “Changingretries” on page 340	Client options file (<code>dsm.opt</code>) or command line.	No effect
<code>compressalways</code> “Compressalways” on page 350	Client options file (<code>dsm.opt</code>) or command line.	Valid
<code>compression</code> “Compression” on page 351	Client options file (<code>dsm.opt</code>) or command line.	Valid
<code>createnewbase</code> “Createnewbase” on page 354	Command line only.	Valid
<code>diffsnapshot</code> “Diffsnapshot” on page 364	Command line only.	Valid
<code>diffsnapshotname</code> “Diffsnapshotname” on page 366	Client options file (<code>dsm.opt</code>) or command line.	Valid
<code>dirsonly</code> “Dirsonly” on page 368	Command line only.	Valid
<code>domain</code> “Domain” on page 371	Client options file (<code>dsm.opt</code>) or command line only.	Valid
<code>enablelanfree</code> “Enablelanfree” on page 392	Client options file (<code>dsm.opt</code>) or command line.	Valid

Table 58. Incremental command: Related options (continued)

Option	Where specified	With snapdiff
encryptiontype "Encryptiontype" on page 393	Client options file (dsm.opt).	Valid
encryptkey "Encryptkey" on page 394	Client options file (dsm.opt).	Valid
exclude.fs.nas "Exclude options" on page 400	Client options file (dsm.opt).	No effect
filelist "Filelist" on page 414	Command line only.	Not valid
filesonly "Filesonly" on page 418	Command line only.	Valid
include.fs.nas "Include options" on page 430	Client options file (dsm.opt) or command line.	No effect
incl excl "Incl excl" on page 429	Client options file (dsm.opt).	Valid, but only when a file change is detected by NetApp.
incrbydate "Incrbydate" on page 445	Command line only.	Not valid
memoryefficientbackup "Memoryefficientbackup" on page 462	Client options file (dsm.opt), server, or command line.	No effect
monitor "Monitor" on page 468	Command line only.	Not valid
nojournal "Nojournal" on page 474	Command line only.	Not valid
postsnapshotcmd "Postsnapshotcmd" on page 486	Client options file (dsm.opt) or with the include.fs option.	Valid
preservelastaccessdate "Preservelastaccessdate" on page 489	Client options file (dsm.opt) or command line.	Valid
presnapshotcmd "Presnapshotcmd" on page 492	Client options file (dsm.opt) or with the include.fs option.	Valid
resetarchiveattribute "Resetarchiveattribute" on page 507	Client options file (dsm.opt).	Valid
skiptpermissions "Skiptpermissions" on page 528	Client options file (dsm.opt) or command line.	Valid
skiptsecuritycrc "Skiptsecuritycrc" on page 528	Client options file (dsm.opt) or command line.	Valid
snapdiffhttps "Snapdiffhttps" on page 534	Command line only.	Valid
snapshotproviderfs "Snapshotproviderfs" on page 535	Client options file (dsm.opt) or with the include.fs option.	Not valid
snapshotproviderimage "Snapshotproviderimage" on page 536	Client options file (dsm.opt) or with the include.image option.	Not valid
snapshotroot "Snapshotroot" on page 537	Command line only.	Not valid
subdir "Subdir" on page 548	Client options file (dsm.opt) or command line.	Not valid
subfilebackup "Subfilebackup" on page 549	Client options file (dsm.opt) or command line.	Not valid
tapeprompt "Tapeprompt" on page 553	Client options file (dsm.opt) or command line.	Valid

Table 58. Incremental command: Related options (continued)

Option	Where specified	With snapdiff
toc "Toc" on page 563	Command line only.	Not valid
useexistingbase "Useexistingbase" on page 569	Command line only.	Valid
virtualfsname "Virtualfsname" on page 575	Command line only.	Not valid

Supported Clients

This option is valid for all Windows clients except IA 64 bit.

Syntax

►►—SNAPDiff—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

Perform a snapshot-differential-incremental backup from a snapshot that is taken of a network share //homestore.example.com/vol/vol1 mounted on drive H:, where homestore.example.com is a file server.

```
incremental -snapdiff H:
```

Perform a snapshot-differential-incremental backup from a snapshot that is taken of a network share //homestore.example.com/vol/vol1 mounted on drive H:, where homestore.example.com is a file server. The -diffsnapshot option value of LATEST means that the operation uses the latest snapshot (the active snapshot) for volume H:.

```
incremental -snapdiff H: -diffsnapshot=latest
```

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 h:
```

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 h:
```

Perform a full incremental backup because you made some include or exclude changes:

```
dsmc incremental -snapdiff -createnewbase=yes h:
```

Related concepts:

"Snapshot differential backup with an HTTPS connection" on page 142

“SnapMirror support for NetApp snapshot-assisted progressive incremental backup (snapdiff)” on page 88

Related tasks:

“Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 83

Related reference:

“Snapdiffhttps”

“Basesnapshotname” on page 337

“Diffsnapshotname” on page 366

“Useexistingbase” on page 569

“Diffsnapshot” on page 364

“Set Password” on page 781

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 all Windows 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 a Windows system with a network share \\netappl\vol1, where netappl is a filer.

```
dsmc incr \\netappl\vol1 -snapdiff -snapdiffhttps
```

Command line:

Issue the following command on a Windows system with a network share \\netappl.example.com\petevol mounted on drive v:, where netappl.example.com is a filer.

```
dsmc incr v: -snapdiff -snapdiffhttps
```

```
IBM Tivoli Storage Manager
Command Line Backup-Archive Client Interface
  Client Version 7, Release 1, Level 0.0
  Client date/time: 10/23/2012 11:36:05
(c) Copyright by IBM Corporation and other(s) 1990, 2013. All Rights Reserved.

Node Name: THINKCENTRE
Session established with server BARKENSTEIN_SERVER1: Windows
  Server Version 5, Release 4, Level 1.0
  Server date/time: 10/23/2012 11:21:44  Last access: 10/23/2012 11:21:14

Incremental by snapshot difference of volume 'v:'
Connected to NetApp Filer netappl.example.com as user pete via HTTPS
NetApp Release 8.1.1RC1 7-Mode: Thu May 31 21:30:59 PDT 2012
Performing a Snapshot Differential Backup of volume
'\\netappl.example.com\petevol'
Creating Diff Snapshot.
Using Base Snapshot 'TSM_THIN5086B9441A1F8_PETEVOL' with timestamp 10/23/2012
15:36:53
Using Diff Snapshot 'TSM_THIN5086B9772AF8_PETEVOL' with timestamp 10/23/2012
15:37:44
Successful incremental backup of '\\netappl.example.com\petevol'
```

Related concepts:

“Snapshot differential backup with an HTTPS connection” on page 142

Related reference:

“Snapdiff” on page 529

Snapshotproviderfs

Use the snapshotproviderfs option to enable snapshot-based file backup and archive operations, and to specify a snapshot provider.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Specify this option in the client options file, `dsm.opt`, to enable snapshots. You can override the client-wide option for a specific operation by specifying this option on the command line for the backup and archive commands. You can also override the client-wide option for a specific file system by using the `include.fs` statement in the `dsm.opt` file. You can also set this option using the Preferences editor.

Syntax

►—SNAPSHOTPROVIDERFS— *value* —►

Parameters

value

Specifies one of the following values:

VSS

Specifies that VSS should be used to provide OFS support.

NONE

Specifies that no snapshot provider should be used; OFS support is turned off. This is the default.

Examples

Options file:

```
snapshotproviderfs VSS
include.fs d: snapshotproviderfs=vss
```

Command line:

```
-SNAPSHOTPROVIDERFS=VSS
```

Related information

For information about configuring open file support, see “Configuring Open File Support” on page 83.

Snapshotproviderimage

Use the `snapshotproviderimage` option to enable snapshot-based image backup, and to specify a snapshot provider.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager API does not support this option. The server can also define this option.

Options File

Specify this option in the client options file, `dsm.opt`, to enable snapshots for all the file systems on the client. You can override the client-wide option for a specific operation by specifying this option on the command line for the **backup image** command. You can also override the client-wide option for a specific file system using the `include.image` statement in the `dsm.opt` file. You can also set this option using the Preferences editor.

Syntax

►►—SNAPSHOTPROVIDERImage— *value* —►►

Parameters

value

Specifies one of the following values:

VSS

Specifies that the VSS should be used to provide online image support.

NONE

Specifies that no snapshot provider should be used. This turns off online image support. This is the default

Examples

Options file:

```
snapshotprovideri VSS
include.image d: snapshotprovideri=vss
```

Command line:

```
-SNAPSHOTPROVIDERImage=NONE
```

Related information

For information about configuring open file support, see “Configuring Open File Support” on page 83.

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.

This option should be used with an incremental backup of a NAS file server volume instead of a simple incremental or incremental with snapshotroot option whenever the NAS file server is running ONTAP V7.3 for performance reasons. The snapdiff and snapshotroot options should not be used together.

The snapshotroot option can be used to back up network share mounted file systems. Both the backup specification (source) and the snapshotroot value can be a network share mounted file specification. For example, the snapshotroot option can be used to back up a network share file system hosted on a network-attached storage (NAS) that supports snapshot.

In the following example, c:\snapshots\snapshot.0 is network share that is mounted from a NAS file server and \\florance\c\$ represents the snapshot that is created at the NAS file server.

```
dsmc incr \\florance\C$ -snapshotroot=c:\snapshots
\snapshot.0
```

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 c: drive and mounts it as the NTFS junction point \\florence\c\$\snapshots\snapshot.0. If you back up this data by using the following command, a unique file space that is called \\florence\c\$\snapshots\snapshot.0 is created on the server.

```
dsmc incremental \\florence\c$\snapshots\snapshot.0
```

However, you might want to associate the snapshot data with the data already processed for the c: drive (\\florence\c\$). Using the snapshotroot option, you can associate the data with the file space corresponding to the c: drive (\\florence\c\$) on the Tivoli Storage Manager server:

```
dsmc incr c: -snapshotroot=\\florence\c$\snapshots\snapshot.0
-or-
dsmc incr \\florence\c$ -snapshotroot=\\florence\c$\snapshots\
snapshot.0
```

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 incr c: -snapshotroot=\\florence\c$\snapshots\snapshot.1
```

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 incr c:\dir1\* -subdir=yes -snapshotroot=\\florence\c$\
snapshots\snapshot.1
dsmc sel c:\dir1\sub1\file.txt -snapshotroot=\\florence\c$\
snapshots\snapshot.1
dsmc archive c:\mydocs\*.doc -snapshotroot=\\florence\c$\
snapshots\snapshot.1
```

If you want to include or exclude specific file specifications, the include and exclude statements should contain the name of the file system that was the source of the snapshot (the c: drive), and not the name of the target of the snapshot (\\florence\c\$\snapshots\snapshot.1). 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 c:\dir1\...\*.txt lyrmgmtclass
exclude \\florence\c$\mydocs\*.doc
```

The following include-exclude statements are not valid because they contain the name of the snapshot:

```
include \\florence\c$\snapshots\snapshot.1\dir1\...\
*.txt lyrmgmtclass
exclude \\florence\c$\mydocs\*.doc
```

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 incr c: -snapshotroot=\\florence\c$\snapshots\snapshot.0
dsmc incr c:\dir1\* -snapshotroot=\\florence\c$\snapshots\
snapshot.0
```

The following command is invalid because it contains two file specifications:

```
dsmc incr c:\dir1\* e:\dir1\* -snapshotroot=\\florence\c$\  
snapshots\snapshot.0
```

The following command is invalid because it contains no file specification:

```
dsmc incr -snapshotroot=\\florence\c$\snapshots\snapshot.0
```

Notes:

1. Ensure that the `snapshotroot` 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 any backup command, such as **backup image**, or **backup systemstate**, and so on.
4. You cannot use the `snapshotroot` option with the `snappdiff` option.
5. Use the `snapshotroot` option with caution if you are using the Tivoli Storage Manager journal-based backup feature. Since there is no coordination between the Tivoli Storage Manager journal and the vendor-acquired snapshot provider (VSS), unwanted behavior can occur with journal notifications received after the snapshot occurs. For example, files might not be backed up, or they might be backed up redundantly to the Tivoli Storage Manager server.
6. 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:

- All Windows clients.

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 incr c: -SNAPSHOTRoot=\\florence\c$\snapshots\snapshot.0
```

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.opt`).

Syntax



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 client options file (dsm.opt) 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 Windows clients that support the image snapshot backup command. The server cannot define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client options file (dsm.opt) for the scheduler. You can set this option on the **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.


Supported Clients

This option is valid for all supported clients.

Options File

Place this option in the client options file (dsm.opt). You can also set this option on the **Communication** tab of the Preferences editor.

Syntax



Parameters

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.

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 42.

“Sslrequired” on page 544

“Tcpport” on page 559

Sslfipsmode

The `ssl fipsmode` 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.

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 59. 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
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.

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 client options (`dsm.opt`) 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. *No* is the default.

Yes Specifies that the backup-archive client requires that all SSL sessions use TLS 1.2 (or higher) protocol.

Examples

Options file:

```
ssldisablelegacytls yes
```

Command line:

Does not apply.

Related information

“Ssl” on page 542.

“Sslrequired” on page 544

“Tcpport” on page 559

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 by Active Directory objects query and restore operations, and by snapshot difference backup. The client uses the stagingdirectory location for temporary files when the client processes files that were migrated with IBM Tivoli Storage Manager HSM for Windows.

Supported Clients

This option is valid for all Windows clients. The server can also define this option.

Options File

Place this option in the client options file (dsm.opt). 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 checks for the existence of the USER environment variables in the following order, and uses the first path found:

1. The path that is specified by the TMP user variable
2. The path that is specified by the TMP system variable
3. The path that is specified by the TEMP user variable
4. The path that is specified by the TEMP system variable
5. The Windows system directory

In Uniform Naming Convention (UNC) format, the path must contain a drive letter. In the following UNC format example, the path contains the drive letter D\$: \\computer7\D\$\temp\tsmstaging.

Examples

Options file:

```
stagingdirectory c:\tsmdata
```

Command line:

```
-stagingdir="e:\tsmdata"
```

Related reference:

“Query Adobjects” on page 696

“Restore Adobjects” on page 738

“Diffsnapshot” on page 364

“Snapdiff” on page 529

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 restore /myfile.txt -pick
```


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 options file (`dsm.opt`).

Syntax



Parameters

No Subdirectories are not processed. This is the default.

Yes

Subdirectories are processed. Because the client program searches all

subdirectories of a directory that is being processed, processing can take longer to complete. Specify *Yes* only when necessary.

If you use the `preservepath` option in addition to `subdir=yes`, it can affect which subdirectories are processed.

Note:

1. When you run the client in interactive mode, and if you use the `-subdir=yes` option, the setting persists for all commands entered in interactive mode, until you end interactive mode, by typing `Quit`.
2. If `subdir=yes` is in effect when you restore multiple files, place a directory delimiter character at the end of the destination file specification. If the delimiter is omitted, the client displays a message indicating that the destination file specification is not valid.
3. It is a best practice to include only the default value for `subdir` (No) in a client options file or a client options set.

Examples

Options file:

```
subdir no
```

Command line:

To restore the structure:

```
\path2\dir1
\path2\dir1\file1
\path2\dir1\dir2
\path2\dir1\dir2\file1
```

enter any of the following commands:

```
rest \path\dir1\* \path2\ -su=yes
rest \path\dir1\file* \path2\ -su=yes
rest \path\dir1\file1* \path2\ -su=yes
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This 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 490

Subfilebackup

The `subfilebackup` option specifies whether to enable adaptive subfile backup.

Note:

1. To use the `subfilebackup` option, the server must have `subfile yes` or `subfile client set`. Ask your Tivoli Storage Manager administrator to set the correct values.
2. You can use open file support with subfile backup operations.
3. The `subfilebackup` option does not work correctly for migrated files. If you use a combination of `subfilebackup yes` and `subfilebackup no` for migrated files, your data might be corrupted on the server.

4. Adaptive subfile backup can be used with selective backups, incremental backups, and on file-level backups of VMware virtual machines (Backup VM commands where `VMBACKUPTYPE=FILE` is specified).
5. Adaptive subfile backup processing occurs only when you back up files; it is not used for archive operations.
6. Restoring files that were backed up with adaptive subfile backup is accomplished by using the **dsmc restore**.


Supported Clients

This option is valid for all 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 options file (`dsm.opt`). You can set this option on the **Backup** tab of the Preferences editor.

Syntax



Parameters

No Specifies that the backup-archive client does not use adaptive subfile backup. This is the default.

Yes Specifies that the backup-archive client uses adaptive subfile backup.

Examples

Options file:

```
subfileb yes
```

Command line:

```
-subfileb=yes
```

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

Subfilecachepath

The `subfilecachepath` option specifies the path where the client cache resides for adaptive subfile backup processing.

If you do not specify a path, the client creates a path called `\cache` under the directory where the backup-archive client executable files reside.

All directories and subdirectories in the path name you specify with the `subfilecachepath` option must exist. For example, if you specify `c:\temp\cache`, the `c:\temp` directory must already exist. You should exclude the `subfilecache` directory from backup.

Supported Clients

This option is valid for all Windows clients. The server can also define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Backup** tab of the Preferences editor.

Syntax

►►—SUBFILECACHEPath— —*path_name*—►►

Parameters

path_name

Specifies the path name where the client cache resides for adaptive subfile backup processing.

Examples

Options file:

subfilecachep c:\temp\cc_tsm

Command line:

-subfilecachep=c:\temp\cc_tsm

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

Subfilecachesize

The subfilecachesize option specifies the client cache size for adaptive subfile backup.

If the cache size is too small, base files for some files will not be cached and subfile processing will not apply for them. However, setting the value too large can take up more disk space than can be spared. The files maintained in the cache should closely reflect the files used on a regular basis.

Supported Clients

This option is valid for all Windows clients. The server can also define this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Backup** tab of the Preferences editor.

Syntax

►►—SUBFILECACHESize— —*size*—►►

Parameters

size

Specifies the size, in megabytes, of the client cache for adaptive subfile backup processing. The range is 1 through 1024 (1 GB); the default is 10.

Examples

Options file:

```
subfilecaches 10
```

Command line:

```
-subfilecaches=10
```

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

Systemstatebackupmethod

Use the `systemstatebackupmethod` option to specify which backup method to use to back up the system writer portion of the system state data. The method you select is used when you backup the system state data.

Supported clients

This option is valid for Windows clients.


Options file

Place this option in the client options file (`dsm.opt`). When specified in the `dsm.opt` file, the option affects system state backups created by **BACKUP SYSTEMSTATE** commands, and system state data backed up by **INCREMENTAL** commands. However, the only command that you can specify this option on is the **BACKUP SYSTEMSTATE** command.

Schedule definitions

You can also specify this option on the options parameter of a schedule definition on schedules that have both `action=backup` and `subaction=systemstate` set. Defining an infrequent schedule with this option set to **FULL** ensures that you periodically perform a full backup of Windows system state data.

Syntax



Parameters

PROgressive

With the **PROgressive** method, the system writer portion of the system state data is backed up using the progressive incremental backup method. That is, if system writer files have not changed since the last system state backup, they are not included in this backup. Only the changed system writer files are backed up. This is the default system state backup method.

This type of system state backup uses the least network bandwidth and Tivoli Storage Manager server storage, but it increases the amount of server database processing required to keep track of the changes.

OPPortunistic

With the OPPortunistic method, if any system writer files have changed since the last system state backup, all system writer files are backed up.

This method, like the PROGressive method, also uses the least network bandwidth and Tivoli Storage Manager server storage if system writer files have not changed since the last system state backup. If any system writer files have changed since the last system state backup then the system writer is backed up in full, which uses more network bandwidth and server storage. With the OPPortunistic method, the amount of server database processing that occurs is less than that caused by the PROGressive method.

FULL

When FULL is specified, all system writer files are backed up, even if they have not changed since the last system state backup.

This type of system state backup uses the most network bandwidth and Tivoli Storage Manager server storage because all system writer files are backed up during each system state backup operation. However, this system state backup method causes little server database processing.

Examples

Options file:

```
SYSTEMSTATEBACKUPMETHOD FULL
```

```
SYSTEMSTATEBACKUPMETHOD OPPORTUNISTIC
```

Command line:

```
backup systemstate -SYSTEMSTATEBACKUPMETHOD=FULL
```

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 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 client options file (`dsm.opt`). 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 options file (`dsm.opt`). 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 client options file (`dsm.opt`).

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 347 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 client options file (`dsm.opt`). You can set this option on the **Scheduler** tab, **Your TCP/IP address** field of the Preferences editor.

Syntax

►►—TCPCLIENTAddress— *—client_address—*—————►►

Parameters

client_address

Specifies the TCP/IP address you want the server to use to contact your client node. Specify a TCP/IP Internet domain name or a numeric IP address. The numeric IP address can be either a TCP/IPv4 or TCP/IPv6 address. You can only use IPv6 addresses if you specified the `commmethod V6Tcpip` option.

Examples

Command line:

```
-tcpclientaddress=192.0.2.0  
-tcpclientaddress=example.mycompany.mydomain.com  
-tcpclientaddress=2001:0DB8:0:0:0:0:0:0
```

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 client options file (dsm.opt). 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 Windows clients.

Options File

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

-tcpport=1501

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

Tcpserveraddress

The tcpserveraddress option specifies the TCP/IP address for a Tivoli Storage Manager server. You can obtain this server address from your administrator.

Supported Clients

This option is valid for all clients.

Options File

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

-tcpserveraddress=129.33.24.99

-tcpserveraddress=2002:92b:111:221:128:33:10:249

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

Tcpwindowsize

Use the tcpwindowsize option to specify, in kilobytes, the size you want to use for the TCP/IP sliding window for your client node.

The sending host cannot send more data until it receives an acknowledgment and a TCP receive window update. Each TCP packet contains the advertised TCP receive window on the connection. A larger window allows the sender to continue sending data and can improve communication performance.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Communication** tab, **Window Size** field of the Preferences editor.

Syntax

►—TCPWindowSize— *window_size* —►

Parameters

window_size

Specifies the size, in kilobytes, to use for your client node TCP/IP sliding window. The range of values is 0 through 2048. A value of 0 allows 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. 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.
- Windows provides a larger TCP receive window size when communicating with hosts that also provide this support, known as RFC1323. In these environments, a value greater than 63 can be useful.

Examples

Options file:

```
tcpwindowsize 63
```

Command line:

```
-tcpw=63
```

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

Timeformat

The timeformat option specifies the format in which you want to display and enter system time.

Use this option if you want to change the default time format for the language of the message repository you are using.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time the client is called. Consult the documentation on your local system for details about setting up your locale definition.

Note: The timeformat option does not affect the web client. The web client uses the time format for the locale that the browser is running in. If the browser is not running in a locale that 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 asr**
- **query backup**
- **query filespace**
- **query image**
- **query nas**
- **query systemstate**
- **restore**
- **restore image**
- **restore nas**
- **restore registry**
- **restore systemstate**
- **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 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.

- 1 23:00:00
- 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`, `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`.

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 valid for all Windows 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 client options file (`dsm.opt`).

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

- 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 query backup -totime=23:59:00 -todate=06/30/2003 c:\mybackups\
```

Txnbytelimit

The `txnbytelimit` option specifies the number of kilobytes the client program buffers before it sends a transaction to the server.

A *transaction* is the unit of work exchanged between the client and server. 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 options file (`dsm.opt`). 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:

txn 25600

txn 2097152

txn 2097152k

txn 2048m

txn 2g

txn 32G

Command line:

-txn=25600

-txn=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 or FILER.

This option is valid for all Windows 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
```

Usedirectory

The `usedirectory` option queries the Active Directory for the communication method and server with which to connect.

This option overrides the `commmethod` parameters specified in the client options file (`dsm.opt`). Optimally, the administrator enables only one server and one specific communication protocol for a given client node. The specification of this information in Active Directory is done using the Tivoli Storage Manager server on Windows, which has a wizard to assist with this configuration. If a node is registered to more than one server published in Active Directory, the first server returned in the Active Directory query is used. If the client cannot contact the server, the client session fails.


Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Communication** tab of the Preferences editor.

Syntax



Parameters

Yes

Specifies that the client ignores `commmethod` parameters set in the client options file and query the Active Directory for the communication method and server with which to connect.

No Specifies that the client uses the communication method specified in the option file. If there is no communication method specified in the option file the default communication method and server are used.

Examples

Options file:

usedirectory no

Command line:

-usedir=yes

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

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 Windows 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 in the client options file (`dsm.opt`).

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 60

Related tasks:

“Configuring the client for automated failover” on page 63

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 c:\relx\dir1\ -v2archive -su=y
```

Archives these files:

```
c:\relx\dir1\file1  
c:\relx\dir1\file2  
c:\relx\dir1\file3  
c:\relx\dir1\dir2\file4  
c:\relx\dir1\dir2\file5
```

Note: The client does not archive `c:\relx\dir1` and `c:\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 Windows clients that are configured to back up VMware vCloud vApps.

Options File

Place this option in the client options file (dsm.opt), 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 MCPRODVAPPS
```

Command line:

```
dsmc backup vapp "org1,vdc1,vapp1" -vappmc=myManagmentClass
```

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 Windows 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 Windows 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 Windows 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.

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 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 Windows 32-bit clients *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 d: -verifyimage

Virtualfsname

Use the `virtualfsname` option with the **backup group** command to specify the name of the virtual file space for the group on which you want to perform the operation. The `virtualfsname` cannot be the same as an existing file space name.

Supported Clients

This option is valid for all Windows clients.

Syntax

►►—VIRTUALFsname =— —fsname————►►

Parameters

fsname

Specifies the name of the container for the group on which you want to perform the operation.

Examples

Command line:

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

Virtualnodename

The `virtualnodename` option specifies the node name of your workstation when you want to restore or retrieve files to a different workstation.

When you use the `virtualnodename` option in your client options file, or with a command:

- You must specify the name you specified with the `nodename` option in your client options file (`dsm.opt`). This name should be different from the name returned by the **hostname** command on your workstation.
- The client prompts for the password assigned to the node that you specify, if a password is required (even when the `passwordaccess` option is set to generate). If you enter the correct password, you have access to all backups and archives that originated from the specified node.

When connecting to a server, the client must identify itself to the server. This login identification is determined in the following ways:

- If the `nodename` and `virtualnodename` options are not specified, or a virtual node name is not specified on the command line, the default login ID is the name returned by the **hostname** command.
- If the `nodename` option is specified, the name specified with the `nodename` option overrides the name returned by the **hostname** command.
- If the `virtualnodename` option is specified, or a virtual node name is specified on a command line, it cannot be the same name as the name returned by the **hostname** command.

Note: The client can use file space information when restoring files. The file space information can contain the name of the computer from which the files were backed up. If you restore from another client node and do not specify a destination for the restored files, the client uses the file space information to restore the files. In such a case, the client attempts to restore the files to the file system on the original computer. If the restoring computer has access to the file system of the original computer, you can restore files to the original file system. If the restoring computer can not access the file system of the original computer, the client can return a network error message. If you want to restore the original directory

structure but on a different computer, specify only the target file system when you restore. This is true when restoring files from another node and when retrieving files from another node.

Supported Clients

This option is valid for all clients.

Options File

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

Syntax

►►—VIRTUALNodename— —*nodename*—————►►

Parameters

nodename

Specifies a 1- to 64-character name that identifies the node for which you want to request Tivoli Storage Manager services. There is no default.

Examples

Options file:

virtualnodename cougar

Command line:

-virtualn=banshee

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

Vmautostartvm

Use the `vmautostartvm` option with the **restore VM** `vmrestoretype=instantaccess` command to specify whether the VM created during instant access processing is automatically powered on.

This option is only valid for VMware virtual machines. The virtual machines must be hosted on VMware ESXi 5.1 servers, or later versions. To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

Supported Clients

This option can be used with supported Windows clients.

Options file

Place this option in the client options file (dsm.opt), or on the command line. This option is only valid when used for an operation where `vmrestoretype=instantaccess`.

Syntax



Parameters

NO The VM created for instant access is not started automatically. The VM must be started manually. This is the default setting. The default provides an opportunity to reconfigure the VM before you power it on, to avoid potential conflicts with existing virtual machines.

YES

The VM created for instant access is started automatically.

Examples

Options file:

```
VMAUTOSTARTvm NO
```

Command line:

```
dsmc restore vm Oslo -VMRESToretype=INSTANTAccess -vmname=Oslo_verify  
-VMAUTOSTARTvm=YES
```

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 Windows clients that are installed on a vStorage backup server or a Hyper-V 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

►—VBACKDir—directory—►

Parameters

directory

Specifies the path where the control files are stored on the backup server.

The default is `c:\mnt\tsmvmbackup\fullvm\`

Examples

Options file:

```
VBACKD c:\mnt\tsmvmbackup\
```

Command line:

```
dsmc backup vm -VBACKUPT=fullvm -VBACKD=G:\virtual_machine\  
control_files\
```

```
dsmc restore vm -VBACKUPT=fullvm -VBACKD=G:\san_temp\
```

Vmbacknodelete

Use the `vmbacknodelete` option with the **backup VM** command to specify that the full VM backup files copied to the backup proxy are not deleted at the end of the backup.

Supported Clients

Use this option to keep one or more of your most recent backups on the backup proxy, which allows for faster restore operations that do not require pulling the backup files from the Tivoli Storage Manager server.

The full VM backup staging directory is `BACKUP CURRENT`. This directory is where VMware exports the full VM files, and where the client performs the backup. When the backup is complete, this directory is either deleted, or if you have specified the `vmbacknodelete` option, this directory is renamed with a date and time stamp of the full VM backup.

For example:

```
C:\mnt\tsmvmbackup\fullvm\myVMname\BACKUP_DATE_yyyy_mm_dd[hh_mm_ss]\
```

where `vmbackdir=c:\mnt\tsmvmbackup` and *myVMname* is the name of your vmware guest.


Important: If you use the `vmbacknodelete` option, delete saved backup copies to ensure that there is sufficient free space available for processing future backups.

This option is valid for Windows clients that are configured as a vStorage backup server of a VMware virtual machine. You can use this option with VCB full VM backups. You cannot use this option VSTOR full VM backups

Options File

Place this option in the client options file (`dsm.opt`), or on the command line.

Syntax



Parameters

no Specifies that the full VM backup files that are copied to the backup proxy is deleted at the end of the backup. This is the default.

yes

Specifies that the full VM backup files that are copied to the backup proxy is not deleted at the end of the backup, and they are saved in a directory with a name that indicates the date and time of the backup.

Examples

Options file:

VMBACKNOD yes

Command line:

dsmc backup vm -vmlist=vm1 -VMBACKUPT=fullvm -VMBACKNOD=yes

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 options file (dsm.opt).

Syntax



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 file-level backup, a VMware full VM backup, or a Hyper-V full VM backup.

Supported Clients

This option is valid on Windows clients that are installed on either a vStorage backup server or on a Microsoft Hyper-V system. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`), or on the command line.

Syntax



Parameters

Fullvm

Specify this value to run a traditional full VM backup of a VMware virtual machine. This is the default backup type for Windows clients that run on Windows server systems, where the Hyper-V server role is not enabled. Contrast with `vmbackuptype=hypervfull`.

File

Specify this value to perform a file-level backup of a VMware virtual machine. To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

HYPERVFULL

Specifies that you are backing up one or more Hyper-V virtual machines. If you enable the Hyper-V server role, this value is the default backup type. If you specify `vmbackuptype=hypervfull`, all of the options that are associated with backing up VMware files from a vStorage backup server are ignored (for example: `VMCHOST`, `VMCUSER`, `VMCPW`, `VMFULLNODELETE`).

Examples

Options file:

```
VMBACKUPT full
```

Command line:

```
dsmc backup vm vm1 -VMBACKUPT=file -vmchost=virtctr  
-vmcuser=virtctr_admin -vmcpw=xxxxx
```

Performs a file-level 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`.

```
dsmc backup vm -VMBACKUPT=file -vmhost=vmsrv1 -vmcuser=root
-vmcpw=xxxxx
Performs a file-level virtual-machine backup of vm1.example.com and
vm2.example.com using the VMware ESX Server machine
vmserv1.example.com, to the Tivoli Storage Manager server, using machine
names vm1 and vm2.
```

```
dsmc backup vm -domain.vmfile=vm1[xxxxxxx]
Performs a file-level virtual-machine backup of vm1 with the IP address of
VM, using the VMware ESX Server machine vmserv1.example.com, to the
Tivoli Storage Manager server with machine name vm1.
```

```
dsmc backup vm vm1 -VMBACKUPT=full -vmhost=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.
```

```
dsmc backup vm -VMBACKUPT=hypervfull -vmlist="VM 1,VM 2"
Performs a full virtual-machine backup of Hyper-V virtual machines
named "VM 1" and "VM 2", to the Tivoli Storage Manager server.
```

Vmbackvcbtransport

Use the `vmbackvcbtransport` option with the **backup VM** command to specify the transport mode to be used with VMware Consolidated Backup (VCB).

The `vmbackvcbtransport` option determines how VCB accesses virtual disk data.

Important:

- Full VM backup uses the VMware VCB support. Use the `vmbackvcbtransport` option to control the transport mode for full VM backups.
- File-level backup uses the vStorage for Data Protection APIs, and determines the best available transport. The `vmbackvcbtransport` option is not necessary for file-level backups.


Supported Clients

This option is valid only on supported Windows clients that are installed on a vStorage backup server. The server can also define this option.

Options File

Place this option in the client options file, or on the command line.

Syntax



Parameters

RETRY

Attempt access using the SAN to the virtual machine data from the backup proxy. If the SAN access is not available, try the operation again using the LAN transport. *RETRY* is the default.

SAN

Use the SAN to access virtual machine data from backup proxy.

LAN

Use the LAN to access virtual machine data from the backup proxy. LAN access is the same as the Network Block Device (NBD) protocol, which can be set on the VCB **vcbmounter** command.

HOTadd

Use the VMWare hotadd transport to access virtual machine data from the backup proxy. For information about the VMWare hotadd transport method, see the VMWare product documentation at <http://www.vmware.com>.

Note: The additional setup that is required to use hotadd transport that is documented in VMware documentation is not required for file-level backups.

Examples

Options file:

```
VMBACKVCBTransport lan
```

Command line:

```
dsmc backup vm vmlocal -vmbackuptype=fullvm -vmbackvcbtransport=lan
```

Related reference:

"Vmbackuptype" on page 581

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.

This option is not supported for Hyper-V backups.

Options File

Place this option in the client options file (**dsm.opt**), or on the command line.

Syntax

►►—VMCHost— —*hostname*—————►►

Parameters

hostname

Specifies the host name of the VMware VirtualCenter or ESX server 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 Windows clients that are installed on a vStorage backup server that is used to backup a VMware virtual machine. This option is not valid for Hyper-V backups.

Options File

Place this option in the client options file (`dsm.opt`), 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 781

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.

This option is valid for clients that act as data mover nodes that protect VMware virtual machines or on clients that protect Microsoft Hyper-V 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 client options file `dsm.opt`.

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.

This option is not valid for Hyper-V backups.

Options File

Place this option in the client options file (`dsm.opt`), or on the command line.

Syntax

►►—VMCUser— *username* —►►

Parameters

username

Specifies the user name of the VMware VirtualCenter or ESX server 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 Windows 64-bit clients.

Options file

You can specify this option in the client options file (`dsm.opt`) 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 658

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 Windows clients that are installed on a vStorage backup server.

Options file

Place this option in the client options file (dsm.opt), 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 590

“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 Windows clients that are installed on a vStorage backup server.

Options file

Place this option in the client options file (dsm.opt), 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 -VMDEFAULTDVPORTGROUP=dvPortGroup
```

Related reference:

“Vmdefaultdvportgroup” on page 588

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 Windows clients that are installed on a vStorage backup server.

Options file

Place this option in the client options file (dsm.opt), 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 588

“Vmdefaultdvswitch” on page 589

Vmdiskprovision

Use the `vmdiskprovision` option to specify a provisioning policy for the virtual disk file that is used to restore VMware virtual machine data. This option is valid only for **restore vm** operations where `vmrestoretype=instantrestore` is specified.

This option is only valid for VMware virtual machines. The virtual machines must be hosted on VMware ESXi 5.1 servers, or later versions. To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

Supported Clients

This option can be used with supported Windows clients.

Options file

Place this option in the client options file (`dsm.opt`), or on the command line.

Syntax



Parameters

THICK

Creates a virtual disk in a default thick format; where the space that is required for the virtual disk is allocated when the virtual disk is created. This setting is the default value.

THIN

Creates a virtual disk in a thin format.

Note: If you are restoring a virtual machine and you specify thin provisioning, the datastore that you restore the VM to must have enough free space to accommodate the total capacity of the VM disk, and not just the amount of disk that is used. For example, if a thin-provisioned VM has 300 GB total capacity for its disk, you cannot restore that VM to a datastore that has less than 300 GB available, even if only a portion of the total capacity is being used.

Examples

Options file:

VMDISKPROvision THIN

Command line:

```
dsmc restore vm Mainz -VMRESToretype=INSTANTRestore  
-VMTEMPDatastore=Temporary Datastore -VMDISKPROvision=THIN
```

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 **vmstortransport** 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 Windows clients.

Options file

You can set this option on the command line, in the 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 658

“Restore VM” on page 758

“Domain.vmfull” on page 381

Vmexpireprotect

This option is used to protect virtual machine snapshots so they cannot be expired while an instant restore or instant access operation (of VMware VMs) or a file-level restore (of a Hyper-V or VMware VM) is in progress.

During a mount or restore operation, the snapshot on the Tivoli Storage Manager server is locked to prevent it from expiring during the operation. Expiration might occur because another snapshot is added to the snapshot sequence. This option specifies whether to prevent or allow snapshot expiration during a mount or a restore operation.

Supported Clients

This option can be used with supported Windows clients that are configured to restore virtual machines.

Options File

For restoring VMware virtual machines, set this option in the client options file (dsm.opt) or on a **restore vm** command where the **vmrestoretype** option is set to **instantaccess** or **instantrestore**.

For restoring file-level backups for virtual machines, specify this option in the client options file, or on the **restore vm** command.

Syntax

►►—VMEXPIREPROTECT—►►

Parameters

Yes

Specify Yes to protect the snapshot from expiration. The snapshot on the Tivoli Storage Manager server is locked and the snapshot is protected from expiration during a mount or a restore operation.

No Specify No to disable expiration protection. This value is the default. The snapshot on the Tivoli Storage Manager server is not locked and the snapshot is not protected from expiration. If the snapshot that is being mounted or restored is expired, the result of the mount or restore operation is unpredictable. For example, the mount point can become unusable or contain errors. However, expiration does not affect the current active copy of the virtual machine. The active copy cannot expire during an operation.

When the snapshot is on a target replication server, the snapshot cannot be locked because it is in read-only mode. A lock attempt by the server causes the mount or restore operation to fail.

To avoid the lock attempt and prevent such a failure, disable expiration protection by specifying No, or by allowing this option to default.

Examples

Client options file:

```
VMEXPIREPROTECT YES
```

Command line:

Run an instant access operation for a VMware virtual machine:

```
dsmc restore vm vm1 -vmname=new_vm1 -vmrestoretype=instantaccess  
-vmexpireprotect=no
```

To restore files from a virtual machine backup, use the Tivoli Storage Manager recovery agent GUI. For information about the Tivoli Storage Manager recovery agent, see the Tivoli Storage Manager for Virtual Environments documentation.

Vmiscsiadapter

This option specifies which iSCSI adapter, on the ESX host, to use for instant restore and instant access operations for VMware virtual machines.

Supported Clients

This option is valid for 64-bit Windows clients that are configured as data movers that backup VMware virtual machines.

Options File

Set this option in the client options file (dsm.opt). You can also specify this option as a command-line parameter on the **restore vm** command that initiates an instant restore or instant access operation. To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

Syntax

►—VMISCSIAdapter=—*iSCSI_adapter_name*—◄

iSCSI_adapter_name

Specifies the name of the iSCSI adapter to connect to on the ESX host. If you do not specify this option, the first iSCSI adapter that is found on the host is used.

Examples

Options file:

```
vmiscsiadapter "vmhba36"
```

Command line:

```
dsmc restore vm "Haifa" -VMRESToretype=INSTANTAccess  
-vmname="Haifa_verify" -VMISCSIAdapter="vmhba36"
```

Vmiscsiserveraddress

Use the `vmiscsiserveraddress` option with the **restore VM** command to specify the host name or the IP address of the iSCSI server that provides the iSCSI targets for instant restore and instant access operations.

The `vmiscsiserveraddress` option is valid for all instant operations (`vmrestoretype=instantaccess` and `vmrestoretype=instantrestore`) for VMware virtual machines.

The virtual machines must be hosted on VMware ESXi 5.1 servers, or later versions. To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

Supported Clients

This option can be used with supported Windows clients.

Options file

Place this option in the client options file (`dsm.opt`), or on the command line.

Syntax

►►—VMISCSIServeraddress— *iSCSI serverhost name or IP address*————►►

Parameters

iSCSI serverhost name or IP address

Specify the host name or IP address of the iSCSI server that supplies the iSCSI target disks. This iSCSI server must connect the data mover machine with all of the ESX hosts that are used for instant restore operations. If `vmiscsiserveraddress` is not specified, the host name or IP address of the data mover machine is used.

For instant restore operations, the IP address of the network card in the data mover machine that is used for the iSCSI transfer should be in the same subnet as the iSCSI adapter on the ESX host.

For file restore mount operations, the Windows and Linux mount proxy systems must be in the same network range.

Examples

Options file:

```
VMISCSIServeraddress 192.168.42.50
```

Command line:

```
dsmc restore vm Oslo -VMRESToretype=INSTANTAccess -vmname=Oslo_verify  
-VMISCSIServeraddress=odin.oslo.no.xyzco.com
```

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 Windows clients.

Options file

This option is valid 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 `vmlimitperdatastore` option is enforced even when virtual machine data exists in two or more datastores. Adjust the values of `vmmaxparallel`, `vmlimitperhost`, and `vmlimitperdatastore` 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 658

“Domain.vmfull” on page 381

“Vmmaxparallel” on page 598

“Vmlimitperhost”

Vmlimitperhost

The `vmlimitperhost` 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 `vmlimitperdatastore` 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 Windows clients.

Options file

This option is valid 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`, `vmlimitperhost`, and `vmlimitperdatastore` 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 658

“Domain.vmfull” on page 381

“Vmmaxparallel” on page 598

“Vmlimitperhost”

Vmlist

Use the `vmlist` option with the **backup VM** command to specify one or more Hyper-V virtual machines to include in the backup operation.

This option is only valid for backing up Hyper-V virtual machines.

Supported Clients

This option is valid for Windows clients that are installed on a Hyper-V host. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`), or on the command line.

Syntax

►►—VMList— —*vmname*—◄◄

Parameters

vmname

Specifies the virtual machine name to back up or a list of virtual machine names. Separate virtual machine names with commas.

Examples

Options file:

```
VMList vm1,vm2,vm3
```

Command line:

```
dsmc backup vm -vmlist="vm1,vm2,vm3" -vmbackuptype=HYPERVFULL
```

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 `vmlimitperhost` and `vmlimitperdatastore` 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 Windows clients.

Options file

This option is valid in the client options file (`dsm.opt`) 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 658

“Domain.vmfull” on page 381

“Vmlimitperhost” on page 597

“Vmlimitperdatastore” on page 596

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 `vmskipmaxvirtualdisks` 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 `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.

Supported clients

This option is valid for 64-bit Windows clients that are configured as data movers that back up VMware virtual machines.

Options file

Set the `vmmaxvirtualdisks` option in the client options file (`dsm.opt`). 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 `vmskipmaxvirtualdisks 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 `vmskipmaxvirtualdisks 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 -vmskipmaxvirtualdisks=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 -vmskipmaxvirtualdisks=no
```

Back up VMDKs that are 8 TB or smaller and exclude VMDKs that are larger than 8 TB:

```
backup vm VM1 -vmaxvirtualdisks=8 -mskipmaxvirtualdisks=yes
```

or

```
backup vm VM1 -vmaxvirtualdisks=999 -mskipmaxvirtualdisks=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 or Hyper-V virtual machines. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`), 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 `myManagementClass`.

```
dsmc backup vm "myVirtualMachine" -vmmc=myManagementClass
```

Vmmountage

Use the `vmmountage` option with the **restore VM "*"**

`-vmrestoretype=mountcleanupall` command to specify the number of hours that a VM file level restore mount must be active to be cleaned up.

Supported Clients

This option is only valid for Windows clients.

Options File

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

Syntax

►►—VMMOUNTAge =— —*hours*—————◄◄

Parameters

hours

Specifies the number of hours that a VM file level restore mount must be active to be cleaned up. All active mount operations that exceed this period are cleaned up.

The value that is specified must be an integer between 0 and 10000. The default is 0.

Examples

Command line:

Clean up all mount operations that are active longer than 24 hours:

```
dsmc restore vm "*" -VMRESToretype=MOUNTCLEANUPALL -VMMOUNTAge=24
```

Clean up all active mount operations:

```
dsmc restore vm "*" -VMRESToretype=MOUNTCLEANUPALL -VMMOUNTAge=0
```

or

```
dsmc restore vm "*" -VMRESToretype=MOUNTCLEANUPALL
```

Vmnoprmdisks

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), 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 604

“Vmprocessvmwithprdm” on page 606

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`), 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 606

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 options file (`dsm.opt`).

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`) 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`” on page 606

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`) 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 604

Vmrestoretype

Use the **vmrestoretype** option with the **query VM** or **restore VM** commands to specify the type of restore operation to perform or query.

This option is only valid for VMware virtual machines. The virtual machines must be hosted on VMware ESXi 5.1 servers, or later versions. To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.


Supported Clients

This option can be used with supported Windows clients.

Options file

This option must be specified on the command line of a **restore vm** or **query vm** command. You cannot set this option in the client options file.

Syntax



Parameters

noninstant

Specifies that classic full VM restore is performed. This is the default restore type.

instantrestore

Specifies that an instant restore is performed. During an instant restore operation, the VM is started during the restore operation. When this restore type is specified on a **query VM** command, the command returns a list of VMs running an instant restore operation.

Important: For instant restore operations, ensure that both the temporary datastore that you specify with the **vmtempdatastore** option and the VMware datastore that is specified by the **datastore** option on the restore VM

command have enough free storage to save the virtual machine that you are restoring, and the snapshot file that contains changes made to the data.

instantaccess

Specifies that a temporary restore of the backed-up VM is performed. Use this restore type when you want to restore a VM temporarily, to test the integrity of a backup, before you run an instant restore. Any changes that are made to the temporary VM are not saved.

When this restore type is specified on a **query vm** command, the command returns a list of VMs that are running an instant access operation.

vmcleanup

Specifies that a cleanup of the selected VM and its components is performed.

For instant access operations, this option removes the temporary VM and all of its components.

For instant restore operations, this option removes only the components that are no longer needed (for example the iSCSI mounts). The virtual machine is not removed. Cleanup operations are not allowed when the VM is still running on the iSCSI disks. To force this behavior see **vmfullcleanup**.

vmfullcleanup

The VM and all its components are removed regardless of the current state. Do not start a full clean up operation while vMotion is still migrating a virtual machine.

alltype

Queries all active instant access and instant restore sessions.

mountcleanupall

Cleans up active VM file level restore mount operations that are older than the period specified with the **vmmountage** option. You must specify **restore vm ""** to use the **mountcleanupall** option.

Examples

Command line:

Perform an instant access of the VM named Oslo. The original VM still exists. As a result, the **-vmname** option is used to assign the new name Oslo_verify.

```
dsmc restore vm Oslo -vmrest=instantaccess -vmname=Oslo_verify
```

Perform an instant restore of the VM named Cologne.

```
dsmc restore vm Cologne -vmrest=instantrestore  
-vmtempdatastore=Verify_datastore
```

Perform a regular (full VM) restore of the virtual machine named San_Jose.

```
dsmc restore vm San_Jose
```

Alternatively, you can also use the following command: **dsmc restore vm San_Jose -vmrest=noni**

Perform an instant restore of the VM named Oslo, with the **-pick** option to choose a specific backup version.

```
dsmc restore vm Oslo -vmrest=instantrestore -pick
```

Perform a cleanup of the VM and all its components. These components include iSCSI mounts, devices, and temporary data that are associated with the VM name, on the ESX host.

```
dsmc restore vm Oslo -VMRESToretype=VMCleanup -vmname=Oslo_Verify
```

Perform a query to find all active instant restore sessions and display an abbreviated status for each.

```
dsmc query vm * -VMRESToretype=INSTANTRestore
```

Perform a query to find all active instant restore mode and instant access mode virtual machines.

```
dsmc query vm * -VMRESToretype=ALLtype
```

Perform a query to find all active instant restore mode virtual machines, and obtain detailed status for each virtual machine.

```
dsmc query vm * -VMRESToretype=INSTANTRestore -Detail
```

Perform a query to find all active instant access sessions.

```
dsmc query vm * -VMRESToretype=INSTANTAccess
```

Perform a mount cleanup of all mount operations that are active longer than 24 hours.

```
dsmc restore vm "*" -vmrestoretype=mountcleanupall -vmmountage=24
```

Related reference:

“Scenarios for running full VM instant access and full VM instant restore from the backup-archive client command line” on page 214

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

This option can be used with supported Windows and Linux 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 Windows clients that are configured as data movers that back up VMware virtual machines.

Options file

Set the vmskipmaxvirtualdisks option in the client options file (dsm.opt). 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 610 for information about setting this option.

Vmstoragetype

Use the `vmstoragetype` option with the **restore VM** command to specify the storage device type from which the snapshot is mounted with Tivoli Storage Manager recovery agent.

You can specify the `vmstoragetype` option with the **restore VM** `-VMRESToretype=INSTANTRestore` or **restore VM** `-VMRESToretype=INSTANTAccess` commands.

When `vmstoragetype` is specified, it is not necessary to set the storage type option in the Tivoli Storage Manager recovery agent GUI. The `vmstoragetype` overwrites the storage type setting in the recovery agent GUI.

Supported Clients

This option is valid on Windows only.

Options File

Place this option in the client options file (`dsm.opt`) on the Windows mount proxy system, or on the command line.

Syntax



Parameters

DISK

The snapshots to be mounted by the recovery agent are on Disk or File storage pools. This value is the default.

VTL

The snapshots to be mounted by the recovery agent are on VTL storage pools.

TAPE

The snapshots to be mounted by the recovery agent are on Tape storage pools.

Examples

Options file:

```
VMSTORAGETYPE TAPE
```

Command line:

Restore a virtual machine that is named Orion by using the following command:

```
dsmc restore vm Orion -Host=esxi.example.com -datacenter=mydatacenter  
-VMTEMPDatastore=temp_datastore -VMRESToretype=INSTANTRestore  
-datastore=mydatastore -VMSTORAGETYPE=VTL
```

This command specifies the name of the virtual machine to restore, the host and data center to restore it to, and the restore type (-VMRESToretype=INSTANTRestore). The -VMSTORAGETYPE=VTL option identifies the snapshot (Orion) that is to be mounted by the recovery agent is on VTL storage pools. The **VMTEMPDatastore** option is a mandatory parameter for instant restore operations.

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

Category	Tag	Tag description
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 higher.
- The VMware Platform Services Controller (PSC) and the vCenter server must be installed on the same host.
- 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.

- A 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, by running the configuration wizard to save the data mover node password, or by using the **dsmc set password** command in the backup-archive command line. Other data movers can be installed elsewhere.

Supported clients

This option can be used with supported Windows 64-bit clients.

Options file

You can specify this option in the client options file (dsm.opt) 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.vmfll client option setting. While the domain.vmfll option defines what virtual machines the client protects, the "Excluded" tag overrides what is defined by the domain.vmfll 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.vmfll" on page 381

"Include.vm" on page 438

"Vmmc" on page 601

"Vappmc" on page 571

| "Vmctlmc" on page 585
| "Set Vmtags" on page 787

Vmtempdatastore

Use the `vmtempdatastore` option with the **restore VM** command to define a temporary datastore on the ESX host for an instant restore operation.

The datastore created with the `vmtempdatastore` option is used to temporarily store the configuration of the VM created during restore processing. This option is required during instant restore operations (`-vmrestoretype=instantrestore`).

This option is only valid for VMware virtual machines. The virtual machines must be hosted on VMware ESXi 5.1 servers, or later versions. To use this option, you must have a license agreement to use Tivoli Storage Manager for Virtual Environments: Data Protection for VMware.

Supported Clients

This option can be used with supported Windows clients.

Options file

Place this option in the client options file (`dsm.opt`), or on the command line.

Syntax

►►—VMTEMPDatastore— *datastore_name*—————►►

Parameters

datastore_name

Specify the name of an existing datastore on the ESX host. The temporary datastore must be different from the original datastore, or the datastore specified by the `datastore` option. The datastore that you specify must be a VMFS datastore.

Examples

Options file:

```
VMTEMPDatastore Verify_Datastore
```

Command line:

```
dsmc restore vm Oslo -VMRESToretype=INSTANTAccess  
-vmname=Oslo_instant_restored -VMTEMPDatastore=Temporary_Datastore
```

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 Windows clients that act as a data mover for VMware guest backups.

Options file

Set this option in the client options file (`dsm.opt`).

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/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
```

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 Windows clients that act as a data mover for VMware guest backups.

Options file

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

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
```

Vmvstortransport

The vmvstortransport 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 the client options file (`dsm.opt`).

Supported clients

This option is valid for Windows 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
`VMSTORTRANSPORT san`

The backup-archive client is running in a virtual machine, but do not use the hotadd transport
`VMSTORTRANSPORT nbdssl:nbd`

Use the LAN transport, even if nbdssl is available, to obtain better performance
`VMSTORTRANSPORT nbd`

The SAN transport is preferred, but use nbd when the SAN is not available, and do not use nbdssl or hotadd
`VMSTORTRANSPORT san:nbd`

Vssaltstagingdir

The `vssaltstagingdir` option specifies the fully qualified path that contains the system exclude cache and temporary data for VSS snapshot operation.

The backup-archive client determines the path for temporary VSS files from the following prioritized choices:

1. The `vssaltstagingdir` option is defined in the `dsm.opt` file.
2. The `c:\adsm.sys` directory exists and is not empty.
3. If the `vssaltstagingdir` is not defined and the `c:\adsm.sys` directory does not exist, the client gets the path from a registry key. The path for temporary VSS files is the `DefaultVssStagingDir` value, and is generated from the `Path` value under the `HKLM\SOFTWARE\IBM\ADSM\CurrentVersion\BackupClient` key. After the `DefaultVssStagingDir` value is created, the value is not changed if the client is reinstalled to a new location.

Supported Clients

This option is valid for all Windows clients.

Options File

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

Syntax

►►—VSSALTSTAGINGDIR—*filepath*—————◄◄

Parameters

filepath

Specify the fully qualified path for temporary files that are related to VSS snapshot operations. If any part of the path does not exist, the backup-archive client attempts to create it. The default value is the client installation directory.

In Uniform Naming Convention (UNC) format, the path must contain a drive letter. In the following UNC format example, the path contains the drive letter `D$`: `\\computer7\D$\temp\snapshot`.

Examples

Options file:

```
vssaltstagingdir "c:\Users\All Users\Tivoli\adsm.sys"
```

Command line:

```
-vssaltstagingdir ="c:\Users\All Users\Tivoli\adsm.sys"
```

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

Vssusesystemprovider

The `vssusesystemprovider` option specifies whether to use the Windows system provider, or to let Windows decide the most suitable provider to use.

Use the `vssusesystemprovider` option for Microsoft Windows Volume Shadow Copy Service (VSS) operations, such as system state backup or Tivoli Storage Manager for Copy Services backups.

Supported Clients

This option is valid for all 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 options file (dsm.opt).

Syntax



Parameters

Yes

Specifies that the Microsoft Windows VSS system provider is used.

No Specifies that the default system provider is used. This provider might or might not be the same as the system provider, depending on what other providers are installed on the system. Use no if you want to use the default system provider and the default system provider is not the Microsoft Windows VSS provider. No is the default.

Examples

Options file:

```
vssusesystemprovider yes
```

Command line:

Does not apply.

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 Windows 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 443

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 service and the web client agent service are required.

If you do not specify this option, the default value, zero (0), is used for both ports. This causes TCP/IP to randomly assign a free port number for the client acceptor service and the web client agent service.

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.opt`). 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 service port number. The range of values is 1000 through 32767. If a value is not specified, the default, zero (0), causes TCP/IP to randomly assign a free port number.

agentport

Specifies the required web client agent service port number. The range of values is 1000 through 32767. If a value is not specified, the default, zero (0), causes TCP/IP to randomly assign a free port number.

Examples

Options file:

```
webports 2123 2124
```

Command line:

```
webports 2123, 2124
```

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

►►—WILDCARDSareliteral—☐no☐yes—►►

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=c:\important_files.txt`, where `important_files.txt` contains the list of files to process.

`important_files.txt` contains the following list of files:

```
c:\gpfs\home\myfiles\file?9000
c:\gpfs\home\myfiles\?file
c:\gpfs\home\myfiles\**README**version2
c:\gpfs\home\myfiles\ABC?file*
```

If both `WILDCARDSARELITERAL` and `QUOTESARELITERAL` are both set to `YES`, the following backups can be successfully processed:

```
c:\gpfs\home\myfiles\"file?
c:\gpfs\home\myfiles\?file'
c:\gpfs\home\myfiles\**"README Tomorrow"**
c:\gpfs\home\myfiles\file*
```

Related information

For information about the `filelist` option, see “[Filelist](#)” on page 414.

For information about syntax for file specifications, see “[Specifying input strings that contain blank spaces or quotation marks](#)” on page 114.

“[Quotesareliteral](#)” on page 497

Chapter 12. 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 628
- “Enter client command names, options, and parameters” on page 629
- “Wildcard characters” on page 632

The following table provides an alphabetical list of the commands and a brief description.

Table 60. Commands

Command	Description
archive “Archive” on page 633	Archives files from a workstation to Tivoli Storage Manager storage.
archive fastback “Archive FastBack” on page 636	Archives volumes specified by the fbpolycname, fbclientname and fbvolumename options for long term retention.
backup fastback “Backup FastBack” on page 639	Backs up volumes specified by the fbpolycname, fbclientname and fbvolumename options for long term retention.
backup group “Backup Group” on page 642	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 644	Creates an image backup of one or more file systems or logical volumes that you specify.
backup nas “Backup NAS” on page 648	Creates an image backup of one or more file systems belonging to a Network Attached Storage (NAS) file server.
backup systemstate “Backup Systemstate” on page 650	Backs up all startable system state and system services components as one object to provide a consistent point-in-time snapshot of the system state. This command is valid for any supported Windows client.
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 658	Backs up virtual machines specified in the vmlist option.
cancel process “Cancel Process” on page 670	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 670	Displays a list of restartable restore sessions from which you can select one to cancel.

Table 60. Commands (continued)

Command	Description
delete access “Delete Access” on page 671	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 672	Deletes archived files from Tivoli Storage Manager server storage.
delete backup “Delete Backup” on page 674	Deletes active and inactive backup files from Tivoli Storage Manager server storage.
delete filespace “Delete Filespace” on page 677	Deletes file spaces in Tivoli Storage Manager server storage.
delete group “Delete Group” on page 678	Deletes a group backup on the Tivoli Storage Manager server.
expire “Expire” on page 680	Inactivates backup objects that you specify in the file specification or with the <code>filelist</code> option.
help “Help” on page 681	Displays a Table of Contents of help topics for the command-line client.
incremental “Incremental” on page 683	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 690	Starts an interactive command session.
macro “Macro” on page 692	Executes commands within a macro file that you specify.
monitor process “Monitor Process” on page 692	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 693	Simulates an archive command without sending data to the server.
preview backup “Preview Backup” on page 694	Simulates a backup command without sending data to the server.
query access “Query Access” on page 695	Displays a list of current authorization rules.
query adobjects “Query Adobjects” on page 696	Displays a list of current authorization rules.
query archive “Query Archive” on page 697	Displays a list of archived files.
query backup “Query Backup” on page 700	Displays a list of backup versions.
query backupset “Query Backupset” on page 703	Queries a backup set from a local file or the Tivoli Storage Manager server. On those clients that support tape devices, this command can query a backup set from a tape device.
query filespace “Query Filespace” on page 706	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 709	Displays information about group backups and their members.

Table 60. Commands (continued)

Command	Description
query image “Query Image” on page 710	Displays information about image backups.
query inclexcl “Query Inclexcl” on page 712	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 713	Displays information about available management classes.
query node “Query Node” on page 714	Displays all the nodes for which an administrative user ID has authority to perform operations.
query options “Query Options” on page 715	Displays all or part of your options and their current settings.
query restore “Query Restore” on page 716	Displays a list of your restartable restore sessions in the server database.
query schedule “Query Schedule” on page 717	Displays information about scheduled events for your node.
query session “Query Session” on page 717	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 718	Gathers Tivoli Storage Manager system information and outputs this information to a file or the console.
query systemstate “Query Systemstate” on page 720	Displays information about the backup of the system state on the Tivoli Storage Manager server. This command is valid for all supported Windows clients.
query vapp Query Vapp	Determines which VMware vApps are backed up.
query vm “Query VM” on page 725	Verifies the successful backups of the virtual machines from the vStorage backup server.
restart restore “Restart Restore” on page 729	Displays a list of restartable restore sessions from which you can one to restart.
restore “Restore” on page 730	Restores copies of backup versions of your files from a Tivoli Storage Manager server.
restore adobjects “Restore Adobjects” on page 738	Restores individual Active Directory objects from the local Active Directory Deleted Objects container.
restore backupset “Restore Backupset” on page 739	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 746	Restores specific members or all members of a group backup.
restore image “Restore Image” on page 748	Restores a file system or raw volume image backup.
restore nas “Restore NAS” on page 751	Restores the image of a file system belonging to a Network Attached Storage (NAS) file server.
restore systemstate “Restore Systemstate” on page 753	Restores a backup of the system state. This command is valid for all supported Windows clients.
restore vapp Restore Vapp	Restores a single VMware vApp.
restore vm “Restore VM” on page 758	Restores a full VM backup, and returns the full VM backup files to the vbackdir directory on the vStorage backup server.

Table 60. Commands (continued)

Command	Description
retrieve “Retrieve” on page 766	Retrieves copies of archived files from the Tivoli Storage Manager server.
schedule “Schedule” on page 770	Starts the client scheduler on the workstation.
selective “Selective” on page 771	Backs up selected files.
set access “Set Access” on page 775	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 778	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.
set password “Set Password” on page 781	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 11, “Processing options,” on page 295

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 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.

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 321

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` 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 c:\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 `c:\project\budget.fin` to the new path `c:\newproj\newbudg.fin`

```
dsmc restore c:\project\budget.fin c:\newproj\newbudg.fin
```

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 `d:\test` directory:

```
set access backup d:\test\* * *  
set access backup d:\test\*\* * *
```

- 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.
- The length of a file specification is limited.
 - The maximum number of bytes for a file name and file path when combined is 6255. However, the file name itself cannot exceed 255 bytes. Furthermore, directory names (including the directory delimiter) within a path are limited to 255 bytes. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.

When using the open file support feature with VSS, the backup-archive client adds the snapshot volume name to the path of the objects being processed. The resulting path (snapshot volume name plus object path) must adhere to the limits shown. The snapshot volume name can be up to 1024 bytes.

- 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.

```
restore /home/mydir/ /away/yourdir
```

The following example illustrates these two rules. Even though `mydir` and `yourdir` are directories, the command will fail because `*` is implied after `mydir`, and `yourdir` is considered a file.

```
restore c:\home\mydir\ c:\away\yourdir
```

- 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.

For example, if the current working directory is `c:\home\me` and the command is `dsmc res c:\fs\dir1\ mydir\`, the complete restore path is this:

```
c:\home\me\mydir
```

- When a file specification contains spaces, it must be enclosed in quotation marks. For example:

```
dsmc sel "x:\dir one\file1"
```

When a file specification ends with a backslash and is enclosed in quotation marks, an extra backslash (`\`) must be added to the end of the file specification. If an extra backslash is not added, the filespec will not be processed correctly, and the operation might cause unexpected results.

The following example is incorrect:

```
dsmc sel "x:\dir one\"
```

The following example is correct:

```
dsmc sel "x:\dir one\\"
```

Here is an example of restoring the contents of one directory to another, when both directory names contain spaces:

```
dsmc rest "x:\dir one\\" "x:\dir two\\"
```

- Microsoft Dfs volumes are accessed using the standard UNC names. The following are examples of valid syntax to access MS Dfs volumes:

```
\\Server_Name\Dfs_Root_Name\path
\\Fault_Tolerant_Name\Dfs_Root_Name\path
```

Related reference:

"Filelist" on page 414

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 61. Wildcard characters

Pattern	Matches	Does not match
<i>Asterisk (*)</i>		
ab*	ab, abb, abxxx	a, b, aa, bb
ab*rs	abrs, abtrs, abrsrs	ars, aabrs, abrss
ab*ef*rs	abefrs, abefghrs	abefr, abers
abcd.*	abcd.c, abcd.txt	abcd, abcdc, abcdtxt
<i>Question Mark (?)</i>		
ab?	abc	ab, abab, abzzz
ab?rs	abrs	abrs, abllrs
ab?ef?rs	abdefrs	abefrs, abdefrs, abefjrs
ab??rs	abcdrs, abzzrs	abrs, abjrs, abkkrs

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.

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.

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 C:\proposals\drafts\ice.doc C:\proposals\drafts\fire.doc, then C:\proposals and C:\proposals\drafts 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 (C:\).

You can specify as many file specifications as available resources or other operating system limits allow.

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 62. Archive command: Related options

Option	Where to use
archmc	Command line only.
autofsrename	Client options file (dsm.opt) only.
changingretries	Client options file (dsm.opt) or command line.
compressalways	Client options file (dsm.opt) or command line.
compression	Client options file (dsm.opt) or command line.
deletefiles	Command line only.
description	Command line only.
dironly	Command line only.
encryptiontype	Client options file (dsm.opt).
encryptkey	Client options file (dsm.opt).
filelist	Command line only.
filesonly	Command line only.
postsnapshotcmd	Client options file (dsm.opt) or with the include.fs option.

Table 62. Archive command: Related options (continued)

Option	Where to use
preservelastaccessdate	Client options file (dsm.opt) or command line.
presnapshotcmd	Client options file (dsm.opt) or with the include.fs option.
skipntpermissions	Client options file (dsm.opt) or command line.
skipntsecuritycrc	Client options file (dsm.opt) or command line.
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.jan in the c:\plan\proj1 directory.

Command: archive c:\plan\proj1\budget.jan

Task Archive all files in the c:\plan\proj1 directory with a file extension of .txt.

Command: archive c:\plan\proj1*.txt

Task Archive all files in the c:\ drive.

Command: archive -subdir=yes c:*.*

Task Archive all files in the Microsoft Dfs volume, MyDfsVolume. You must specify *subdir=yes* to archive *all* files in the volume.

Command: archive \\myserver\mydfsroot\mydfsvolume*.* -subdir=yes

Task Assuming that you initiated a snapshot of the C:\ drive and mounted the snapshot as \\florence\c\$\snapshots\snapshot.0, archive the c:\dir1\sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name C:.

Command: dsmc archive c:\dir1\sub1* -subdir=yes
-snapshotroot=\\florence\c\$\snapshots\snapshot.0

Related tasks:

“Configuring Open File Support” on page 83

Related reference:

“Include options” on page 430

“Snapshotproviderfs” on page 535

Open file support

If open file support has been configured, the backup-archive client performs a snapshot backup or archive of files that are locked (or “in use”) by other applications.

The snapshot allows the archive to be taken from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to

the file system are not included in the archive. You can set the **snapshotproviderfs** parameter of the `include.fs` option to `none` to specify which drives do not use open file support.

Use VSS as the snapshot provider for open file support.

Note:

1. You can use the `include.fs` option to set snapshot options on a per file system basis.
2. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) that are formatted with NTFS or ReFS file systems. This support includes SAN-attached volumes that meet these requirements.
3. To enable open file support in a cluster environment, all systems in the cluster should have VSS configured.

Archive FastBack

Use the **archive fastback** command to archive Tivoli Storage Manager FastBack volumes specified by the `fbpolicyname`, `fbclientname` and `fbvolumename` options for long-term retention.

Before using this command, configure the client to back up and archive Tivoli Storage Manager FastBack data. Also, before you issue this command, at least one snapshot should exist in the FastBack repository for the FastBack policy being archived or backed up.

If a policy specification contains both Windows and Linux FastBack clients, only the Windows volumes will be backed up or archived to the Tivoli Storage Manager server by the Windows backup-archive client.

You can use Tivoli Storage Manager FastBack options to archive the latest snapshots of the following volumes:

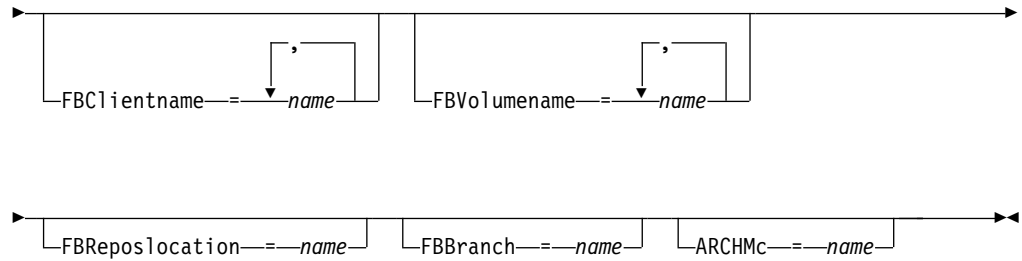
- All clients and volumes associated with a FastBack policy or a list of FastBack policies.
- All volumes associated with a specific FastBack client or a list of FastBack clients for a given FastBack policy.
- A specific volume or volumes associated with a FastBack client for a given FastBack policy.

Supported Clients

This option is valid for all Windows clients that are configured as FastBack dedicated proxies. This command is also valid for Windows clients that are installed on a FastBack server workstation, or the Tivoli Storage Manager FastBack Disaster Recovery (DR) Hub.

Syntax

```
➤➤ ARCHIVE 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 always specify the FBReposLocation option for Linux.

Parameters

Table 63. Archive FastBack command: Related options

Option	Where to use
fbpolicyname "Fbpolicyname" on page 409	Command line and scheduler.
fbserver "Fbserver" on page 411	Command line and scheduler.
fbclientname "Fbclientname" on page 407	Command line and scheduler.
fbvolumename "Fbvolumename" on page 412	Command line and scheduler.
fbreposlocation "Fbreposlocation" on page 410	Command line and scheduler.
fbbranch "Fbbranch" on page 406	Command line and scheduler.
archmc "Archmc" on page 323	Command line and scheduler.

Examples

Command line:

The backup-archive client is installed on a FastBack server. Use this command to archive all FastBack volumes for all Windows FastBack clients that are defined for FastBack policy1:

```
dsmc archive fastback -fbpolicyname=Policy1  
-fbserver=myfbserver
```

The repository location is not required. If you provide the repository location, it is ignored.

The FastBack server name, -myfbserver, is the short host name of the FastBack server where the client is running.

Command line:

The backup-archive client is installed on a FastBack disaster recovery hub. Use this command to archive all FastBack volumes for all FastBack clients that are found in the policy named Policy 1:

```
dsmc archive fastback -fbpolicyname="Policy 1"  
-fbserver=myFbServer -fbbranch=branch1
```

The repository location is not required. If you provide the repository location, it is ignored.

The parameter myFbServer specifies the short host name of the FastBack Server whose FastBack branch is specified using the FBBranch option

Command line:

The backup-archive client is installed on a dedicated proxy machine with Tivoli Storage Manager FastBack administrative command line and FastBack mount. The client is connecting to a FastBack server repository. Use this command to archive all FastBack volumes for all FastBack clients that are found in the policy named Policy 1:

```
dsmc archive fastback -fbpolicyname="Policy 1" -fbserver=myFbServer  
-fbreposlocation=\\myFbServer.company.com\REP
```

The repository location is required.

The short host name of the machine where the FastBack server is installed is myFbServer.

Command line:

The backup-archive client is installed on a dedicated proxy machine with Tivoli Storage Manager FastBack administrative command line and FastBack mount. The client is connecting to a remote branch repository on a FastBack disaster recovery hub. Use this command to archive all FastBack volumes for all FastBack clients that are found in the policy named Policy 1:

```
dsmc archive fastback -fbpolicyname="Policy 1" -fbserver=myFbServer  
-fbreposlocation=\\myfbdrhub.company.com\REP  
-fbbranch=aFbServerBranch
```

The repository location is required.

The myFbServer value specified with the -fbserver option is the short host name of the FastBack Server whose FastBack branch is specified using the FBBranch option.

The fbbranch option specifies the branch ID of the FastBack server on the disaster recovery hub.

Command line:

Archive all volumes protected by FastBack policy named policy1 from the FastBack server named basil, and apply management class "my_tsm_mgmt_class" to the archived volumes.

```
dsmc archive fastback -Fbpolicyname=policy1
-FBServer=basil -ARCHMC="my_tsm_mgmt_class"
```

Related concepts:

"Configuring the client to back up and archive Tivoli Storage Manager FastBack data" on page 67

Backup FastBack

Use the **backup fastback** command to back up Tivoli Storage Manager FastBack volumes specified by the fbpolycname, fbclientname and fbvolumename options for long-term retention.

Before using this command, configure the client to back up and archive Tivoli Storage Manager FastBack data. Also, before you issue this command, at least one snapshot should exist in the Tivoli Storage Manager FastBack repository for the Tivoli Storage Manager FastBack policy being archived or backed up.


If a policy specification contains both Windows and Linux FastBack clients, only the Windows volumes will be backed up or archived to the Tivoli Storage Manager server by the Windows 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 all Windows clients that are configured as Tivoli Storage Manager FastBack dedicated proxies. This command is also valid for Windows clients that are installed on a Tivoli Storage Manager FastBack server workstation, or a Tivoli Storage Manager FastBack Disaster Recovery (DR) Hub.

Syntax



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.

Table 64. Backup FastBack command: Related options

Option	Where to use
fbpolicyname "Fbpolicyname" on page 409	Command line and scheduler.
fbserver "Fbserver" on page 411	Command line and scheduler.
fbclientname "Fbclientname" on page 407	Command line and scheduler.
fbvolumename "Fbvolumename" on page 412	Command line and scheduler.
fbreposlocation "Fbreposlocation" on page 410	Command line and scheduler.
fbbranch "Fbbranch" on page 406	Command line and scheduler.
backmc "Backmc" on page 335	Command line and scheduler.

Examples

Command line:

The backup-archive client is installed on a FastBack server. Use this command to back up all Tivoli Storage Manager FastBack volumes for all Windows FastBack clients that are defined for Tivoli Storage Manager FastBack policy1:

```
dsmc backup fastback -fbpolicyname=Policy1
-fbserver=myfbserver
```

The repository location is not required. If you provide the repository location, it is ignored.

The FastBack server name, -myfbserver, is the short host name of the FastBack server where the client is running.

Command line:

The backup-archive client is installed on a FastBack disaster recovery hub. Use this command to back up all FastBack volumes for all FastBack clients that are found in the policy named Policy 1:

```
dsmc backup fastback -fbpolicyname="Policy 1"  
-fbserver=myFbServer -fbbranch=branch1
```

The repository location is not required. If you provide the repository location, it is ignored.

The FastBack server name, myFbServer, is the short host name of the FastBack server whose FastBack branch is specified using the FBBranch option

Command line:

The backup-archive client is installed on a dedicated proxy machine with FastBack administrative command line and FastBack mount. The client is connecting to a FastBack server repository. Use this command to back up all FastBack volumes for all FastBack clients that are found in the policy named Policy 1:

```
dsmc backup fastback -fbpolicyname="Policy 1" -fbserver=myFbServer  
-fbreposlocation=\\myFbServer.company.com\REP
```

The repository location is required.

The short host name of the machine where the FastBack server is installed is myFbServer.

Command line:

The backup-archive client is installed on a dedicated proxy machine with FastBack administrative command line and FastBack mount. The client is connecting to a remote branch repository on a FastBack disaster recovery hub. Use this command to back up all FastBack volumes for all FastBack clients that are found in the policy named Policy 1:

```
dsmc backup fastback -fbpolicyname="Policy 1" -fbserver=myFbServer  
-fbreposlocation=\\myfbdrhub.company.com\REP  
-fbbranch=aFbServerBranch
```

The repository location is required.

The myFbServer value specified with the -fbserver option is the short host name of the FastBack Server whose FastBack branch is specified using the FBBranch option.

The fbbranch option specifies the branch ID of the FastBack server on the disaster recovery hub.

Command line:

Back up all volumes protected by FastBack policy named policy1 from the FastBack server named basil, and apply the management class "my_tsm_mgmt_class" to the backed up volumes:

```
dsmc backup fastback -Fbpolicyname=policy1  
-FBServer=basil -BACKMC="my_tsm_mgmt_class"
```

Related concepts:

"Configuring the client to back up and archive Tivoli Storage Manager FastBack data" on page 67

Backup Group

Use the **backup group** command to create and back up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.

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

Syntax

►►Backup GRoup— *options* —————►►

Parameters

Table 65. Backup Group command: Related options

Option	Where to use
filelist "Filelist" on page 414	Command line only.
groupname "Groupname" on page 422	Command line only.
mode "Mode" on page 463	Command line only.
snapshotproviderfs "Snapshotproviderfs" on page 535	Client options file (dsm.opt) or with the <code>include.fs</code> option.
snapshotproviderimage "Snapshotproviderimage" on page 536	Client options file (dsm.opt) or with <code>include.image</code> option.
virtualfsname "Virtualfsname" on page 575	Command line only.

Examples

Task Perform a full backup of all the files in the `c:\dir1\filelist1` file to the virtual file space name `\virtfs` containing the group leader `group1` file.

Command:

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

Related information

“Include options” on page 430

“Query Group” on page 709

“Restore Group” on page 746

“Delete Group” on page 678

“Query Filespace” on page 706

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.

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.

Note:

1. The account that is running the backup-archive client must have administrator authority to successfully perform any type of image backup.
2. The API must be installed to use the **backup image** command.


Use the **include.image** option to include a file system or logical volume for image backup, or to specify volume-specific options for image backup.

The **backup image** command uses the **compression** option.

Supported Clients

This command is valid for all Windows platforms.

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.

Image backup is only supported on a volume that has a mount or a drive letter assigned to it. A volume without a drive letter or mount point cannot be backed up.

Table 66. Backup Image command: Related options

Option	Where to use
compressalways "Compressalways" on page 350	Client options file (dsm.opt) or command line.
compression "Compression" on page 351	Client options file or command line.
imagegapsize "Imagegapsize" on page 427	Use with the backup image command, the include.image option, or in the options file.
mode "Mode" on page 463	Command line only.
postsnapshotcmd "Postsnapshotcmd" on page 486	Use with the backup image command, the include.image option, or in the options file.
presnapshotcmd "Presnapshotcmd" on page 492	Use with the backup image command, the include.image option, or in the options file.
snapshotproviderimage "Snapshotproviderimage" on page 536	Client options file or with include.image option.

Examples

Task Back up a volume that has no drive letter but is mounted as a mount point.

```
dsmc backup image m:\mnt\myntfs
```

Task Back up the h drive by using an image incremental backup. An image incremental backup backs up files that are new or changed since the last full image backup.

```
dsmc backup image h: -mode=incremental
```

Task Perform an offline image backup of the f drive.

```
dsmc backup image f: -snapshotproviderimage=none
```

Task Perform an online image backup of the f drive.

```
dsmc backup image f: -snapshotproviderimage=VSS
```

Task Back up the f drive, which is mapped to a volume that has not been formatted with a file system.

```
dsmc backup image f:
```

Related information

"Imagegapsize" on page 427

"Snapshotproviderimage" on page 536

"Configuring Open File Support" on page 83

"Image backup" on page 153

"Mode" on page 463

"Comparing methods 1 and 2" on page 157 To decide which method is appropriate for your environment.

Offline and online image backup

The traditional offline image backup prevents write access to the volume by other system applications during the operation.

If open file support has been configured, the backup-archive client performs a snapshot backup or archive of files that are locked (or "in use") by other applications.

Use VSS as the snapshot provider for open file support.

The following considerations apply to offline and online image backups:

- If you create an image of the system drive, you cannot restore it to the original location. Restore of any image requires that the client have an exclusive lock of the volume you are restoring to, so the system drive cannot be restored since the client is unable to lock the system drive. You can restore an image backup of the system drive to an alternate location.
- Because of different system component configurations, the system image not be consistent across components (such as Active Directory). Some of these components can be configured to use different volumes where parts are installed on the system drive and others to non-system volumes.
- Install the Tivoli Storage Manager client program on the system drive. The client cannot restore an image to the same volume where the client program is installed.

- Image backup is only supported on volumes that have a mount point or a drive letter assigned. the client will not back up a volume without a mount point or drive letter.
- If bad disk sectors are detected on the source drive during a LAN-free or LAN-based image backup, data corruption occur. In this case, bad sectors are skipped when sending image data to the Tivoli Storage Manager server. If bad disk sectors are detected during the image backup, a warning message is issued after the image backup completes.

Utilizing image backup to perform file system incremental backup

There are two methods of utilizing image backups to perform efficient incremental backups of your file system. These backup methods allow you to perform point-in-time restore of your file systems and improve backup and restore performance.

You can perform the backup only on formatted volumes; not on raw logical volumes. You can 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 h:`
- To perform an image backup of the same file system: `dsmc backup image h:`
- To periodically perform incremental backups: `dsmc incremental h:`

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 h: -incremental -deletefiles`.

During the restore, the client does the following:

- Restores the most recent image on the server.
- Deletes all of the files restored in the previous step which are inactive on the server. These are files which existed at the time of the image backup, but were subsequently deleted and recorded by a later incremental backup.
- Restores new and changed files from the incremental backups.

If you do not follow the steps exactly, two things can occur:

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 h:`
- To perform an incremental image backup of the file system: `dsmc backup image h: -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 h:`

- To restore the image: `dsmc restore image h: -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.

NAS file system designations on the command line require brace delimiters {} around the file system names, such as: {/vol/vol0}.

Supported Clients

This command is valid for all Windows clients.

Syntax



Parameters

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 67. Backup NAS command: Related options

Option	Where to use
mode "Mode" on page 463	Command line only.
monitor "Monitor" on page 468	Command line only.
nasnodename "Nasnodename" on page 471	Client options file (dsm.opt) or command line.
toc "Toc" on page 563	Command line or with the <i>include.fs.nas</i> option in your client options file (dsm.opt).

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=netapps {/vol/vol0}
-toc=yes

Related information

"Nasnodename" on page 471

"Toc" on page 563

"Mode" on page 463

"Monitor" on page 468

"Cancel Process" on page 670

"Domain.nas" on page 375

Backup Systemstate

Use the **backup systemstate** command to back up all bootable system state and system services components as a single object, to provide a consistent point-in-time snapshot of the system state.

Bootable system state components can include the following:

- Active Directory (domain controller only)
- System volume (domain controller only)
- Certificate Server Database
- COM+ database

- Windows Registry
- System and boot files
- ASR writer

System services components can include the following:

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

The list of bootable system state and system services components is dynamic and can change depending on service pack and operating system features installed. The backup-archive client allows for the dynamic discovery and backup of these components.

System state is represented by several VSS writers of type "bootable system state" and "system service". Of these, the System Writer is the largest part of the system state in terms of number of files and size of data. By default, the System Writer backup is incremental. You can use the `systemstatebackupmethod` option to perform full backups of the System Writer. For more information, about this option, see "Systemstatebackupmethod" on page 552. The client always backs up all other writers in full.

This command also backs up ASR data for Windows clients; BIOS and UEFI boot architectures are supported.

Note:

1. The system and boot files component of system state is backed up only if a member (file) of that component has changed since the last backup. If a member changes, the entire group of files that comprise that component are backed up.
2. The backup-archive client on Windows does not allow the backup of any individual component.
3. By default, system state backups are bound to the default management class. To bind them to a different management class, use the `include.systemstate` option; specify **all** as the pattern, and specify the name of the new management class.
For example: `include.systemstate ALL BASVT2`.
4. Use the **query systemstate** command to display information about a backup of the system state on the Tivoli Storage Manager server.
5. Use the **restore systemstate** command to restore the backup of the system state from the Tivoli Storage Manager server.

Supported Clients

This command is valid for all supported Windows clients.

Syntax

►►Backup SYSTEMState◄◄

Parameters

There are no parameters for this command.

Examples

Task Back up the system state.

Command: backup systemstate

Related information

“Preparation for Automated System Recovery” on page 151

“Query Systemstate” on page 720

“Restore Systemstate” on page 753

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 Windows 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 68. Backup VAPP command: Related options

Option	Where to use
asnodename	Client user-options file (dsm.opt), on the command line, or on the General tab in the Preferences editor.
include.vapp	Client options file.

Table 68. Backup VAPP command: Related options (continued)

Option	Where to use
mode	Client user-options file (dsm.opt), or on the command line. Specify IFFull or IFIncremental. The default for vApp backups is IFIncremental.
vappmc	Client user-options file (dsm.opt), 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 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 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 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 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 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 on the command line. This option can also be included on the server in a client option set.
vmctlmc	Client user-options file (dsm.opt), or on the command line. This option can also be included on the server in a client option set.

Table 68. Backup VAPP command: Related options (continued)

Option	Where to use
vmvstortransport	Client user-options file (dsm.opt), 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 324
- “Mode” on page 463
- “Vmbackuptype” on page 581
- “Vmchost” on page 583
- “Vmcpw” on page 584
- “Vmcuser” on page 586

"Vmlimitperdatastore" on page 596

"Vmlimitperhost" on page 597

"Vmmaxparallel" on page 598

"Vmmc" on page 601

"Vmctlmc" on page 585

"Vmvstortransport" on page 618

Backup VM

Use the **backup vm** command to start a full backup of a virtual machine.

The Backup VM command can be used to back up both Microsoft Hyper-V virtual machines and VMware virtual machines. The information for each hypervisor is presented in its own heading. If you are backing up a virtual machine that is part of a Hyper-V setup, you do not need to read the *Backing up VMware virtual machines* text. If you are backing up a VMware virtual machine, you do not need to read the *Backing up Hyper-V virtual machines* text.

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 167.

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.

For more information about compression, see “Compression, encryption, and adaptive subfile backup processing” on page 435.

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 on supported Windows 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 can be the DNS host name or 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.

If you do not specify *vmname*, you can identify the virtual machine with one of the following options:

- `domain.vmfull`
- `domain.vmfile`

: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 381, “Exclude.vmdisk” on page 405, “Include.vmdisk” on page 440.

-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: <ul style="list-style-type: none"> 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 a file-level backup of a VM named vm1.

```
dsmc backup vm vm1 -vmbackuptype=file
```

Perform a file-level virtual machine backup of vm1.example.com by using the VMware VirtualCenter machine named virtctr.example.com.

```
dsmc backup vm vm1 -vmbackuptype=file -vmchost=virtctr
```

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 725
- “**Restore VM**” on page 758
- “Domain.vmfile” on page 378
- “Domain.vmfull” on page 381
- “Include.vm” on page 438
- “Mbobjrefreshthresh” on page 460
- “Mbpctrefreshthresh” on page 461
- “Mode” on page 463
- “Vmbackdir” on page 578
- “Vmbacknodelete” on page 579
- “Vappmc” on page 571
- “Vmbackuptype” on page 581
- “Vmchost” on page 583
- “Vmctlmc” on page 585
- “Vmcpw” on page 584
- “Vmcuser” on page 586
- “Vmlimitperdatastore” on page 596
- “Vmlimitperhost” on page 597
- “Vmmc” on page 601
- “Vmmaxparallel” on page 598
- “Vmtagdatamover” on page 612
- “**Set Vmtags**” on page 787
- “Virtual machine exclude options” on page 404
- “Virtual machine include options” on page 437

Backing up Microsoft Hyper-V virtual machines

Use the **backup vm** command to back up Hyper-V virtual machines. You can back up Hyper-V guests that exist on a local disk, a SAN-attached disk, a cluster shared volume, or guests that exist on a remote file server share. Remote file server shares must be on a Windows Server 2012 (or newer) system.

You specify the backup mode to use when you want to back up a virtual machine by adding the **-mode** parameter on the command line, or you can set the mode option in the client options file. Any of the following modes can be specified:

Full Full mode. In this mode, an image backup is created of all objects on a virtual machine’s disks, including configuration information, snapshots, and all of the disks. If you are not licensed to use Tivoli Storage Manager for Virtual Environments, you must specify a full VM backup.

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. The backup includes configuration information, and all of the disks. You must have a license for Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V to use this mode. This mode can be used only on Windows clients on Hyper-V hosts that are running in Windows Server 2012 or 2012 R2 environments.

IFIncremental

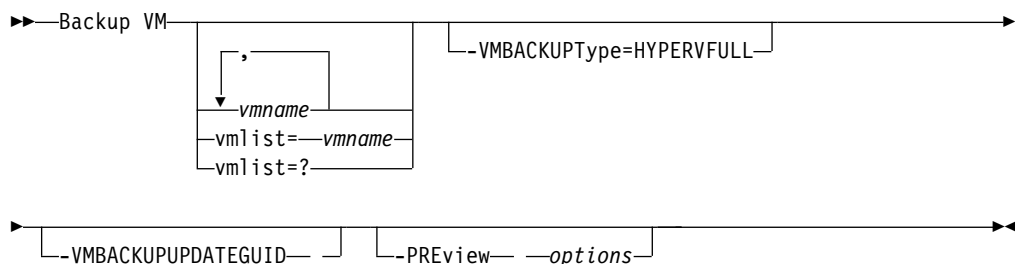
Incremental-forever-incremental. In this mode, a snapshot is created of the blocks that were changed since the last backup. The backup includes configuration information, and all of the disks. You must have a license for Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V to use this mode. This mode can be used only on Windows clients on Hyper-V hosts that are running in a Windows Server 2012 or 2012 R2 environments.

For information about the incremental-forever backup strategy, see Incremental-forever backup strategy.

Supported Clients

This command is valid on supported Windows clients that are installed on a Microsoft Hyper-V host server.

Syntax



Parameters

vmname

Specify the name of the virtual machine that you want to back up; the name is case sensitive. If you specify multiple virtual machine names, separate the names with commas.

Note: 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 must be unique. To prevent errors during processing, ensure that 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

If you do not specify a virtual machine name, and if you specify `-mode=ifincremental` or `-mode=iffull`, the `domain.vmfull` option is used to determine which virtual machines to include in the backup operation.

- 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 a virtual machine name, and if you specify `-mode=full`, the `vmist` option is used to determine which virtual machines to include in the backup operation.

vmist=vmname_1,vmname_2,vmname_n

Specifies a list of virtual machine names to back up. Separate the virtual machine names with commas. Wildcard characters can be used in the virtual machine name to include virtual machines that match the wildcard pattern. The `vmname` and `vmist=` syntaxes are mutually exclusive; you can specify a virtual machine name or a list of virtual machine names, but not both.

vmist=?

The behavior of this syntax varies, depending on which backup mode is specified on the `mode` option:

- If the `mode` option is set to `MODE=FULL`, this syntax returns a list of virtual machines that are defined to the Hyper-V host. An error occurs if you specify an asterisk (*), instead of a question mark.
- If the `mode` option is set to `MODE=IFFULL` or `MODE=IFINCREMENTAL`, the question mark (?) in this syntax is interpreted as a wildcard, and the backup operation is performed for Hyper-V virtual machines that match the wildcard. That is, any Hyper-V virtual machines that have a one-character name are backed up.

-VMBACKUPType=HYPERVFULL

This parameter indicates that the backup operation that is being performed is for a Microsoft Hyper-V virtual machine. If you enable the Hyper-V server role, **HYPERVFULL** is the default backup type.

-VMBACKUPUPDATEGUID

You must be licensed to use Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V to use this option.

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 causes this command to display additional information about a virtual machine, including the virtual disks in the virtual machine. This option applies only to Windows Server 2012, and subsequent versions of Windows Server operating systems. If you specify this option, the output is displayed, but the backup operation does not start.

The following shows example output when you specify the **-preview** option. This example output is abbreviated to show only information about one virtual machine and one disk.

```
tsm> backup vm -vmbackuptype=hypervfull -mode=ifincr -preview
Backup VM command started. Total number of virtual machines to process: 6
```

1. VM Name: VM2

```
Domain Keyword:      all-vm
Mode:                Incremental Forever - Incremental
Target Node Name:    SPC
Data Mover Node Name: SPC
Cluster Resource:    no
```

```
Disk[1]
Name: \\beta\d$\Hyper-V\VM2\VM2.vhdx
Capacity:            40.00 GB
Size:                9.13 GB
Full Backup:         included
Incremental Backup:  excluded
Disk Type:           VHDX
Number of Subdisk:   1
```

Use both the **-preview** option and the **-detail** option on the command to display information about subdisks that are included when the backup is run. A subdisk is the AVHDX file that is created when a snapshot is taken of a VHDX file. This example output is abbreviated to show only information about one virtual machine and one disk.

```
tsm> backup vm -vmbackuptype=hypervfull -mode=ifincr -preview -detail
Backup VM command started. Total number of virtual machines to process: 6
```

1. VM Name: VM2

```
Domain Keyword:      all-vm
Mode:                Incremental Forever - Incremental
Target Node Name:    SPC
Data Mover Node Name: SPC
Cluster Resource:    no
```

```
Disk[1]
Name: \\beta\d$\Hyper-V\VM2\VM2.vhdx
Capacity:            40.00 GB
Size:                9.13 GB
Full Backup:         included
Incremental Backup:  excluded
```



```

Disk Type:          VHDX
Number of Subdisk:  1

Subdisk[1]
Name: \\beta\d$\Hyper-V\VM2\VM2_88ABC7B4-E59F-87B9-20D064E549C4.avhdx
Capacity:          40.00 GB
Size:              4.00 MB
Full Backup:       included
Incremental Backup: included
Disk Type:         AVHDX

```

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: <ul style="list-style-type: none"> 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.

Microsoft Hyper-V backup examples

Start a full VM backup of a Hyper-V virtual machine, by specifying the virtual machine name "VM1".

```
dsmc backup vm VM1 -vmbackuptype=hypervfull -mode=full
```

Start an incremental-forever-incremental backup of a Hyper-V virtual machine that is named "VM1".

```
dsmc backup vm VM1 -vmbackuptype=hypervfull -mode=ifincremental
```

Start an incremental-forever-incremental backup of all Hyper-V virtual machines that are specified on the domain.vmfull option.

```
dsmc backup vm -vmbackuptype=hypervfull -mode=ifincremental
```

Use the **vm1ist=** keyword to back up two Hyper-V virtual machines, named vm1 and vm5. This command example applies only to backup-archive client configurations where the Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V licensed product is not installed.

```
dsmc backup vm -vmbackuptype=hypervfull -mode=full -vm1ist=vm1,vm5
```

List all virtual machines that are defined to the Hyper-V host. This command example applies only to backup-archive client configurations where the Tivoli Storage Manager for Virtual Environments: Data Protection for Microsoft Hyper-V licensed product is not installed.

```
dsmc backup vm -mode=full -vm1ist=?
```

Hyper-V options file examples

This example specifies individual virtual machines in the client options file, and shows the `backuptype` and `vm1ist` options. The preferred way to include Hyper-V virtual machines in backup operations is to use the `domain.vmfull` option, instead of the `vm1ist` option.

```
backuptype hypervfull
vm1ist vm1,vm2,vm5
```

Related links for backing up Hyper-V virtual machines

- “Detail” on page 363
- “Domain.vmfull” on page 381
- “Mbojrefreshthresh” on page 460
- “Mbpctrefreshthresh” on page 461
- “Mode” on page 463
- “Query VM” on page 725
- “Restore VM” on page 758
- “Vmbackdir” on page 578
- “Vmbackuptype” on page 581
- “Vm1ist” on page 598

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 all Windows clients.

Syntax

►►—Cancel Process—►►

Parameters

There are no parameters for this command.

Examples

Task Cancel current NAS image backup or restore processes.

Command: `cancel process`

Cancel Restore

The **cancel restore** command displays a list of your restartable restore sessions in the server database.

You can 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.
- You want to cancel restartable restore sessions.
- 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	daisy	c:\dev\proja\list.c
2	Archive	node3	marm	c:\fin\budg\depta.jan
3	Backup	node4	susie	c:\plan\exp\deptc.feb
4	Archive	node5	susies	c:\mfg\invn\parta.wip

Enter Index of rule(s) to delete, or quit to cancel:

To delete the authorization rules that allow marm and susies to access your files, type 2 4 or 2,4, then press Enter.

Delete Archive

The **delete archive** command deletes archived files from Tivoli Storage Manager server storage. Your administrator must give you 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] [—{filespace—}{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*, do not include a drive letter in the file specification.

{filespace}

Specifies the file space (enclosed in braces) on the server that contains the file you want to delete. This is the name on the workstation drive from which the file was archived.

Use the *filespace* 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.

You can specify a UNC name; drive label names are only used for removable media.

You must specify a mixed or lowercase NTFS or ReFS file space name that is enclosed in quotation marks and braces. For example, {"NTFSDrive"}. Single quotation marks are valid in loop mode. For example, {"NTFSDrive"} and

{'NTFSDrive'} are both valid. In batch mode, only single quotation marks are valid. The single quotation mark requirement is a restriction of the operating system.

Table 69. Delete Archive command: Related options

Option	Where to use
dateformat "Dateformat" on page 356	Client options file (dsm.opt) or command line.
description "Description" on page 362	Command line only.
filelist "Filelist" on page 414	Command line only.
noprompt "Noprompt" on page 474	Command line only.
numberformat "Numberformat" on page 476	Client options file (dsm.opt) or command line.
pick "Pick" on page 481	Command line only.
subdir "Subdir" on page 548	Client options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 553	Client options file (dsm.opt) or command line.
timeformat "Timeformat" on page 562	Client options file (dsm.opt) or command line.

Examples

Task Delete files from file space abc in the proj directory.

```
dsmc delete archive {"abc"}\proj\*
```

Task Delete a file that is named budget.

```
dsmc delete archive c:\plan\proj1\budget.jan
```

Task Delete all files that are archived from the c:\plan\proj1 directory with a file extension of .txt.

```
delete archive c:\plan\proj1\*.txt
```

Task Delete files that are archived from the c:\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 c:\project\* -pick
```

Task Delete selected files from the group of files that are archived with the description "Monthly Budgets 2013" located in c:\projects and its subdirectories.

```
dsmc delete ar c:\projects\* -description="Monthly Budgets 2013"
-pick -subdir=yes
```

Related information

"Filelist" on page 414

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.

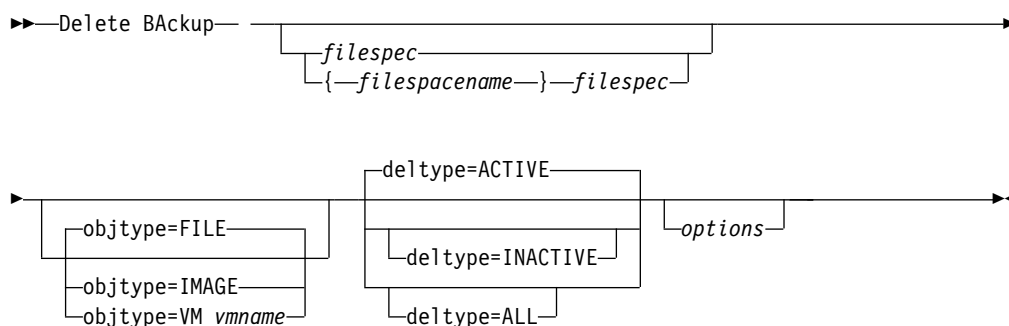
When you delete files, the Tivoli Storage Manager server takes all of the backed up files that meet the *filespec* and *deltype* options that are specified and deactivates them. The server also assigns a deactivation date of *infinite-minus* so that the files are no longer available for restore and are 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 deltype=ALL, you cannot use the pick option because deltype=ALL and the pick option are mutually exclusive.

Table 70. Delete Backup command: Related options

Option	Where to use
description "Description" on page 362	Command line only.
filelist "Filelist" on page 414	Command line only.
fromdate "Fromdate" on page 420	Command line, and in the GUI find function.
fromtime "Fromtime" on page 421	Command line, and in the GUI find function.
noprompt "Noprompt" on page 474	Command line only.
pick "Pick" on page 481	Command line only.
pitdate "Pitdate" on page 482	Command line, and in the GUI find function.
pittime "Pittime" on page 483	Command line, and in the GUI find function.
subdir "Subdir" on page 548	Client options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 553	Client options file (dsm.opt) or command line.
timeformat "Timeformat" on page 562	Client options file (dsm.opt) or command line.
today "Today" on page 564	Command line, and in the GUI find function.
totime "Totime" on page 565	Command line, and in the GUI find function.

Examples

Task Delete all active file objects from file space abc in the proj directory.

Command: delete backup {abc}\proj*

Task Delete all inactive files with a name that ends with .txt that were backed up from the c:\plan\proj1 directory, and its subdirectories.

Command: delete backup c:\plan\proj1*.txt -deltype=inactive -subdir=yes

Task Delete selected active files that are backed up from the c:\project directory. Use the -pick option to display a list of backup copies that match the file specification. From the list, you can select which versions to delete.

Command: delete backup c:\project* -pick

Task Delete all active and inactive versions of files and subdirectories in c:\user\myproject.

Command: delete backup c:\user\myproject* -deltype=all

Note: The backup versions of directory object `c:\user\myproject` are not deleted.

Task Delete the active backup of a virtual machine that is named `vm1`.

Command: `delete backup -objtype=vm vm1`

Note: If one or more inactive versions of this backup exist, the most recent becomes the active version.

Task Delete one or more backup versions of a virtual machine that is named `vm_test`.

Command: `delete backup -objtype=vm -inactive vm_test`

Note: All versions of backups for this VM node are displayed in a list; you select the versions to delete.

Related reference:

“Filelist” on page 414

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.

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 UNC 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 71. Delete Filespace command: Related options

Option	Where to use
class "Class" on page 341	Command line only.
detail "Detail" on page 363	Command line only.
nasnodename "Nasnodename" on page 471	Client options file or command line.
scrolllines "Scrolllines" on page 521	Client options file or command line.
scrollprompt "Scrollprompt" on page 522	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 471

"Class" on page 341

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 Windows clients.

Syntax



Parameters

filespec

Specifies the virtual file space name and the group name that you want to delete from the server storage.

Table 72. Delete Group command: Related options

Option	Where to use
inactive “Inactive” on page 428	Command line only.
noprompt “Noprompt” on page 474	Command line only.
pick “Pick” on page 481	Command line only.
pitdate “Pitdate” on page 482	Command line only.
pittime “Pittime” on page 483	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 428

“Pick” on page 481

“Noprompt” on page 474

“Query Filespace” on page 706

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 service.

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 73. Expire command: Related options

Option	Where to use
dateformat "Dateformat" on page 356	Client options file (dsm.opt) or command line.
filelist "Filelist" on page 414	Command line only.
noprompt "Noprompt" on page 474	Command line only.
numberformat "Numberformat" on page 476	Client options file (dsm.opt) or command line.
pick "Pick" on page 481	Command line only.
timeformat "Timeformat" on page 562	Client options file (dsm.opt) or command line.

Examples

Task Deactivate the letter1.txt file in the home directory.

Command: expire c:\home\letter1.txt

Task Deactivate all files in the admin\mydir directory.

Command: expire c:\admin\mydir*

Task Deactivate all files that are named in the c:\avi\filelist.txt file.

Command: expire -filelist=c:\avi\filelist.txt

Task Deactivate the current backup of the virtual machine that is named vm_test.

Command: expire -objtype=VM vm_test

Help

Use the **help** command to display information about commands, options, and messages.

Tip: If you use the **help** command on the initial command line, no server contact is made and no password is needed.

Supported Clients

This command is valid for all clients.

Syntax

»» Help	command-name [subcommand-name]	««
	option-name	
	TOC-section-number	
	[ANS]message-number	

Entering the **help** command with no arguments causes help to display the complete table of contents. Either with the initial command or when HELP

displays a prompt, you can enter the following parameters.

Parameters

command-name [subcommand-name]

Specifies a command name and, optionally, a subcommand name or their abbreviation, for example: **backup image**, or **b i**. In 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

Command: dsmc help ANS1036

Incremental

The **incremental** command backs up all new or changed data in the locations that you specify, unless you exclude them from backup services.

You can back up all new or changed files or directories in the default client domain or from file systems, directories, or files.

To incrementally back up selected files or directories, enter a file specification in the command. If you do not enter a file specification, the default is to back up files or directories in the default domain.

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 322.

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 or the journal database. 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 wildcards to select a group of files or all the files in a directory. You can specify as many file specifications as available resources or other operating system limits permit. Separate file specifications with a space. You can also use the **filelist** option to process a list of files. The backup-archive client opens the file that you specify with this option and processes the list of files within according to the specific command. If you do not specify a file specification, the **domain** option determines what to backup.

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 74. Incremental command: Related options

Option	Where to use
absolute "Absolute" on page 322	Command line only.
autofsrename "Autofsrename" on page 333	Client options file (dsm.opt) only.

Table 74. Incremental command: Related options (continued)

Option	Where to use
changingretries “Changingretries” on page 340	Client options file (dsm.opt) or command line.
compressalways “Compressalways” on page 350	Client options file (dsm.opt) or command line.
compression “Compression” on page 351	Client options file (dsm.opt) or command line.
detail “Detail” on page 363	Command line only.
diffsnapshot “Diffsnapshot” on page 364	Command line only.
dirsonly “Dirsonly” on page 368	Command line only.
domain “Domain” on page 371	Client options file (dsm.opt) or command line.
encryptiontype “Encryptiontype” on page 393	Client options file (dsm.opt).
encryptkey “Encryptkey” on page 394	Client options file (dsm.opt).
filelist “Filelist” on page 414	Command line only.
filesonly “Filesonly” on page 418	Command line only.
incrbydate “Incrbydate” on page 445	Command line only.
memoryefficientbackup “Memoryefficientbackup” on page 462	Client user-options file (dsm.opt), server, or command line.
nojournal “Nojournal” on page 474	Command line only.
postsnapshotcmd “Postsnapshotcmd” on page 486	Client options file (dsm.opt) or with include.fs option.
preservelastaccessdate “Preservelastaccessdate” on page 489	Client options file (dsm.opt) or command line.
presnapshotcmd “Presnapshotcmd” on page 492	Client options file (dsm.opt) or with include.fs option.
resetarchiveattribute “Resetarchiveattribute” on page 507	Client options file (dsm.opt).
skipntpermissions “Skipntpermissions” on page 528	Client options file (dsm.opt) or command line.
skipntsecuritycrc “Skipntsecuritycrc” on page 528	Client options file dsm.opt) or command line.
snapdiff “Snapdiff” on page 529	Command line only.
snapshotproviderfs “Snapshotproviderfs” on page 535	System-options file (dsm.sys) within a server stanza or with the include.fs option.
snapshotproviderimage “Snapshotproviderimage” on page 536	Client options file (dsm.opt) or with the include.image option.
snapshotroot “Snapshotroot” on page 537	Command line only.
subdir “Subdir” on page 548	Client options file (dsm.opt) or command line.
tapeprompt “Tapeprompt” on page 553	Client options file (dsm.opt) or command line.

Examples

Task Run an incremental backup of the default client domain that is specified in your client options file (dsm.opt).

```
Incremental
```

Run an incremental backup of the domain that is specified in your client user options file. Adding the **-absolute** option forces a backup of all files in the domain, even if they were not changed since the last incremental backup.

```
Incremental -absolute
```

Task Run an incremental backup of the C, D, and E drives.

```
incremental c: d: e:
```

Task Run an incremental backup of the \home\ngai directory and its contents on the current drive.

```
i \home\ngai\
```

Task Assuming that you initiated a snapshot of the C drive and mounted the snapshot as \\florence\c\$\snapshots\snapshot.0, run an incremental backup of all files and directories under the local snapshot and manage them on the Tivoli Storage Manager server under the C:\ drive file space name.

```
dsmc inc c: -snapshotroot=\\florence\c$\snapshots\snapshot.0
```

Task Run a **snappdiff** incremental backup from a snapshot taken of a network share //homestore.example.com/vol1 mounted on drive H, where homestore.example.com is a file server.

```
incremental -snappdiff H:
```

Task Run a **snappdiff** incremental backup from a snapshot taken of a network share //homestore.example.com/vol1 mounted on drive H, where homestore.example.com is a file server. The **-diffsnapshot** option value of LATEST means that the operation uses the latest snapshot (the active snapshot) for volume H.

```
incremental -snappdiff H: -diffsnapshot=LATEST
```

Related information

“Absolute” on page 322

“Journal-based backup” on page 687

“Selective” on page 771

“Include options” on page 430

“Incrthreshold” on page 446

Chapter 11, “Processing options,” on page 295 for adaptive subfile backups.

Open file support

If open file support has been configured, the backup-archive performs a snapshot backup or archive of files that are locked (or “in use”) by other applications.

Use VSS as the snapshot provider; set **snapshotproviderimage** or **snapshotproviderfs** to VSS.

Note:

1. You can use the `include.fs` option to set snapshot options on a per file system basis.
2. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with NTFS file systems. This support includes SAN-attached volumes that meet these requirements.
3. If the client is unable to create a snapshot, failover to non-OFS backup occurs; the same backup support that would be done if the OFS feature was not configured.
4. To enable open file support in a cluster environment all systems in the cluster should have the OFS feature configured.

Journal-based backup

If the journal engine service is installed and running, then by default the **incremental** command performs a journal-based backup on file systems that are being monitored by the journal engine service.

The backup-archive client does not use the journaling facility inherent in Windows NTFS or ReFS file systems or any other journaled file system.

The journal engine service records changes to an object or its attributes in a journal database. During a journal-based backup, the client obtains a list of files that are eligible for backup from the journal database. Performing backups regularly maintains the size of the journal.

Journal-based backup can increase backup performance. With journal-based backup, the client does not scan the local file system or obtain information from the server to determine which files to process. Journal-based backup also reduces network traffic between the client and server.

The client filters the list by using the current include-exclude list. Tivoli Storage Manager processes, expires, and updates the resulting files according to policy constraints, such as serialization. The management-class copy frequency attribute is ignored during journal-based backup.

The journal engine service excludes specific system files (pagefile, registry, and so on) 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 service configuration file `tsmjbbd.ini`, which is in the backup-archive client installation directory, for specific system files that are excluded.

To support journal-based backup, you must install the journaling engine service. Install this service by using the **dsmcutil** command or the GUI Setup wizard.

If the file specification on the **incremental** command is a file space, the client processes any journal entries for that file space. The client processes directories and file specifications that contain wildcards in the same way. The client uses the domain list if you do not specify a file specification.

Note: Journal-based backup might not fall back to the traditional incremental backup if the policy domain of your node is changed on the server, depending on

when the policy set within the domain was last updated and the date of the last incremental. In this case, you must force a full traditional incremental backup to rebind the files to the new domain. Use the `nojournal` option with the **incremental** command to specify that you want to perform a traditional full incremental backup, instead of the default journal-based backup.

When a user deletes a file with a long name, the Windows operating system might supply a short (compressed) name to the journal engine service. After the object is deleted, the compressed name can be reused and the deletion notice might no longer identify a unique object. During a journaled incremental backup, the attempt to expire the file fails because the compressed name is not known to the server. When this failure occurs, a record is placed in the journal, which indicates that the current directory is not exactly represented at the server. Use the `incrthreshold` option to specify what action is taken when this occurs.

The journal database is considered invalid and the client reverts to the traditional full incremental backup when any of the following events occur:

- A journaled file space name changes.
- The client node name changes.
- The client contacts a different server to do the backup.
- A policy changes occurs (new policy set activation).
- The journal is corrupted (out of space conditions, disk error).
- The journal service is not running.
- The journal service is stopped or started for any reason, even if it is restarted because the system is rebooted.

Journal-based backup differs from the traditional full incremental backup in the following ways:

- Tivoli Storage Manager does not enforce non-default copy frequencies (other than 0).
- Attribute changes to an object require a backup of the entire object.

You can use the `nojournal` option with the **incremental** command to perform a traditional full incremental backup instead of the default journal-based backup.

Multiple journal-based backup sessions are possible.

Backing up NTFS or ReFS volume mount points

If you perform an incremental backup of a file system on which a volume mount point exists, IBM Tivoli Storage Manager backs up the directory (junction) where the volume is mounted, but it does not traverse or back up the data on the mounted volume.

For example, if `C:\mount` is a mount point, then an incremental backup of the `C:\` drive backs up only the junction (`C:\mount`), and not the data under `C:\mount`.

Related concepts:

“Restoring NTFS or ReFS volume mount points” on page 735

“Restoring data on NTFS mounted volumes” on page 735

“Backing up data on NTFS or ReFS mounted volumes”

Backing up data on NTFS or ReFS mounted volumes

Backing up a volume from the mount point is especially useful for volumes that have no drive letter assignment. If the volume mounted on the mount point can also be referenced by drive letter, then the volume does not have to be backed up over the mount point.

For example, if the F:\ drive is mounted on C:\mount then the data can be backed up by including C:\mount or the F:\ drive in the domain. In this case, duplicate backups can be avoided by configuring the domain to back up C:\mount or the F:\ drive, but not both.

To back up the data on the mounted volume, run an incremental backup of the mount point by using the **incremental** command:

```
dsmc incremental c:\mount
```

You can also add C:\mount to the DOMAIN option to back up the data on the mount point as part of a domain incremental backup operation. For example, to back up the system state, the C:\ drive, and the data on the volume that is mounted on C:\mount as part of a scheduled incremental backup, configure a DOMAIN statement as follows:

```
domain c: c:\mount systemstate
```

If you use `exclude.dir` to exclude C:\mount, then:

- The C:\mount directory is not backed up during an incremental backup of the C:\ drive.
- Nothing is backed up during an attempt to back up C:\mount; a message is displayed indicating that C:\mount is excluded.

Related concepts:

“Restoring NTFS or ReFS volume mount points” on page 735

“Restoring data on NTFS mounted volumes” on page 735

“Backing up NTFS or ReFS volume mount points” on page 688

Back up Microsoft Dfs root

If you perform an incremental backup of Microsoft Dfs root with `dfsbackupmntpnt=yes` specified, the backup-archive client backs up only the junction points, *not* the subtree under the junctions.

If you want to traverse the Dfs tree and back up the files and subdirectories of any junction it encounters, specify the `dfsbackupmntpnt=no` option. If you want to backup both the Dfs tree structure and the data contained in the Dfs tree you must run two backups: one with `dfsbackupmntpnt=yes` and one with `dfsbackupmntpnt=no`.

This option has no effect if you are backing up individual junctions. The *exclude.dir* option behavior for Dfs junctions is same as for mounted virtual volumes.

Note: If a Dfs root is added or modified, the client will not back it up. You must specify the Dfs root in the `domain` option in the client options file (`dsm.opt`) regardless of whether DOMAIN ALL-LOCAL is specified.

Incremental-by-Date

An incremental-by-date backup backs up new and changed files with a modification date later than the date of the last incremental backup stored at the server, unless the files are excluded from backup by an **exclude** statement.

If an incremental-by-date is performed on only part of a file system, the date of the last full incremental is not updated, and the next incremental-by-date will back up

these files again. Use the **query filespace** command to determine the date and time of the last incremental backup of the entire file system.

To perform an incremental-by-date backup, use the **incrbydate** option with the **incremental** command.

Unlike a full incremental, an incremental-by-date does not maintain current server storage of *all* your workstation files 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.

Adaptive subfile backups

If you plan to perform a backup over a network device with limited bandwidth, such as a modem, you can reduce the network traffic by using the **subfilebackup** option.

If you plan to use this option, ensure that you specify the **subfilebackup**, **subfilecachepath** and **subfilecachesize** options during the initial backup of your file. The options specify whether to perform an adaptive subfile backup along with the path and size of the data to be transferred.

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 options file (`dsm.opt`). This value remains in effect for the entire interactive session unless overridden by a different value on a given interactive command. For example, if you set the `subdir` option to *yes* in your client options file (`dsm.opt`), and you specify `subdir=no` on the initial command line, the `subdir=no` setting remains in effect for the entire interactive session unless overridden by the `subdir=yes` value on a given interactive command. However, the `subdir=yes` value only affects the command it is entered on. When that command completes, the value reverts back to `subdir=no`, the value at the beginning of the interactive session.

You can enter all valid commands in interactive mode *except* the **schedule** and **loop** commands.

There are some options that you cannot use in the interactive session created by the **loop** command and are identified in the option description by this statement: *This option is valid only on the initial command line. It is not valid in interactive mode.*

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.

To end an interactive session, enter `quit`

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 11, “Processing options,” on page 295 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:

- c:\devel\project\proja
- c:\devel\project\projb
- c:\devel\project\projc

Command: macro backabc.mac

Where backabc.mac contains the following statements:

```
selective c:\devel\project\proja\*.*
selective c:\devel\project\projb\*.*
selective c:\devel\project\projc\*.*
```

Monitor Process

The **monitor process** command displays a list of current NAS (if NDMP support is enabled) image backup and restore processes for which the administrative user has authority. 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 all Windows clients.

Syntax

►►—MONitor Process—◄◄

Parameters

There are no parameters for this command.

Examples

Task Monitor current NAS image backup or restore processes.

Command: monitor process

Preview Archive

The **preview archive** command simulates an archive command without sending data to the server.

The **preview archive** command generates a tab-delineated text file that can be imported into a spreadsheet program. The preview contains information such as whether the file is excluded or included. If the file is excluded, the pattern, or reason, that the file is excluded is listed, along with the source for the pattern.

Supported Clients

This command is valid for all clients.

Syntax

►►—PREview—Archive—filespec—
└─filter=ALL─┐
└─filter=INCL─┐
└─filter=EXCL─┐
└─FILENAME= filename─┐
◄◄

└─CONsole─┐
└─TRAverse=Yes─┐
└─TRAverse=No─┐
◄◄

Parameters

filespec

Specifies the path and file name that you want to archive. Use wildcard characters to select a group of files or all the files in a directory.

-filter Specifies the output to display. 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.

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 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.

Query Adobjects

Use the **query adobjects** command to display information about the deleted objects that are located on the local Active Directory domain.

On Windows Server operating system clients, Active Directory object information can also be displayed from full system-state backups on the server.

Supported Clients

This command is valid for Windows Server OS clients only.

Syntax

```

>> Query ADOBJects [—sourcepathspec] [options]
  
```

Parameters

sourcepathspec

Specifies the Active Directory object or container to query. You can specify an asterisk (*) as a wildcard character. You can specify either the full distinguished name of an object or a container, or just the name attribute (cn or ou), where the wildcard might be used. You can also specify object GUID enclosed in braces ({}). The following special characters require an escape character, the backslash, (\), if any of them are contained in the name:

```

\
#
+
=
<
>
  
```

For example, "cn=test#" is entered as "cn=test\#".

The client cannot display any object names that contain an asterisk (*) as part of the name.

Table 75. Query Adobjects command: Related options

Option	Where to use
adlocation "Adlocation" on page 323	Command line only.
dateformat "Dateformat" on page 356	Client options file (dsm.opt) or command line.
detail "Detail" on page 363	Command line only.
pitdate (option is ignored when adlocation is not specified) "Pitdate" on page 482	Command line only.

Table 75. Query Adobjects command: Related options (continued)

Option	Where to use
pittime (option is ignored when adlocation is not specified) "Pitttime" on page 483	Command line only.
scrolllines "Scrolllines" on page 521	Client options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 522	Client options file (dsm.opt) or command line.
timeformat "Timeformat" on page 562	Client options file (dsm.opt) or command line.

Examples

Task Query all local deleted objects.

Command: query adobjects

Task Query all local deleted objects for a user with the name starting with Fred.

Command: query adobjects "cn=Fred*" -detail

Task Query all objects that are located in the Users container of the bryan.test.example.com domain from the server (Windows Server 2008 only).

Command: query adobjects "cn=Users,DC=bryan,DC=test,DC=ibm,DC=com" -adloc=server

Task Query all local deleted objects for organizational unit testou.

Command: query adobjects "ou=testou"

Task Query the local deleted object with a GUID of E079130D-3451-4C69-8349-31747E26C75B.

Command: query adobjects {E079130D-3451-4C69-8349-31747E26C75B}

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
- Creation 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 \\halley\m$\tsm620c.0901fa\debug\bin\
winnt_unicode\dsm.opt 03/03/2016
Archive Date: 03/03/2016
RetInit:STARTED Obj
Held:NO
Modified: 03/03/2016 19:43:00 Created: 03/01/2016 15:31:23
Compression Type: LZ4 Encryption Type: None Client-deduplicated: YES

```

For more information about the compression type, see “Compression” on page 351.

Supported Clients

This command is valid for all clients.

Syntax

```

>> Query ARchive [—options—] [—filespec—]
                  [—{filespace name—}—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 include *filespace name*, do not include a drive letter in the file specification. Drive label names are only used for removable media.

{filespace name}

Specifies the file space (enclosed in braces) on the server that contains the file that you want to query. The file space is the name on the workstation drive from which the file was archived. The following example is valid for specifying a UNC name: {'\\machine\C\$'}.

Use the *filespace name* if the name was changed or if you are querying files that were archived from another node with drive labels that are different from yours.

Note: You must specify a mixed or lowercase NTFS *filespace name* that is enclosed in quotation marks within braces, for example, {"NTFSDrive"}. Single quotation marks or double quotation marks are valid in loop mode. For example: {'NTFSDrive'} and {"NTFSDrive"} are both valid. In batch mode, only single quotation marks are valid.

Table 76. Query Archive command: Related options

Option	Where to use
dateformat “Dateformat” on page 356	Client options file (dsm.opt) or command line.
description “Description” on page 362	Command line only.
detail “Detail” on page 363	Command line only.
dirsonly “Dirsonly” on page 368	Command line only.

Table 76. Query Archive command: Related options (continued)

Option	Where to use
filelist "Filelist" on page 414	Command line only.
filesonly "Filesonly" on page 418	Command line only.
fromdate "Fromdate" on page 420	Command line only.
fromnode "Fromnode" on page 420	Command line only.
fromtime "Fromtime" on page 421	Command line only.
numberformat "Numberformat" on page 476	Client options file (dsm.opt) or command line.
querysummary "Querysummary" on page 495	Command line only.
scrolllines "Scrolllines" on page 521	Client options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 522	Client options file (dsm.opt) or command line.
subdir "Subdir" on page 548	Client options file (dsm.opt) or command line.
timeformat "Timeformat" on page 562	Client options file (dsm.opt) or command line.
today "Today" on page 564	Command line only.
totime "Totime" on page 565	Command line only.

Examples

Task Display a list of all your archived files in the c:\proj directory.

Command: q ar c:\proj*

Task Display a list of archived files from your c: drive with the description "January Ledgers".

Command: query archive c:\ -su=y -descr="January Ledgers"

Task Display a list of all your archived files in the c:\proj directory. Use the dateformat and timeformat options to reformat the dates and times.

Command: q ar -date=5 -time=4 c:\proj*

Task Display a list of all your archived files in the c:\dir1 directory. Use the detail option to display the last modification date and the creation date of each file.

Command: q ar -detail c:\dir1*

Task Display a list of archived files in the c:\proj directory that contains a file extension of .dev. Use the dateformat and timeformat options.

Command: q ar -date=5 -time=4 c:\proj*.dev

Task Recently you changed the label of your c:\ drive to store and archive some files. Then, yesterday the label was changed to dev and some more files were archived. Display a list of all the files you archived in the c:\proj directory when the label was store.

Command: q ar {store}\proj*

Task Recently you archived files from a diskette labeled docs. Display a list of all the files you archived.

Command: q ar {docs}*

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
- Creation date
- Compression type
- Encryption type
- Client-side data deduplication

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
----	-----	-----	-----
1,000,000 B	03/15/2016 14:33:17	DEFAULT	A \\eighth\n\$\testdir\myfile1.txt
Modified: 03/15/2016 14:31:42		Created: 03/15/2016 14:31:41	
Compression Type: LZ4 Encryption Type: None Client-deduplicated: YES			

For more information about the compression type, see “Compression” on page 351.

Supported Clients

This command is valid for all clients.

Syntax



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. Do not use wildcard characters when you query NAS file system images with `-class=nas` option setting.

If you include *filespace-name*, do not include a drive letter in the file specification. Drive label names are only used for removable media.

You can also use the following value for *filespec*:

systemstate

Displays the list of backup versions of Windows system state.

{filespace-name}

Specifies the file space, enclosed in braces, on the server that contains the file you want to query. This is the drive label or UNC name on the workstation drive from which the file was backed up. The following example shows how to specify a UNC name: `{ '\\machine\C$' }`.

Use the *filespace-name* if the name has changed, or if you want to query files backed up from another node with drive label names that are different from yours.

You must specify a mixed or lowercase NTFS or ReFS file space name that is enclosed in quotation marks and braces. For example, `{"NTFSDrive"}`. Single quotation marks or double quotation marks are valid in loop mode. For example: `{"NTFSDrive"}` and `{'NTFSDrive'}` are both valid. In batch mode, only single quotation marks are valid.

Table 77. Query Backup command: Related options

Option	Where to use
backupsetname “Backupsetname” on page 336	Command line only.
class “Class” on page 341	Command line only.
dateformat “Dateformat” on page 356	Client options file (dsm.opt) or command line.
detail “Detail” on page 363	Command line only.
dirsonly “Dirsonly” on page 368	Command line only.
filelist “Filelist” on page 414	Command line only.
filesonly “Filesonly” on page 418	Command line only.
fromdate “Fromdate” on page 420	Command line only.

Table 77. Query Backup command: Related options (continued)

Option	Where to use
fromowner "Fromnode" on page 420	Command line only.
fromtime "Fromtime" on page 421	Command line only.
inactive "Inactive" on page 428	Command line only.
nasnodename "Nasnodename" on page 471	Client options file (dsm.opt) or command line.
numberformat "Numberformat" on page 476	Client options file (dsm.opt) or command line.
pitdate "Pitdate" on page 482	Command line only.
pittime "Pittime" on page 483	Command line only.
querysummary "Querysummary" on page 495	Command line only.
scrolllines "Scrolllines" on page 521	Client options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 522	Client options file (dsm.opt) or command line.
subdir "Subdir" on page 548	Client options file (dsm.opt) or command line.
timeformat "Timeformat" on page 562	Client options file (dsm.opt) or command line.
todate "Todate" on page 564	Command line only.
totime "Totime" on page 565	Command line only.

Examples

```
dsmc query backup c:\* -subdir=yes -querysummary
```

```
dsmc query archive c:\* -subdir=yes -querysummary
```

Task Query files from the abc file space proj directory.

```
dsmc query backup {"abc"}\proj\*.*
```

Task Display a list of all active and inactive backup versions that were backed up from the c:\proj directory.

```
dsmc q backup -ina c:\proj\*
```

Task Display a list of all your backups in the c:\dir1 directory. Use the detail option to display the last modification date and the creation date of each file.

```
dsmc q backup -detail c:\dir1\*
```

Task Display a list of all active and inactive backup versions that were backed up from the c:\proj directory. Use the `dateformat` and `timeformat` options to reformat the dates and times.

```
dsmc q b -date=5 -time=4 -ina c:\proj\*
```

Task Last week you backed up files from a diskette labeled **docs**. Display a list of those files.

```
dsmc q b {docs}\*
```

Task Query file system images from the nas2 NAS file server.

```
dsmc query backup -nasnodename=nas2 -class=nas
```

Task Display a list of all files from your c drive that are contained in the backup set `weekly_accounting_data.32145678`.

```
dsmc query backup c:\* -subdir=yes  
-backupsetname=weekly_accounting_data.32145678
```

Task Display information about all the active and inactive backup versions of the system state on the server.

```
dsmc query backup -ina systemstate
```

Related information

“Restore data from a backup set” on page 206

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 options file (`dsm.opt`). The value in the client 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 341

“Nasnodename” on page 471

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.

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.

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 78. Query Backupset command: Related options

Option	Where to use
description "Description" on page 362	Command line only.
scrolllines "Scrolllines" on page 521	Client options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 522	Client 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 `monthy_financial_data` from the Tivoli Storage Manager server.

Command: `query backupset
-backupsetname=monthy_financial_data.12345678`

Task Query the backup set in the file `c:\budget\weekly_budget_data.ost`.

Command: `query backupset -backupsetname=c:\budget\
weekly_budget_data.ost loc=file`

Task Query the backup set from the `\\.\tape0` tape device.

Command: `dsmc query backupset -backupsetname=\\.\tape0 -loc=tape`

Related information

“Restore data from a backup set” on page 206

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 79. Query Backupset command: Related options

Option	Where to use
description "Description" on page 362	Command line only.
scrolllines"Scrolllines" on page 521	Client options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 522	Client 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: `\budget\weekly_budget_data.ost`.

Command: `query backupset c:\budget\weekly_budget_data.ost loc=file`

Task Query the backup set from the `\\.\tape0` tape device.

Command: `dsmc query backupset \\.\tape0 -loc=tape`

Related information

"Restore data from a backup set" on page 206

Query Filespace

The **query filespace** 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.

A Unicode file space name might not display correctly if the server is unable to display the Unicode name. In this case, use the file space identifier (fsID) to identify these file spaces on the server. Use the **query filespace** command with the **detail** option to determine the fsID of a file space.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

filespaceName
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 80. Query Filespace command: Related options

Option	Where to use
backupsetname "Backupsetname" on page 336	Command line only.
class "Class" on page 341	Command line only.
dateformat "Dateformat" on page 356	Client options file (dsm.opt) or command line.
detail "Detail" on page 363	Command line only.
fromnode "Fromnode" on page 420	Command line only.
nasnodename "Nasnodename" on page 471	Client options file (dsm.opt) or command line.
scrolllines "Scrolllines" on page 521	Client options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 522	Client options file (dsm.opt) or command line.
timeformat "Timeformat" on page 562	Client 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 filepace -date=5 -time=4
```

Query a file space from the `nas2` NAS file server.

```
query filepace -nasnodename=nas2 -class=nas
```

Display the `\\florence\c$` file space.

```
query filepace \\florence\c$
```

Display all of the file space names on the server with a file space name that ends in '\$' belonging to system named `florence`.

```
query filepace \\florence\*$
```

Display file spaces in the backup set named `monthly_accounting.23456789`.

```
query filepace -backupsetname=monthly_accounting.23456789
```

Display detailed file space information that shows the replication status during a failover.

Command:

```
query filepace -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 206

“Automated client failover overview” on page 60

Related tasks:

“Determining the status of replicated client data” on page 65

Related reference:

“Nasnodename” on page 471

“Class” on page 341

“Nrtablepath” on page 475

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 options file (`dsm.opt`). The value in the client options file is the default, but this value can be overridden on the command line. If the `nasnodename` option is not specified in the client options file, it must be specified on the command line when processing NAS file systems.

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 file space** 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.

These backups inactivate any previous full and differential backups:

```
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.

Syntax



Parameters

filespec

Specifies the virtual file space name (enclosed in braces) and the group name on the server that you want to query.

Table 81. Query Group command: Related options

Option	Where to use
fromnode "Fromnode" on page 420	Command line only.
inactive "Inactive" on page 428	Command line only.

Table 81. Query Group command: Related options (continued)

Option	Where to use
pitdate "Pitdate" on page 482	Command line only.
pittime "Pittime" on page 483	Command line only.
showmembers "Showmembers" on page 526 (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 706

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. Because image backup allows you to back up only used blocks in a file system, the stored image size on the Tivoli Storage Manager server could be smaller than the volume size. For online image backups, the stored image can be larger than the file system based on the size of the cache files.
- File system type
- Backup date and time
- Management class 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 all Windows clients.

Syntax



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 82. Query Image command: Related options

Option	Where to use
backupsetname "Backupsetname" on page 336	Command line only.
dateformat "Dateformat" on page 356	Client option file (dsm.opt) or command line.
fromnode "Fromnode" on page 420	Command line only.
inactive "Inactive" on page 428	Command line only.
numberformat "Numberformat" on page 476	Client option file (dsm.opt) or command line.
pitdate "Pitdate" on page 482	Command line only.
pittime "Pittime" on page 483	Command line only.
scrolllines "Scrolllines" on page 521	Client options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 522	Client options file (dsm.opt) or command line.
timeformat "Timeformat" on page 562	Client option file (dsm.opt) or command line.

Examples

Task Display all backed up images.

Command: q image

Task Display active and inactive version of the h: image.

Command: q im h: -inactive

Task Display all images that are contained within the backup set weekly_backup_data.32145678.

Command: query image -backupsetname=weekly_backup_data.32145678

Related information

“Restore data from a backup set” on page 206

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.

Use the **display** option to display the files that are included or excluded from a file system back up operation.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

test pattern

Use for testing the validity of patterns you want to use in your include-exclude list. When you use a test pattern with this command, the following occurs:

- The internal include-exclude list is not displayed
- The pattern is processed as if it 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.

-DISPLAY=basic | vssexcl | all

-DISPLAY=basic displays the files and directories that have been included or excluded by one of the following methods:

- The objects were included or excluded in the client options file.
- The objects were included or excluded in a server-side client option set.
- The objects were excluded by the operating system because they are contained in the HKEY_LOCAL_MACHINES\SYSTEM\CurrentControlSet\BackupRestore\FilesNotToBackup registry key.
- The objects were explicitly excluded by the client.

This is the default if a display value is not specified.

-DISPLAY=vssexcl displays a list of files that are excluded from a file system backup, because they are included when a system state backup is performed. Files that are backed up by a **backup systemstate** operation are protected by the VSS writer; you cannot include these files in a file system backup by adding them to an include statement in the dsm.opt file, or client option set.

-DISPLAY=all displays all files that are included or excluded during a file system backup.

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 inclexcl
```

Task Display a list of files that are excluded from file system backups because the VSS writer includes them in system state backups.

```
query inclexcl -display=vssexcl
```

Task Display a list of include-exclude statements. Display the management class that is associated with each statement.

```
query inclexcl -detail
```

Task Test the validity of this pattern: ..\?x?*.log

```
query inclexcl ..\?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 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 83. Query Mgmtclass command: Related options

Option	Where to use
detail "Detail" on page 363	Command line only.
fromnode "Fromnode" on page 420	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



Parameters

Table 84. Query Node command: Related options

Option	Where to use
type "Type" on page 567	Command line only.
scrolllines "Scrolllines" on page 521	Client options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 522	Client options file (dsm.opt) or command line.

Examples

Task Display all NAS nodes.
 Command: query node -type=nas

Related information

"Type" on page 567


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 85. Query Options command: Related options

Option	Where to use
scrolllines "Scrolllines" on page 521	Client options file (dsm.opt) or command line.
scrollprompt "Scrollprompt" on page 522	Client 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 60

Related tasks:

"Determining the status of replicated client data" on page 65

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 The following example displays the output when you use **query restore**:

```
--- Restartable Restore Information ---
Restartable Session: 1
  Start date/time: 10/17/2001 15:18:22
    Source: {"\\ers\c$"}\data\proposals\*
    Destination: - not specified by user -

Restartable Session: 2
  Start date/time: 10/17/2001 15:20:01
    Source: {"\\ers\c$"}\data\spreadsheets\*
    Destination: - not specified by user -
```

Query Schedule

The **query schedule** command displays the events 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 Session—◄◄

Parameters

There are no parameters for this command.

Examples

Task Display your session information.

Command: query session

A sample **query session** display follows:

```
Server Name.....: HALLEY_SERVER1
Server Type.....: Windows
Archive Retain Protect...: "No"
Server Version.....: Ver. 6, Rel. 2, Lev. 0.0
Last Access Date.....: 09/03/2009 09:08:13
Delete Backup Files.....: "No"
Delete Archive Files.....: "Yes"
Deduplication.....: "Server Only"
```

```
Node Name.....: HALLEY
User Name.....:
```

Possible client-side deduplication values:

- None
 - Displayed when connected to a pre V6.1 Tivoli Storage Manager server
- Server Only
- Client Or Server

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.
- ENV - Environment variables.
- ERRORLOG - The client error log file.
- FILE - Attributes for the file name that you specify.
- FILESNOTTOBACKUP - Enumeration of Windows Registry key:

```
HKEY_LOCAL_MACHINE\
  SYSTEM\
    CurrentControlSet\
      BackupRestore\
        FilesNotToBackup
```

This key specifies those files that are not to be backed up. The **query inclexcl** command indicates that these files are excluded per the operating system.

- INCLEXCL - Compiles a list of include-exclude in the order in which they are processed during backup and archive operations.
- KEYSNOTTORESTORE - Enumeration of Windows Registry key:

```
HKEY_LOCAL_MACHINE\
  SYSTEM\
    ControlSet001\
      BackupRestore\
        KeysNotToRestore
```

This key specifies those Windows Registry keys that are not to be restored.

- MSINFO - Windows system information (output from MSINFO32.EXE).
- OPTIONS - Compiled options.
- OSINFO - Name and version of the client operating system
- POLICY - Policy set dump.
- REGISTRY - Tivoli Storage Manager-related Windows Registry entries.
- SCHEDLOG - The contents of the schedule log (usually dsmsched.log).
- SFP - The list of files that are protected by Windows System File Protection, and for each file, indicates whether that file exists. These files are backed up as part of the SYSFILES system object.
- SFP=<filename> - Indicates whether the specified file (*filename*) is protected by Windows System File Protection. For example:

```
SFP=C:\WINNT\SYSTEM32\MSVCRT.DLL
```
- SYSTEMSTATE - Windows system state information.
- CLUSTER - Windows cluster information.
- ENCRYPT - Available encryption methods.

Note:

1. Use the filename option to specify a file name in which to store the information gathered from the items you specify. If you do not specify a file name, by default the information is stored in the dsminfo.txt file.
2. Use the console option if you want to output the information to the console.

Table 86. Query Systeminfo command: Related options

Option	Where to use
console "Console" on page 352	Command line only.
filename "Filename" on page 417	Command line only.

Examples

Task Gather and store the contents of the dsm.opt file and the Tivoli Storage Manager error log file in the tsminfo.txt file.

Command: query systeminfo dsmoptfile errorlog
-filename=tsminfo.txt

Related information

"Filename" on page 417

"Console" on page 352

Query Systemstate

Use the **query systemstate** command to display information about a backup of the system state on the Tivoli Storage Manager server, or system state inside a backup set from the Tivoli Storage Manager server, when the backupsetname option is specified.

The output indicates whether the object is active ("A") or inactive ("I"). Only active objects are listed unless the inactive option is specified with the command. The backup-archive client on Windows supports standard and detailed format.

Supported Clients

This command is valid for supported Windows clients only.

Syntax

►► Query SYSTEMState options ►►

Parameters*Table 87. Query Systemstate command: Related options*

Option	Where to use
backupsetname "Backupsetname" on page 336	Command line only.

Table 87. Query Systemstate command: Related options (continued)

Option	Where to use
dateformat "Dateformat" on page 356	Client option file (dsm.opt) or command line.
inactive "Inactive" on page 428	Command line only.
numberformat "Numberformat" on page 476	Client option file (dsm.opt) or command line.
pitdate "Pitdate" on page 482	Command line only.
pittime "Pittime" on page 483	Command line only.
showmembers "Showmembers" on page 526	Command line only.
timeformat "Timeformat" on page 562	Client option file (dsm.opt) or command line.
detail "Detail" on page 363	Command line only.

Examples

Task Display information about the active backup of the system state on the Tivoli Storage Manager server.

Command: query systemstate

Task Display information about the active backup of the system state on the Tivoli Storage Manager server.

Command: query systemstate -detail

Task Display information about the active backup of the system state that is contained within the backup set daily_backup_data.12345678.

Command: query systemstate
-backupsetname=daily_backup_data.12345678

Task To display information about Active Directory, enter the following command: query systemstate -detail.

Locate information that is related to Active Directory in the output.

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 Windows clients that are configured to perform vCloud Director vApp backups.

Syntax



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=vc1,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 88. Query VAPP command: Related options

Option	Where to use
asnodename	Client user-options file (dsm.opt), on the command line, or on the General tab in the Preferences editor.
detail	Command line only. Shows more detail about each vApp.
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: Support_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

Table 89. Query VM command: Related options for VMware virtual machine queries.

Option	Where to use
detail "Detail" on page 363 Valid for vmbackuptype=fullvm Valid for -vmrestoretype	Command line.
inactive "Inactive" on page 428 Valid for vmbackuptype=fullvm	Command line.
pitdate "Pitdate" on page 482 Valid for vmbackuptype=fullvm	Command line.
pittime "Pittime" on page 483 Valid for vmbackuptype=fullvm	Command line.
vmbackuptype "Vmbackuptype" on page 581	Command line or client options file.
vmchost "Vmchost" on page 583	Command line or client options file.
vmcpw "Vmcpw" on page 584	Command line or client options file.
vmcuser "Vmcuser" on page 586	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   DEFAULT    VMFULL  A    vm1
   11:34:14
2  09/26/2010   DEFAULT    VMFULL  A    vm2
   11:34:37
3  09/26/2010   DEFAULT    VMFULL  A    vm3
   11:34:49
4  09/26/2010   DEFAULT    VMFULL  I    vm2
   12:35:08
5  09/26/2010   DEFAULT    VMFULL  I    vm3
   12:35:18

```

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

```

File Level

```

QUERY VM -vmbackuptype=FILE

Filespace Query for Virtual Machine 'vm1'

Accessing as node: vm1
#      Last Incr Date      Type      File Space Name
-      -      -      -
  1  09/25/2013 12:19:23  NTFS      \\vm1\c$  fsID: 2
  2  09/25/2013 12:19:25  NTFS      \\vm1\e$  fsID: 3
  3  09/25/2013 12:19:22  NTFS      \\vm1\z$  fsID: 1

Filespace Query for Virtual Machine 'vm2'

Accessing as node: vm2
#      Last Incr Date      Type      File Space Name
-      -      -      -
  1  09/25/2013 10:49:10  NTFS      \\vm2\c$  fsID: 2
  2  09/25/2013 10:49:09  NTFS      \\vm2\h$  fsID: 1

Filespace Query for Virtual Machine 'vm3'

Accessing as node: vm3
#      Last Incr Date      Type      File Space Name
-      -      -      -
  1  09/25/2013 10:49:22  NTFS      \\vm3\c$  fsID: 2
  2  09/25/2013 10:49:21  NTFS      \\vm3\h$  fsID: 1

```

The following command returns a list of VMs that are running an instant restore operation.

```
q vm * -vmrestoretype=instantrestore
```

Query all VMware virtual machines that were backed up using -vmbacktype=fullvm:

```
q vm * -vmbackuptype=fullvm
```


Query VM for Microsoft Hyper-V virtual machines

Use the **query vm** command to determine which Hyper-V virtual machines were backed up.

Supported Clients

This command is valid on Windows clients that are installed on a Hyper-V host system.

Syntax



Parameters

vmname
Specifies the virtual machine host name that you want to query. The virtual machine name is case-sensitive. If you specify a VM name on the command, the name cannot contain wildcard characters.

If you omit the virtual machine name, the command displays all VM backups on the Tivoli Storage Manager server.

Table 90. Query VM command: Related options for Hyper-V virtual machine queries.

Option	Where to use
detail "Detail" on page 363	Command line. Displays the details of each disk (label, name) and its status (protected or excluded), and incremental-forever backup performance statistics.
inactive "Inactive" on page 428 Valid for vmbackuptype=hypervfull	Command line.
pitdate "Pitdate" on page 482 Valid for vmbackuptype=hypervfull	Command line.
pittime "Pittime" on page 483 Valid for vmbackuptype=hypervfull	Command line.
vmbackuptype=hypervfull "Vmbackuptype" on page 581	Command line or client options file.

Examples

Task List all virtual machines that have been backed up by the backup-archive client on the Hyper-V host.

```
dsmc query vm -vmbackuptype=hypervfull
```

Query VM examples (Hyper-V)

The following is an example of **query VM** command that displays information about Hyper-V virtual machines that have been backed up.

```
dsmc query vm -vmbackuptype=hypervfull
```

#	Backup Date	Mgmt Class	Type	A/I	Virtual Machine
1	07/24/2014	DEFAULT	Full	A	Windows 2008
2	07/25/2014	DEFAULT	Full	A	Windows 2008
3	08/04/2014	DEFAULT	Full	A	Windows 2008

The next example shows output from a **query VM** that includes the **-detail** option. This command includes a virtual machine name so the output is only displayed for the specified virtual machine. The detailed output includes the type of backup that was performed, the size of the virtual machine, information about its disks, and statistics.

```
tsm> q vm HPV2VM1 -detail
```

Periodic Full - Full

Query Virtual Machine for Full VM backup

#	Backup Date	Mgmt Class	Size	Type	A/I	Virtual Machine
1	03/14/2014 09:58:44	STANDARD	60.00 GB	IFINCR	A	HPV2VM1

The size of this incremental backup: 8.62 MB
The number of incremental backups since last full: 4
The amount of extra data: 0
The TSM objects fragmentation: 3
Backup is represented by: 174 TSM objects
Application protection type: n/a
Disk[1]Label: Hard Disk 1
Disk[1]Name: na
Disk[1]Status: Protected

All averages are calculated only for incremental forever backups displayed above.
The average size of incremental backup: 8.62 MB
The average number of incremental backups since last full: 4
The average overhead of extra data: 0
The average TSM objects fragmentation: 3
The average number of TSM objects per backup: 174

The following example shows the syntax to use to list all Hyper-V virtual machines that have been backed up by this node:

```
q vm -vmbackuptype=hypervfull
```

Related concepts:

“Back up virtual machines on a Hyper-V system” on page 178

Related tasks:

“Preparing the environment for full backups of VMware virtual machines” on page 167

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. When you restore directory, its modification date and time is set to the date and time of the restore, not to the date and time the directory had when it was backed up. This is because the client restores the directories first, then adds the files to the directories.
2. An error occurs if you attempt to restore a file whose name is the same the short name of an existing file. For example, if you attempt to restore a file that you specifically named ABCDEF~1.DOC into the same directory where a file named abcdefghijk.doc exists, the restore fails because the Windows operating system equates the file named abcdefghijk.doc to a short name of ABCDEF~1.DOC. The restore function treats this as a duplicate file.

If this error occurs, perform any of the following actions to correct it:

- Restore the file with the short file name to a different location.
- Stop the restore and change the name of the existing file.
- Disable the short file name support on Windows.

- Do not use file names that would conflict with the short file naming convention; for example, do not use ABCDEF~1.DOC.

If you set the **subdir** option to yes when 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.

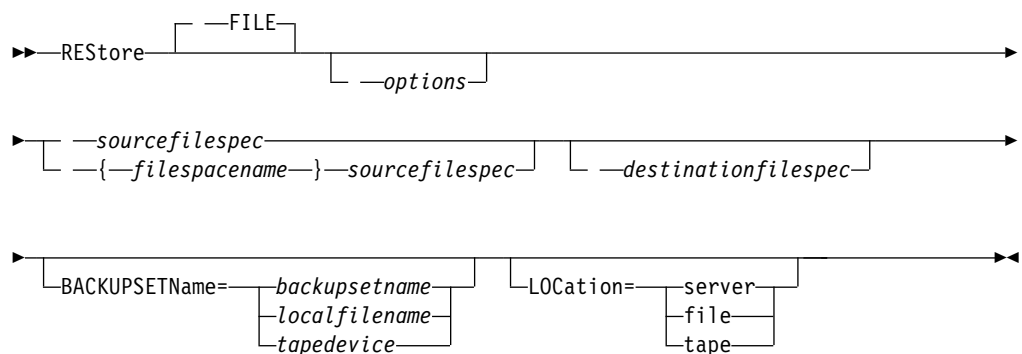
For more information, see the Microsoft Knowledge Base article Q121007, entitled *How to Disable the 8.3 Name Creation on NTFS Partitions*, for more information.

If the **restore** command is tried again because of a communication failure or session loss, the transfer statistics display the bytes that the client attempted to transfer across all command attempts. Therefore, the statistics for bytes transferred might not match file statistics, such as those for file size.

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.

Note: If you include *filespace*, do not include a drive letter in the file specification.

{filespace}

Specifies the file space (enclosed in braces) on the server that contains the files you want to restore. This is the name on the workstation drive from which the files were backed up.

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.

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

destinationfilespec

Specifies the path and file name where you want to place the restored files. If you do not specify a destination, the client restores the files to the original source path.

When you enter the *destinationfilespec*, consider the following points:

- 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.

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 91. Restore command: Related options

Option	Where to use
asrmode "Asrmode" on page 327	Command line only.
dateformat "Dateformat" on page 356	Client options file (dsm.opt) or command line.
dirsonly "Dirsonly" on page 368	Command line only.
filelist "Filelist" on page 414	Command line only.
filesonly "Filesonly" on page 418	Command line only.
fromdate "Fromdate" on page 420	Command line only.
fromnode "Fromnode" on page 420	Command line only.
fromtime "Fromtime" on page 421	Command line only.
ifnewer "Ifnewer" on page 426	Command line only.
inactive "Inactive" on page 428	Command line only.
latest "Latest" on page 456	Command line only.
numberformat "Numberformat" on page 476	Client 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 481	Command line only.
pitdate "Pitdate" on page 482	Command line only.
pittime "Pittime" on page 483	Command line only.
preservepath "Preservepath" on page 490	Command line only.
replace "Replace" on page 499	Client options file (dsm.opt) or command line.
skipntpermissions "Skipntpermissions" on page 528	Client options file (dsm.opt) or command line.
skipntsecuritycrc "Skipntsecuritycrc" on page 528	Client options file (dsm.opt) or command line.

Table 91. Restore command: Related options (continued)

Option	Where to use
subdir "Subdir" on page 548	Client options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 553	Client options file (dsm.opt) or command line.
timeformat "Timeformat" on page 562	Client options file (dsm.opt) or command line.
todate "Todate" on page 564	Command line only.
totime "Totime" on page 565	Command line only.

Examples

- Task** Restore a single file named budget.fin.
- ```
restore c:\devel\projecta\budget.fin
```
- Task** Restore a single file named budget.fin, which exists in the current directory.
- ```
restore file budget.fin
```
- Task** Restore files from the abc file space proj directory.
- ```
rest {"abc"}\proj*.*
```
- Task** Restore all files with a file extension of .c from the c:\devel\projecta directory.
- ```
rest c:\devel\projecta\*.c
```
- Task** Restore all files with an extension of .c from the \devel\projecta directory that is located in the winnt file space.
- ```
rest {winnt}\devel\projecta*.c
```
- Task** Restore all files with a file extension of .c from the c:\devel\projecta directory to the c:\newdevel\projectn\projecta directory. If the projectn or projectn\projecta directory does not exist, it is created.
- ```
restore c:\devel\projecta\*.c c:\newdevel\projectn\
```
- Task** Restore files in the c:\project directory. Use the pick and inactive options to select active and inactive backup versions.
- ```
restore c:\project* -pi -ina
```
- Task** Restore all files in the c:\mydir directory to their state as of 1:00 PM on August 17, 2002.
- ```
restore -pitd=8/17/2002 -pitt=13:00:00 c:\mydir\
```
- Task** Restore a file from the renamed file space \\your-node\h\$_OLD to its original location. Enter both the source and destination as follows:
- ```
res \\your-node\h$_OLD\docs\myresume.doc h:\docs\
```
- Task** Restore all files in the c:\mydir directory to their state as of 1:00 PM on August 17, 2002.
- ```
restore -pitd=8/17/2002 -pitt=13:00:00 c:\mydir\
```

Task Restore a single file named `budget.fin` contained within the backup set `daily_backup_data.12345678`.

```
restore c:\projecta\budget.fin  
-backupsetname=daily_backup_data.12345678 -location=server
```

Related information

“Restore data from a backup set” on page 206

“Preservepath” on page 490

Restoring NTFS or ReFS volume mount points

When restoring a file system that contains a volume mount point, only the mount point (directory) is restored. The data on the volume mounted on that directory is not restored.

A mount point can also be restored individually. For example, `C:\mount` is a mount point and has been backed up as part of the `C:\` drive on the system named STORMAN. The following command can be used to restore this mount point:

```
dsmc restore {\storman\c$}\mount
```

The braces (`{` and `}`) are required if you also backed up the data on the mounted volume from the mount point. Without the braces, the client restores data from the file space with the longest name that matches the file specification. If you backed up the data through the mount point, the backups are stored in a file space named `\\storman\c$\mount`. The braces are used to specify that the data be restored from the `\\storman\c$` file space.

The mount point cannot be restored if any of the following conditions is true:

- The mount point already exists.
- A non-empty directory matching the mount point name exists.
- A file matching the mount point name exists.

Related concepts:

“Restoring data on NTFS mounted volumes”

“Backing up NTFS or ReFS volume mount points” on page 688

“Backing up data on NTFS or ReFS mounted volumes” on page 688

Restoring data on NTFS mounted volumes

The mount point must exist before the data on the mounted volume can be restored to its original location.

If the mount point does not exist then you can restore it as described in “Restoring NTFS or ReFS volume mount points.”

For example, `C:\mount` is a mount point and it has been backed up as part of the `C:\` drive on a system called STORMAN. The data on the mounted volume has also been backed up. After ensuring that the mount point has been restored, the following command can be used to restore the data:

```
dsmc restore c:\mount\* -subdir=yes
```

Important: If the mount point does not exist, then the data will instead be restored to the root of the mount point's file system. For example, the following objects exists on `C:\mount`:

- C:\mount\projects\2009plan.doc
- C:\mount\projects\2010plan.doc
- C:\mount\master_list.xls

If the restore command (shown previously) is issued, but the mount point does not exist, then these objects are restored to the root of the C:\ drive as follows:

- C:\projects\2009plan.doc
- C:\projects\2010plan.doc
- C:\master_list.xls

Note: When using the GUI client and web client to view objects in a file space that contains a mount point, the mount point is displayed as an empty directory. Objects from the data on the mounted volume can be viewed and restored by viewing the file space for that mounted volume.

Related concepts:

“Restoring NTFS or ReFS volume mount points” on page 735

“Backing up NTFS or ReFS volume mount points” on page 688

“Backing up data on NTFS or ReFS mounted volumes” on page 688

Restore Microsoft Dfs junctions

To restore Microsoft Dfs junctions, you must restore Microsoft Dfs root.

If you select the junction point itself, the backup-archive client restores data under junction, but not the junction itself. If you select a junction point that no longer exists under Dfs root, the client creates a local directory under Dfs root with the same name as the junction before restoring data.

Restore active files

When restoring active and inactive versions of the same file using the replace option, only the most recently restored file is replaced.

Universal Naming Convention restores

The client stores files on the Tivoli Storage Manager server using the Windows Universal Naming Convention (UNC), not the drive letter. The UNC name is the network name for the file. The system name is a part of the UNC name. For example, if your system name is STAR and you have a file named c:\doc\h2.doc, the UNC name is \\star\c\$\doc\h2.doc.

When you restore files on the same system from which they were backed up, you can use the local drive letter or the UNC name to refer to the file. For example, either of the following will restore c:\doc\h2.doc to its original location:

```
dsmc restore c:\doc\h2.doc
dsmc restore \\star\c$\doc\h2.doc
```

When you restore files on a system with a different name, then you must use the UNC name to refer to the file. This is true even if you are restoring to the same physical system, but the system name has changed since the backup occurred.

For example, if you back up c:\doc\h2.doc on system STAR and you want to restore it to system METEOR then you must use the UNC name to refer to the file. You must also specify a destination restore location. This is because the default

behavior is to restore the file to its original location, which would be on system STAR. To restore the file to system METEOR, you can run either of the following on METEOR:

```
dsmc restore \\star\c$\doc\h2.doc c:\
dsmc restore \\star\c$\doc\h2.doc \\meteor\c$\
```

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.

If you want to restore from file spaces that are not Unicode-enabled, you must specify the source on the server and a destination on the client. For example, you backed up your H disk named \\your-node\h\$ prior to installing the Unicode-enabled client. After the installation, you issue the following command for a selective backup:

```
sel h:\logs\*.log
```

Before the backup takes place, the server renames the file space to \\your-node\h\$_OLD. The backup continues placing the data specified in the current operation into the Unicode-enabled file space named \\your-node\h\$. That file space now contains only the \logs directory and the *.log files. If you want to restore a file from the (old) *renamed* file space to its original location, you must enter both the source and destination as follows:

```
restore \\your-node\h$_OLD\docs\myresume.doc h:\docs\
```

Restore named streams

The backup-archive client restores named streams on a file basis only.

Windows directories can contain named streams. Named streams attached to a directory will always be overwritten (regardless of the value of the prompt option) during a restore operation.

Restore sparse files

When restoring sparse files to a non-NTFS or non-ReFS file system, set the Tivoli Storage Manager server communication time-out value (idletimeout) to the maximum value of 255 to avoid client session timeout.

The backup-archive client is restricted to restoring sparse files that are less than 4 gigabytes in size.

The following issues apply if more data is restored than the Microsoft disk quota allows:

- If the user who performs the restore has a disk quota (for example, the user belongs to the Backup Operator Group), the client does not restore any data that exceeds the disk quota of the restore user and displays a "Disk Full" message.
- If the user who performs the restore does not have a disk quota (for example, the user belongs to the Administrator Group), the client does restore all data and transfers ownership of the files that exceed the disk quota of the original owner to the user who is performing the restore (in this case, the Administrator).

Restore Adobjects


Use the **restore adobjects** command to restore individual Active Directory objects from the local Deleted Objects container.

Backup-archive clients that run on Windows Server platforms can restore individual Active Directory objects from full system-state backups stored on the Tivoli Storage Manager server.

Supported Clients

This command is valid for Windows Server OS clients.

Syntax



Parameters

sourcepathspec

Specifies the Active Directory object or container to restore. If a container is specified, its contents are also restored. You can either specify the full distinguished name of an object or a container, or just the name attribute ('cn' or 'ou'), where the wildcard might be used. The following special characters require an escape character, the backslash, (\), if any of them are contained in the name:

$$\backslash \# + = < >$$

For example, "cn=test#" is entered as "cn=test\#".

The client cannot display any object names that contain an asterisk (*) as part of the name.

Do not use wildcards when you specify a distinguished name.

Table 92. Restore Aobjects command: Related options

Option	Where to use
adlocation “Adlocation” on page 323	Command line only.
dateformat (the option is ignored when adlocation is not specified) “Dateformat” on page 356	Client options file (dsm.opt) or command line.
pitdate (the option is ignored when adlocation is not specified) “Pitdate” on page 482	Command line only.

Table 92. Restore Adobjects command: Related options (continued)

Option	Where to use
pittime (the option is ignored when adlocation is not specified) "Pittime" on page 483	Command line only.
replace "Replace" on page 499	Client options file (dsm.opt) or command line.
timeformat (the option is ignored when adlocation is not specified) "Timeformat" on page 562	Client options file (dsm.opt) or command line.

Examples

Task Restore a specific deleted Active Directory object.

Command: restore adobj
"CN=Administrator,CN=Users,DC=bryan,DC=test,DC=ibm,DC=com"

Task Restore all deleted objects that were originally located in the Users container.

Command: restore adobj "CN=Users,DC=bryan,DC=test,DC=ibm,DC=com"

Task Restore individual Active Directory objects from the Tivoli Storage Manager server. Use the pitdate and pittime options to select from a list of more recent and less recent backup versions (Windows Server 2008 only).

Command: restore adobj "cn=guest" -adloc=server
-pitdate=03/17/2008 -pittime=11:11:11

Task Restore all deleted users with the name starting with Fred.

Command: restore adobjects "cn=Fred*"

Task Restore all deleted organizational units with the name testou.

Command: restore adobjects "ou=testou"


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.

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

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 93. Restore Backupset command: Related options

Option	Where to use
dironly "Dironly" on page 368	Command line only.
filesonly "Filesonly" on page 418	Command line only.
ifnewer "Ifnewer" on page 426	Command line only.
preservepath "Preservepath" on page 490	Command line only.
quiet "Quiet" on page 496	Client options file (dsm.opt) or command line.
replace "Replace" on page 499	Client options file (dsm.opt) or command line.
skiptpermissions "Skiptpermissions" on page 528	Client options file (dsm.opt) or command line.
subdir "Subdir" on page 548	Client 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 from the `\\.\tape0` device.

```
dsmc restore backupset
-backupsetname=\\.\tape0 -loc=tape
```

Task Restore groups from the backup set mybackupset.12345678 on the Tivoli Storage Manager server to the c:\newdevel\projectn directory. The groups' virtual file space is accounting.

```
dsmc restore backupset {accounting}\*
c:\newdevel\projectn\
-backupsetname=mybackupset.12345678
-loc=server -subdir=yes
```

Task Restore the entire backup set contained in the file: c:\budget\weekly_budget_data.ost.

```
dsmc restore backupset
-backupsetname=c:\budget\weekly_budget_data.ost
-loc=file
```

Task Restore the \budget\ directory and subdirectories from the backup set contained in the file: c:\budget\weekly_budget_data.ost.

```
dsmc restore backupset m:\budget\*
-backupsetname=c:\budget\weekly_budget_data.ost
-loc=file -subdir=yes
```

Task Restore the file \budget\salary.xls from the backup set contained in the file: c:\budget\weekly_budget_data.ost.

```
dsmc restore backupset m:\budget\salary.xls
-backupsetname=c:\budget\weekly_budget_data.ost
-loc=file -subdir=yes
```

Task Restore the system state that is contained in the backup set file: \recovery\backupsetfile.name. Replace all components; do not prompt for each component.

```
dsmc restore backupset systemstate
-BACKUPSetname=c:\recovery\backupsetfile.name
-loc=file -replace=all
```

Related information

"Restore data from a backup set" on page 206

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.
- 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 456

"Restore" on page 730

"Restore Image" on page 748

"Restore Backupset" on page 739

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.
- 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

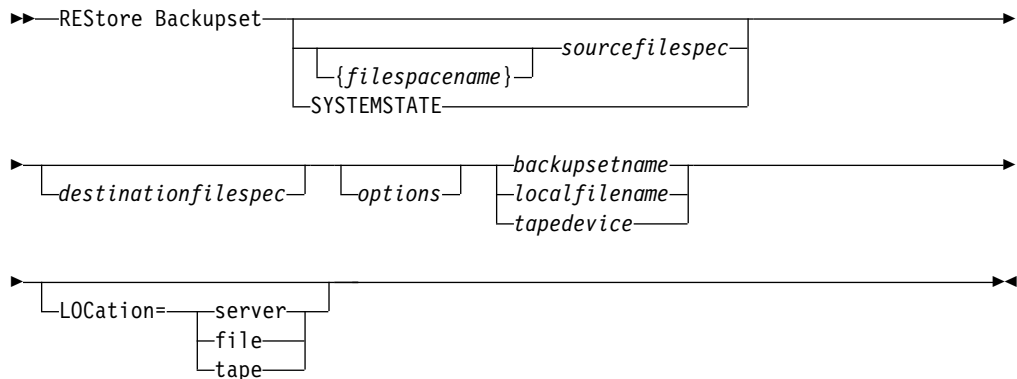
The **restore backupset** command can be used without the **backupsetname** parameter.

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.

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

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 from the `\\.\tape0` device.

```
dsmc restore backupset \\.\tape0 -loc=tape
```

Task Restore groups from the backup set `mybackupset.12345678` on the Tivoli Storage Manager server to the `c:\newdevel\projectn` directory. The groups' virtual file space is accounting.

```
dsmc restore backupset mybackupset.12345678 {accounting}\*  
c:\newdevel\projectn\ -loc=server -subdir=yes
```

Task Restore the entire backup set contained in the file: c:\budget\weekly_budget_data.ost.

```
dsmc restore backupset c:\budget\weekly_budget_data.ost -loc=file
```

Task Restore the \budget\ directory and subdirectories from the backup set contained in the file: c:\budget\weekly_budget_data.ost.

```
dsmc restore backupset c:\budget\weekly_budget_data.ost m:\budget\  
-loc=file -subdir=yes
```

Task Restore the file \budget\salary.xls from the backup set contained in the file: c:\budget\weekly_budget_data.ost.

```
dsmc restore backupset c:\budget\weekly_budget_data.ost  
m:\budget\salary.xls -loc=file -subdir=yes
```

Task Restore the system state contained in the backup set file: \recovery\backupsetfile.name. Replace all components; do not prompt for each component.

```
dsmc restore backupset c:\recovery\backupsetfile.name systemstate  
-loc=file -replace=all
```

Related information

“Restore data from a backup set” on page 206

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.

Syntax

```
➤➤ REStore GRoup options source destination ➤➤
```

Parameters

source

Specifies the virtual file space name (enclosed in braces) 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 94. Restore Group command: Related options

Option	Where to use
backupsetname "Backupsetname" on page 336	Command line only.
fromdate "Fromdate" on page 420	Command line only.
fromnode "Fromnode" on page 420	Command line only.
fromtime "Fromtime" on page 421	Command line only.
ifnewer "Ifnewer" on page 426	Command line only.
inactive "Inactive" on page 428	Command line only.
latest "Latest" on page 456	Command line only.
pick "Pick" on page 481	Command line only.
pitdate "Pitdate" on page 482	Command line only.
pittime "Pittime" on page 483	Command line only.
preservepath "Preservepath" on page 490	Command line only.
replace "Replace" on page 499	Client options file (dsm.opt) or command line.
showmembers "Showmembers" on page 526	Command line only.
skiptpermissions "Skiptpermissions" on page 528	Client options file (dsm.opt) or command line.
skiptsecuritycrc "Skiptsecuritycrc" on page 528	Client options file (dsm.opt) or command line.
subdir "Subdir" on page 548	Client options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 553	Client options file (dsm.opt) or command line.
todate "Todate" on page 564	Command line only.
totime "Totime" on page 565	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 739

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. The account that runs the backup-archive client must have administrator authority to successfully perform any type of image restore.
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 Tivoli Storage Manager API must be installed to use the **restore image** command.
- You can restore an NTFS or ReFS file system to a FAT32 volume or vice versa.
- The destination volume to which you restore must exist and be the same size or larger than the source volume.
- The physical layout of the target volume (striped, mirrored) can differ.

- The target volume is overwritten with data contained in the image backup.
- You do not have to format a target volume before you restore an image backup that contains a file system.
- The client requires an exclusive lock to destination volume you are restoring. The client locks, restores, unlocks, unmounts, and mounts the volume during a restore operation. During the restore process, the destination volume is not available to other applications.
- If you use the **pick** option, the following information is displayed for file system images that are backed up by the client:
 - Image Size
 - Stored Size - This value is the actual image size that is stored on the server. The **imagegapsize** option can be set so only used blocks in a file system are backed up. So, the stored image size on the server might be smaller than the volume size. For online image backups, the stored image can be larger than the file system based on the size of the cache files.
 - File system type
 - Backup date and time
 - Management class that is assigned to image backup
 - Whether the image backup is an active or inactive copy
 - The image name
- If a restored image is corrupted, use the **chkdsk** utility to check for and repair any bad sectors or data inconsistencies (unless the restored volume is RAW).

Supported Clients

This command is valid for all Windows clients.

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. You can restore an NTFS or ReFS file system to a FAT32 volume or vice versa.

Table 95. Restore Image command: Related options

Option	Where to use
backupsetname "Backupsetname" on page 336	Command line only.

Table 95. Restore Image command: Related options (continued)

Option	Where to use
dateformat "Dateformat" on page 356	Client option file (dsm.opt) or command line.
deletefiles "Deletefiles" on page 361	Command line only.
fromnode "Fromnode" on page 420	Command line only.
imagetofile "Imagetofile" on page 428	Command line only.
inactive "Inactive" on page 428	Command line only.
incremental "Incremental" on page 446	Command line only.
noprompt "Noprompt" on page 474	Command line only.
pick "Pick" on page 481	Command line only.
pitdate "Pitdate" on page 482	Command line only.
pittime "Pittime" on page 483	Command line only.
timeformat "Timeformat" on page 562	Client option file (dsm.opt) or command line.
verifyimage "Verifyimage" on page 575	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. The destination volume must be mapped to a drive letter. 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. If the destination volume is on a dynamic disk, the lost space can be recovered by increasing the size of the volume. Increasing the size of the volume also increases the size of the restored volume.

Examples

Task Restore the e: drive to its original location.

Command: `dsmc rest image e:`

Task Restore the h: drive to its original location and apply the changes from the last incremental backup of the original image that is recorded on the server. The changes include deletion of files.

Command: `dsmc restore image h: -incremental -deletefiles`

Task Restore the d: drive to its original location. Use the **verifyimage** option to enable detection of bad sectors on the target volume.

Command: `dsmc restore image d: -verifyimage`

Task If bad sectors present on the target volume, use the **imagetofile** option to restore the d: drive to the e:\diskD.img file to avoid data corruption.

Command: `dsmc restore image d: e:\diskD.img -imagetofile`

Task Restore the e: drive from the backup set `weekly_backup_data.12345678` to its original location.

Command: `restore image e:
-backupsetname=weekly_backup_data.12345678`

Related information

“Verifyimage” on page 575

“Imagetofile” on page 428

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 options file (`dsm.opt`). The value in the client 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.

A NAS file system specification uses the following conventions:

- Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.
- NAS file system designations on the command line require brace delimiters {} around the file system names, such as: {/vol/vol0}.

Supported Clients

This command is valid for all Windows clients.

Syntax

```

➔➔—REStore NAS—[—options—]—sourcefilespec—[—destinationfilespec—]➔➔

```

Parameters

sourcefilespec

Specifies the name of the NAS file system image you want to restore. This parameter is required unless you use the pick option to display a list of NAS images from which to choose. You cannot use wildcard characters when 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 96. Restore NAS command: Related options

Option	Where to use
dateformat "Dateformat" on page 356	Client option file (dsm.opt) or command line.
inactive "Inactive" on page 428	Command line only.
mode "Mode" on page 463	Command line only.
monitor "Monitor" on page 468	Command line only.
nasnodename "Nasnodename" on page 471	Client options file (dsm.opt) or command line.
numberformat "Numberformat" on page 476	Client option file (dsm.opt) or command line.
pick "Pick" on page 481	Command line only.

Table 96. Restore NAS command: Related options (continued)

Option	Where to use
pitdate "Pitdate" on page 482	Command line only.
pittime "Pittime" on page 483	Command line only.
timeformat "Timeformat" on page 562	Client 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 471

"Monitor" on page 468

"Cancel Process" on page 670

Restore Systemstate

Use the **restore systemstate** command to restore the backup of the system state stored on the Tivoli Storage Manager server, or inside a backup set from the Tivoli Storage Manager server, when the backupsetname option is specified.

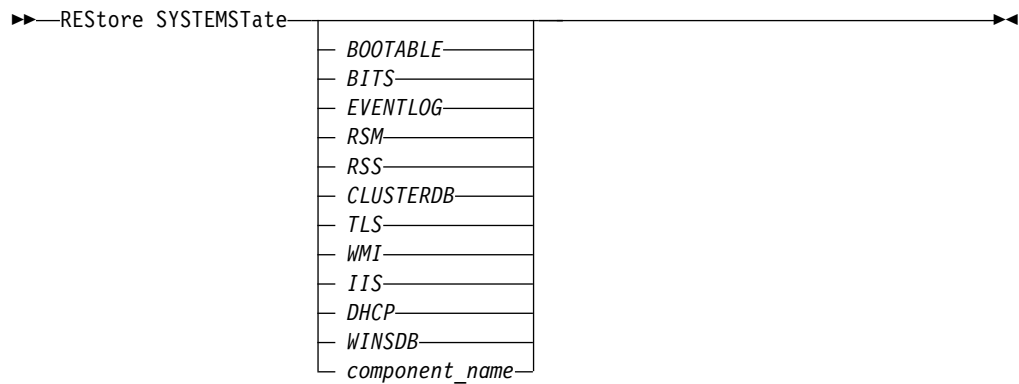
The Windows backup-archive client does not support the current commands for restoring individual startable system state components.

This command also restores ASR files for Windows clients that are running in a Windows Preinstallation Environment (WinPE); both BIOS and UEFI boot architectures are supported.

Supported Clients

This command is valid for all supported Windows clients.

Syntax



Parameters

Table 97. Restore Systemstate command: Related options

Option	Where to use
asrmode "Asrmode" on page 327	Command line only.
backupsetname "Backupsetname" on page 336	Command line only.
dateformat "Dateformat" on page 356	Client option file (dsm.opt) or command line.
numberformat "Numberformat" on page 476	Client option file (dsm.opt) or command line.
pitdate "Pitdate" on page 482	Command line only.
pittime "Pittime" on page 483	Command line only.
replace "Replace" on page 499	Client option file (dsm.opt) or command line.
timeformat "Timeformat" on page 562	Client option file (dsm.opt) or command line.

Examples

- Task** Restore the entire system state object.
Command: restore systemstate
- Task** Restore the startable system state component.
Command: restore systemstate bootable
- Task** Restore the system service WMI component.
Command: restore systemstate wmi
- Task** Restore the system state object contained within the backup set daily_backup_data.12345678. Replace all components; do not prompt for each component.
Command: restore systemstate
-backupsetname=daily_backup_data.12345678 -replace=all

Task Restore the WMI component contained within the backup set `daily_backup_data.12345678`.

Command: `restore systemstate wmi
-backupsetname=daily_backup_data.12345678 -replace=all`

Related information

“Recovering a computer when the Windows OS is not working” on page 202

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 Windows 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 98. Restore VAPP command: Related options

Option	Where to use
asnodename	Client user-options file (dsm.opt), 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 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 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 324

“Pick” on page 481

“Pitdate” on page 482

“Pittime” on page 483

“Vmchost” on page 583

“Vmcpw” on page 584

“Vmcuser” on page 586

Restore VM

Use the **restore vm** command to restore a virtual machine that was previously backed up.

The **restore VM** command can be used to restore both Microsoft Hyper-V virtual machines and VMware virtual machines. The information for each type of restore is presented in its own heading. If you are restoring a virtual machine that is part of a Hyper-V setup, you can skip over the *Restore VM for VMware virtual machines* text. If you are restoring a VMware virtual machine, you do not need to read the *Restore VM for Hyper-V virtual machines* text.

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.

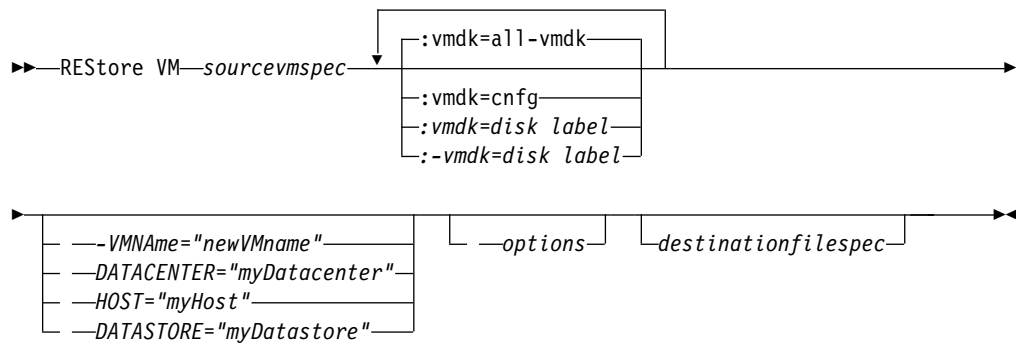
Backup management tags that were backed up with the run **backup vm** command are restored with the virtual machine. Backup management tags are used to exclude virtual machines from backups and to specify the retention policy of backups.

Full virtual machine backups that were previously created by using VMware Consolidated Backup (VCB) can still be restored by using the original VCB restore steps. To restore full virtual machine backups that were created by VCB, see “Restoring full VM backups that were created with VMware Consolidated Backup” on page 222. If you use VCB to restore a virtual machine, use the VMware converter program on the client to move the restored files to a running state on a VMware server. If the backup-archive client is running in a virtual machine, and if you performed a file-level backup of the virtual machine’s files, you can restore the backed-up files to the virtual machine by using the command-line interface or the Java GUI.

Supported Clients

This command is valid on supported Windows 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.

destinationfilespec

This parameter is for VMware VCB restores only. It specifies the location where VCB full virtual machine image files are restored. If this option is not specified the **vmbackdir** option is used.

: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 99. 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.
vmautostartvm This parameter is only valid when instantaccess is specified as the vmrestoretype value.	Command line or client options file.
vmbackdir	Command line or client options file.
vmbackuptype	Command line or client options file.
vmdiskprovision This parameter is only valid when instantrestore is specified for the vmrestoretype value.	Command line or client options file.
vmiscsiserveraddress This parameter is only valid when either instantaccess or instantrestore is specified for the vmrestoretype value.	Command line or client options file.
vmrestoretype	Command line.
vmtempdatastore This parameter is only valid when instantrestore is specified for the vmrestoretype value.	Command line or client options file.
vmvstortransport	Command line or client options file. This parameter is not valid for restoring VMware virtual machines that were backed up using VCB.

Examples

Task To perform an instant restore or instant access operation from the command line, see Scenarios for running instant access and instant restore from the backup-archive client command line.

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).

Use the **Restore VM** command to restore Hyper-V guests. You can restore Hyper-V guests to a local disk, a SAN-attached disk, to a cluster shared volume, or to a remote file server share. Remote file server shares must be on a Windows Server 2012 (or newer) system.

Even though the client shuts down the virtual machine before it deletes it, manually shutting down the virtual machine before you run the **Restore VM** command is a good practice to follow, to bring any in-progress application activities to an orderly stop. Then, use the **Restore VM** command to restore the virtual machine such that its content and configuration is identical to what it was when it was backed up.

This command is valid on supported Windows clients that are installed on a Hyper-V host system.

```

➤➤--REStore VM-- --sourcevmspec-----➤➤
➤└--vmname=--new_vm_name-- --targetpath=--path--┐┌-----┐
└-----┘└--options--┘└-----┘➤➤

```

Note: Any parameter that contains spaces must be enclosed in quotation (" ") marks.

Required. Specifies the name of the virtual machine that was backed up. The virtual machine name is case-sensitive. You cannot use wildcards in the virtual machine name.

Optional. This parameter is only valid for restoring virtual machines that were backed up by using either `mode=iffull` or `mode=ifincr`. Specify a new name for the virtual machine. The name can contain 1-100 characters, and the following characters are not valid: `\ / : ; , * ? " ' < > |`

If this parameter is not specified, the virtual machine is restored and it has the same name it had when it was backed up. If you do specify a new name for a virtual machine, a target path is required.

This parameter is required if you specify a new virtual machine name

(-vmname=). Use this parameter to restore the virtual machine to an alternative location. Specify the disk drive and path that you want to restore the virtual machine to.

Table 100. Restore VM command: Related options when restoring Hyper-V virtual machines

Option	Where to use
inactive	Command line
pick	Command line
pitdate	Command line
pittime	Command line
replace	Command line, client options file, or client preferences editor.
vmbackdir	Command line, client options file.
vmbackuptype	Command line or client options file. To restore Hyper-V virtual machines, this must be set to vmbackuptype hypervfull

Examples

Task Restore the most recent backup version of a virtual machine named VM1 to the drive and path it was in when it was backed up.

```
dsmc restore vm VM1
```

Task Restore the most recent backup version of a virtual machine named vm1 to the drive and path it was in when it was backed up. Replace the existing virtual machine without prompting.

```
dsmc restore vm vm1 -replace=yes
```

Task Restore the backed up virtual machine named VM1 to a new name (vm2):

```
dsmc restore vm VM1 -VmName=vm2
```

Task Restore the backed up virtual machine named vm1, and assign it a new name (vm2). Issue a prompt before overwriting vm2, if that virtual machine already exists.

```
dsmc restore vm vm1 -VmName=vm2 -replace=prompt
```

Task Restore the virtual machine named vm1 to a specific drive and path, without renaming the virtual machine:

```
dsmc restore vm vm1 -targetpath="E:\New Path"
```

Task Restore the virtual machine named vm1, but rename it as vm2 and restore it to a new path:

```
dsmc restore vm vm1 -VmName=vm2 -targetpath=F:\NewPath
```

Task Use -pick and -inactive to show active and inactive backups for a virtual machine that is named vm1. You choose the backup to restore from a list:

```
dsmc restore vm vm1 -pick -inactive
```

Related concepts:

“Virtual machine exclude options” on page 404

“Virtual machine include options” on page 437

Related tasks:

“Preparing the environment for full backups of VMware virtual machines” on page 167

Related reference:

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:

1. When a directory is retrieved, its modification date and time is set to the date and time of the retrieve, not to the date and time the directory had when it was archived. This is because the backup-archive client retrieves the directories first, then adds the files to the directories.
2. An error occurs if you attempt to retrieve a file whose name is the same as the short name of an existing file. For example, if you attempt to retrieve a file you specifically named ABCDEF~1.DOC into the same directory where a file named abcdefghijk.doc exists, the retrieve fails because the Windows operating system equates the file named abcdefghijk.doc to a short name of ABCDEF~1.DOC. The retrieve function treats this as a duplicate file.

If this error occurs, perform any of the following actions to correct it:

- Retrieve the file with the short file name you specified to a different location.
- Stop the retrieval, and change the name of the existing file.
- Disable the short file name support on Windows.
- Do not use file names that conflict with the short file naming convention. For example, do not use ABCDEF~1.DOC.

The workstation name is part of the file name. Therefore, if you archive files on one workstation and you want to retrieve them to another workstation, you must specify a destination. This is true even if you are retrieving to the same physical workstation, but the workstation has a new name. For example, to retrieve the c:\doc\h2.doc file to its original directory on the workstation, named star, you would enter:

```
dsmc retrieve c:\doc\h2.doc \\star\c$\
```

The workstation named star has been renamed and the new name is meteor. To retrieve the c:\doc\h2.doc file to meteor, you would enter:

```
dsmc retrieve c:\doc\h2.doc \\meteor\c$\
```

You could also enter:


```
dsmc retrieve c:\doc\h2.doc \\star\c$\
```

You can enter the command in either of the preceding ways because if the workstation name is not included in the specification, the local workstation is assumed (meteor, in this case).

Supported Clients

This command is valid for all clients.

Syntax



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.

Note: If you include *filespace name*, do not include a drive letter in the file specification.

{filespace name}

Specifies the file space (enclosed in braces) on the server that contains the files you want to retrieve. This name is the drive label on the workstation drive from which the files were archived.

Use the file space name if the drive label name has changed, or if you are retrieving files that were archived from another node that had drive label names that are different from yours.

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

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.

When you enter the *destinationfilespec* string, consider the following points:

- If the *sourcefilespec* names a single file, the *destinationfilespec* can be a file or a directory.
- If the *sourcefilespec* is wildcarded or if you specify the *subdir=yes* option, the *destinationfilespec* must be a directory and end with a directory delimiter (\).

Note: If the destination path or any part of it does not exist, the client creates it.

Table 101. Retrieve command: Related options

Option	Where to use
dateformat "Dateformat" on page 356	Client options file (dsm.opt) or command line.

Table 101. Retrieve command: Related options (continued)

Option	Where to use
description "Description" on page 362	Command line only.
dirsonly "Dirsonly" on page 368	Command line only.
filelist "Filelist" on page 414	Command line only.
filesonly "Filesonly" on page 418	Command line only
fromdate "Fromdate" on page 420	Command line only
fromnode "Fromnode" on page 420	Command line only.
fromtime "Fromtime" on page 421	Command line only
ifnewer "Ifnewer" on page 426	Command line only
pick "Pick" on page 481	Command line only.
preservepath "Preservepath" on page 490	Command line only.
replace "Replace" on page 499	Client options file (dsm.opt) or command line.
skiptpermissions "Skiptpermissions" on page 528	Client options file (dsm.opt) or command line
skiptsecuritycrc "Skiptsecuritycrc" on page 528	Client options file (dsm.opt) or command line
subdir "Subdir" on page 548	Client options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 553	Client options file (dsm.opt) or command line.
timeformat "Timeformat" on page 562	Client options file (dsm.opt) or command line.
todate "Todate" on page 564	Command line only.
totime "Totime" on page 565	Command line only.

Examples

Task Retrieve a single file named `budget.fin`.

```
ret c:\devel\projecta\budget.fin
```

Task Retrieve all files with an extension of `.c` from the `c:\devel\projecta` directory.

```
ret c:\devel\projecta\*.c
```

- Task** Retrieve all files with a file extension of .c from the \devel\projecta directory on the winnt file space.
- ```
ret {winnt}\devel\projecta*.c
```
- Task** Retrieve all files in the c:\devel directory.
- ```
ret c:\devel\*
```
- Task** Retrieve files from the abc file space proj directory.
- ```
ret {abc}\proj*.*
```
- Task** Retrieve all files with a file extension of .c from the c:\devel\projecta directory to the c:\newdevel\projectn\projecta directory. If the \projectn or the \projectn\projecta directory does not exist, it is created.
- ```
ret c:\devel\projecta\*.c c:\newdevel\projectn\
```
- Task** Retrieve files in the c:\project directory. Use the **pick** option.
- ```
ret c:\project* -pick
```
- Task** Retrieve a file from the renamed file space \\your-node\h\$\_OLD to its original location. Enter both the source and destination as follows:
- ```
ret \\your-node\h$_OLD\docs\myresume.doc h:\docs\
```

Related information

“Client options reference” on page 321

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.

If you want to retrieve archives from file spaces that were renamed by the Unicode-enabled client, you must specify the source on the server and a destination on the client. For example, you archived files from your H-disk, named \\your-node\h\$ prior to installing the client. After the installation, you issue the following archive command:

```
arc h:\logs\*.log
```

Before the archive takes place, the server renames the file space to \\your-node\h\$_OLD. The archive continues placing the data specified in the current operation into the Unicode-enabled file space named \\your-node\h\$. That file space now contains only the \logs directory and the *.log files. If you want to retrieve a file from the (old) *renamed* file space to its original location, you must enter both the source and destination as follows:

```
retrieve \\your-node\h$_OLD\docs\myresume.doc h:\docs\
```

Retrieve named streams

The backup-archive client retrieves named streams on a file basis only.

Directories in Windows systems can contain named streams. Named streams attached to a directory will always be overwritten (regardless of the value of the prompt option) during the retrieve.

Retrieve sparse files

When retrieving sparse files to a non-NTFS or non-ReFS file system, set the server communication time-out value (**IDLETIMEOUT**) to the maximum value of 255 to avoid client session timeout.

The following issues apply if more data is restored than the Microsoft disk quota allows:

- If the user who is performing the retrieve has a disk quota (for example, the user belongs to the Backup Operator Group), the backup-archive client does not retrieve any data that exceeds the disk quota of the retrieve user and displays a "Disk Full" message.
- If the user who is performing the retrieve does not have a disk quota (for example, the user belongs to the Administrator Group), the backup-archive client retrieves all data and transfers ownership of the files that exceed the disk quota of the original owner to the user who is performing the retrieve (in this case, the Administrator).

Schedule

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 `manageservices` option is set to `schedule` (does not apply to Mac OS X).
2. This command is valid only on the initial command line. It is not valid in interactive mode or in a macro file.

If the `schedmode` option is set to `polling`, the client scheduler contacts the server for scheduled events at the hourly interval you specified with the `querschedperiod` option in your client 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 options file (`dsm.opt`) or on the **schedule** command.

After you start the client scheduler, it continues to run and to start scheduled events until you press **Ctrl+Break**, restart the workstation, or turn off the workstation to end it.

Note: You *cannot* enter this command in interactive mode.

Supported Clients

This command is valid for all clients.

Syntax

```
►►—Schedule—┐
                └—options—┘
```

Parameters

Table 102. Schedule command: Related options

Option	Where to use
maxcmdretries "Maxcmdretries" on page 459	Client options file (dsm.opt) or command line.
password "Password" on page 478	client options file (dsm.opt)
querschedperiod "Querschedperiod" on page 494	Client options file (dsm.opt) or command line.
retryperiod "Retryperiod" on page 511	Client options file (dsm.opt) or command line.
schedlogname "Schedlogname" on page 517	Client options file (dsm.opt) or command line.
schedmode "Schedmode" on page 519	Client options file (dsm.opt) or command line.
sessioninitiation "Sessioninitiation" on page 523	Client options file (dsm.opt) or command line.
tcpclientport "Tcpclientport" on page 557	Client options file (dsm.opt) or command line.

Examples

Task Start the client scheduler.

Command: `dsmc sch -password=notell`

When you run the **schedule** command, all messages that regard scheduled work are sent to the `dsmsched.log` file or to the file you specify with the `schedlogname` option in your client options file (`dsm.opt`). If you do not specify a directory path with the file name in the `schedlogname` option, the `dsmsched.log` resides in the current working 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

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 c:\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.

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 `C:\proposals\drafts\ice.doc` `C:\proposals\drafts\fire.doc`, then `C:\proposals` and `C:\proposals\drafts` 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 (`C:\`).

You can specify as many file specifications as available resources or other operating system limits allow.

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 103. Selective command: Related options

Option	Where to use
changingretries "Changingretries" on page 340	Client options file (dsm.opt) or command line.
compressalways "Compressalways" on page 350	Client options file (dsm.opt) or command line.
compression "Compression" on page 351	Client options file (dsm.opt) or command line.
dironly "Dironly" on page 368	Command line only.
filelist "Filelist" on page 414	Command line only.
filesonly "Filesonly" on page 418	Command line only.
postsnapshotcmd "Postsnapshotcmd" on page 486	Client options file (dsm.opt) or with the include.fs option.
preservelastaccessdate "Preservelastaccessdate" on page 489	Client options file (dsm.opt) or command line.
presnapshotcmd "Presnapshotcmd" on page 492	Client options file (dsm.opt) or with the include.fs option.
skiptpermissions "Skiptpermissions" on page 528	Client options file (dsm.opt) or command line.
skiptsecuritycrc "Skiptsecuritycrc" on page 528	Client options file (dsm.opt) or command line.
snapshotproviderfs "Snapshotproviderfs" on page 535	Client options file (dsm.opt) or with the include.fs option.
snapshotroot "Snapshotroot" on page 537	Command line only.
subdir "Subdir" on page 548	Client options file (dsm.opt) or command line.
tapeprompt "Tapeprompt" on page 553	Client options file (dsm.opt) or command line.

Examples

Task Back up the proja.dev file in the c:\devel directory.

Command: sel c:\devel\proj.a.dev

Task Back up all files in the c:\devel directory whose file names begin with proj.

Command: sel c:\devel\proj*.*

Task Back up all files in the c:\devel directory whose file names begin with proj. Back up all files with a file extension of .fin in the c:\planning directory.

Command: sel c:\devel\proj* c:\planning*.fin

Task Assuming that you initiated a snapshot of the C:\ drive and mounted the snapshot as \\florence\c\$\snapshots\snapshot.0, run a selective backup of the c:\dir1\sub1 directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name C:\.

Command: dsmc sel c:\dir1\sub1* -subdir=yes -snapshotroot=\\florence\c\$\snapshots\snapshot.0

Related information

"Autofsrename" on page 333

"Include options" on page 430

"Subfilebackup" on page 549

Open file support

If open file support is configured, the backup-archive client performs a snapshot backup or archive of files that are locked (or "in use") by other applications.

Use VSS as the snapshot provider; set **snapshotproviderimage** or **snapshotproviderfs** to VSS.

Note:

1. You can use the include.fs option to set snapshot options on a per file system basis.
2. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with NTFS or ReFS file systems. This support includes SAN-attached volumes that meet these requirements.
3. If the client is unable to create a snapshot, failover to non-OFS backup occurs; the same backup support that would be done if the OFS feature was not configured.
4. To enable open file support in a cluster environment, all systems in the cluster should have the OFS feature configured.

Adaptive subfile backups

If you perform a backup over a network device with limited bandwidth, such as with a modem, you can reduce network traffic by using the subfilebackup option.

In order to use this feature, you must first specify the adaptive subfile backup options during a normal backup.

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.

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.

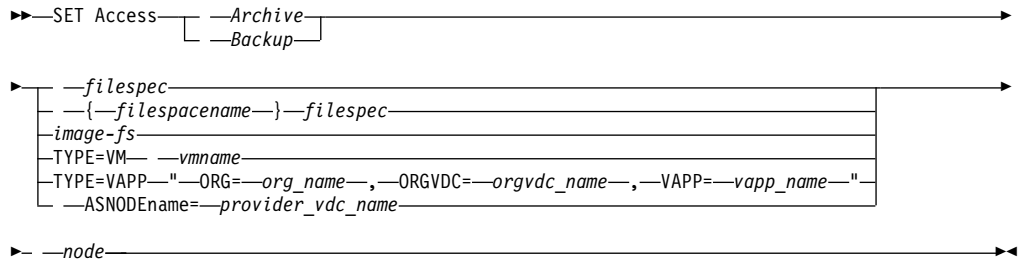
Note: You cannot give access to both archives and backups using a single command.

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 drive. However, you cannot use a wildcard to specify all drives. Use a single asterisk "*" for the file spec to give access to all files or images owned by you and backed up on the server. When the command `set access backup "*" node` is entered, no check is made with the server; it is assumed you have at least one object backed up.

If you give access to a branch of the current working directory, you only need to specify the branch. If you give access to objects that are not in a branch of the current working directory, you must specify the complete path. The file spec to which you gave access must have at least one backup version or archive copy object (file or directory) on the server.

To specify all files in a named directory, enter `d:\test\mine\proj1*` on the command line.

To give access to all objects below a certain level, use an asterisk, directory delimiter, and an asterisk at the end of your file spec. For example, to give access to all objects below `d:\test` use file spec `d:\test**`.

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.

Note:

1. Use the file space name if the drive label name has changed.

2. If you include *filespace*name, do not include a drive letter in the file specification.

For example:

- Your directory structure is multilevel: d:\test\sub1\subsub1.
- The d:\test directory contains the h1.txt and h2.txt files.
- The d:\test\sub1 directory contains file s1.htm.
- The d:\test\sub1\sub2 directory contains the ss1.cpp file.

To allow access to all files in the d:\test\sub1\sub2 directory, enter:

```
set access backup d:\test\sub1\sub2\* * *
```

To allow access to only those files in the d:\test directory, enter:

```
set access backup d:\test\* * *
```

To allow access to all files in all directories in and below the d:\test directory, enter:

```
set access backup d:\test\* * *  
set access backup d:\test\*\* * * *
```

{filespacename}

Specifies the file space name (enclosed in braces) on the server that contains the files to which you are giving access. This name is the drive label name on the workstation drive from which the file was backed up or archived. Use the file space name if the drive label name has changed.

image-fs

The name of the image file system to be shared. This can be specified as an asterisk (*) to allow access to all images owned by the user granting access.

-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 -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.

Examples

Task Give 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 -ASNODE=ABC
```

Task Give the user at node_2 authority to restore all files with an extension of .c from the c:\devel\proja directory.

```
set access backup c:\devel\proja\*.c node_2
```

Task Give the user at node_3 authority to retrieve all files in the c:\devel directory, but do not permit access to files in subdirectories of c:\devel, such as c:\devel\proj.

```
set access archive c:\devel\* node_3
```

Task Give all nodes whose names end with bldgb the authority to restore all backup versions from all directories on the d: drive. The d: drive has the file space name of project.

```
set ac b {project}\*\* *bldgb
```

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 | Activateretention --<filespec>
► --filelist=<filespec> --description=--pick
```

Parameters

TYPE=

Specifies the event type setting. This parameter must be specified.

hold

Prevents the object from being deleted regardless of expiration policy.

release

Allows normal event-controlled expiration to take place.

activateretention

Signals the server that the controlling event 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 \\user\c\$\tsm521\debug\bin\winnt_unicode\dsm.opt, with objects rebound (as opposed to archived or some other notation).

```
Rebinding--> 274 \\user\c$\tsm521\debug\
bin\winnt_unicode\dsm.opt
Rebinding--> 290 \\user\c$\tsm521\debug\
bin\winnt_unicode\dsm.opt
```

```
Total number of objects inspected:      2
Total number of objects archived:       0
Total number of objects updated:        0
Total number of objects rebound:       2
Total number of objects deleted:        0
```

```

Total number of objects expired:      0
Total number of objects failed:      0
Total number of bytes transferred:    0 B
Data transfer time:                   0.00 sec
Network data transfer rate:           0.00 KB/sec
Aggregate data transfer rate:         0.00 KB/sec
Objects compressed by:                0%
Elapsed processing time:               00:00:02

```

Task The `-pick` option used with the `set event` command `set event type=activate \user\c$\tsm521\common\winnt` 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\c\$\tsm521 \common\winnt
2.	08/01/2003 10:38:11	766 B	\\user\c\$\tsm521 \common\winnt
3.	08/05/2003 08:47:46	5.79 KB	\\user\c\$\tsm521 \common\winnt
4.	08/01/2003 10:38:11	5.79 KB	\\user\c\$\tsm521 \common\winnt
5.	08/05/2003 08:47:46	10.18 KB	\\user\c\$\tsm521 \common\winnt

Related information

“Dateformat” on page 356

“Numberformat” on page 476

“Noprompt” on page 474

“Subdir” on page 548

“Timeformat” on page 562

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 Windows clients that perform 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 85

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 Windows systems:

Enclose the command parameters in quotation marks (").

Command line example:

dsmc set password "t67@#\$\$%^&" "pass2><w0rd"


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=DOMAIN** to set the Windows domain administrator credentials to enable users to log in to a remote Windows proxy node (the file restore interface), for file restore operations. This option requires a Tivoli Storage Manager for Virtual Environments: Data Protection for VMware license.

Use the following format for the **set password -type=domain** command:

```
set password -type=domain -validate administrator_name password
```

where:

VALIDate

Validates the Windows domain administrator credentials before the credentials are stored in the Windows registry. If the validation fails, the credentials are not stored in the registry, and users are not able to log in to the file restore interface. The **validate** parameter is valid only with the **TYPE=DOMAIN** parameter.

administrator_name

Specifies the account name of a domain administrator. The account name must contain the Windows domain name and the administrator ID. The account name must be in the following format:

domain_name\administrator_ID

password

Specifies the password that is associated with the specified domain administrator account.

For more information about configuration requirements for remote mount proxy nodes, see the Tivoli Storage Manager for Virtual Environments: Data Protection for VMware documentation.

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 TSM.PWD 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 **vmchost** option. That is, if **vmchost** 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 **vmchost**, **vmcuser**, and **vmcpw** options.

You can also set the **vmchost** 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
myvmchost.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$word"
```

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 Set up the Windows domain administrator credentials that are necessary for users to log in to the file restore interface and save the Windows domain credentials to the registry. In this example, the Windows domain in which all user accounts are registered is called example_domain. Kev_the_admin is the Windows domain administrator ID and pas\$word! is the corresponding password for the administrator.

```
dsmc set password -type=domain -val "example_domain\Kev_the_admin"  
"pas$word!"
```

Related reference:

"Snapdiff" on page 529

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 higher.
- The vmhost 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 vmhost 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 on supported Windows 64-bit 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 268

| **Related reference:**

| “Vmchost” on page 583

| “Vmtagdatamover” on page 612

| **“Set Password”** on page 781

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

Numerics

128-bit AES encryption support 134

256-bit AES encryption support 134

A

absolute mode 272

absolute option 322

accessibility features 791

active backup versions

displaying 118, 194, 700

restoring 194

Active Directory

query for communication method and server with which to connect 568

active directory objects

modifying CAD and agent services 228

restore restrictions and limitations 226

restoring 224, 226, 228

restoring from system state backup 224

restoring using GUI and command line 225

adaptive subfile backup 549

excluding files from 400

including files for 431

restore permissions 108

specifying client cache path 550

specifying client cache size 551

adlocation option 323

administrative client

allowing secure sessions within a private network 554

archive

assign description to 239, 242, 362

associating local snapshot with server file space 242, 537

binding management class to 242, 323

case sensitive file name conflicts 339

command 633

compressing files 351

configuring client-node proxy support 242

copy group attributes 269

copy mode 272

delete files after 361

delete individual archives from server file space 244, 672

deleting file spaces 232

directory tree 240

display the last modification date and last access

datecreation date 697

files only 570

grace period retention 268

how managed 267

improving speed using share memory 297

include files for 431

information, query 697

list of files 242, 414

more than one file specification 242

number of attempts to archive open files 340

only files; not directories 242

overriding management class during 274

packages 239

primary tasks 239

process directories only (not files) 368

archive (*continued*)

query user access 695

retrieving

using command line 245

shared data on multiple clients under a single node

name 242, 324

specifying whether to include subdirectories for 242

starting a web client session 115

summary of options 297

suppress confirmation prompt before deleting 474

using commands 241

archive copy group 268

archive fastback

command 636

archive maximum file size 135

archmc option 323

asnodename option 324

asnodename session settings 326

ASR

backup system drive

Windows 152

dsm.opt

Windows 151

preparation

Windows 151

restore procedure

Windows 202

WinPE CD

Windows 202

asrmode option 327, 620

attributes

preserving in tombstone objects 227

auditlogging option 328

auditlogname option 330

authentication

Tivoli Storage Manager client 108

authorization

authorizing user to restore or retrieve your files 229

options 311

auto-update 332

autodeploy option 332

autofsrename option 333

automated client failover

configuration and use 60

configuring 63

determining the replication status 65

force failover 67

other components 63

overview 60

preventing 66

requirements 61

restoring data 228

restrictions 62

retrieving data 228

testing the connection 67

Automated System Recovery (ASR)

back up 150

automating backup services

displaying scheduled work 254, 256

options for 260

process commands after backup 484

- automating backup services (*continued*)
 - process commands before backup 487
 - starting client scheduler 39

B

- back up 149
 - configuring support for client node proxy backups 145
 - include-exclude list 134
 - Microsoft Dfs trees and files 190
 - network-attached storage (NAS) 648
 - new or changed files 136
 - NTFS file spaces 188
 - number of attempts to back up open files 340
 - parallel 596, 597, 598
 - ReFS file spaces 188
 - removable media using drive label 187
 - VM templates 592
 - with limited bandwidth 133
- backing up
 - in parallel sessions 174
- Backing up cluster groups 71
- backing up data 167
- backing up network shares 109
- backing up virtual machines on a Hyper-V system 178
- backing up volume mount points
 - NTFS 688
 - ReFS 688
- backmc option 335
- backup 125
 - automating using Client Service Configuration Utility 280
 - case sensitive file name conflicts 339
 - copy mode 272
 - displaying processing status 182
 - excluding domains 127
 - excluding system state object from 149
 - file-level VMware 173
 - grace period retention 268
 - image 153
 - client domain 374, 378
 - with incremental backup 647
 - image, offline and online 153
 - improving speed using share memory 297
 - incremental
 - associating local snapshot with server file space 690
 - client command line 127
 - command line 127
 - incremental-by-date
 - client command line 127
 - command line 127
 - multi-session, send files contiguously to the server 346
 - Net Appliance CIFS share definitions 182
 - one server session per file specification 346
 - overview 121
 - primary tasks 121
 - process directories only (not files) 368
 - query user access 695
 - selective
 - associating local snapshot with server file space 775
 - back up list of files 127
 - client command line 127
 - command line 127
 - shared data on multiple clients under a single node
 - name 324
 - starting a web client session 115
 - subdirectories 127
 - summary of options 297

- backup (*continued*)
 - system state, assigning management class 149
- backup and restore
 - NAS file servers using CIFS 163
- backup chain integrity checks 615, 617
- backup comparison: incremental, journal-based, incremental-by-date 141
- backup copy group 122, 268
 - attributes 269
- backup fastback command 639
- backup files
 - assigning management class 273
- backup group command 642
- backup image command 644
- backup maximum file size 135
- backup nas command 648
- Backup NTFS or ReFS data on mounted volumes 689
- backup operators group
 - required user security rights for backup and restore 105
- backup planning 121
- backup set
 - enabling GUI for local restore 206
 - restore 193, 206
 - restore systemstate from 739, 744
 - restoring in a SAN environment 743
- backup sets
 - restore considerations 208, 742
- backup systemstate command 650
- backup vapp command 652
- backup vm command 658
- backup with client node proxy
 - agent node 145
 - overview 145
 - target node 145
- backup-archive client
 - overview 1
- backup-archive scheduler service
 - install 279
- backupset
 - enabling GUI for local restore of 457
- backupsetname option 336
- basesnapshotname option 337
- batch mode 628
 - starting a session 112
- bottom-up processing
 - include-exclude list 102
- bypass processing of Windows file system security information 528

C

- c-mode 85
- C++ redistributable files
 - forced reboot 10
- cadlistenonport option 338
- cancel process command 670
- cancel restore command 671
- casesensitiveaware option 339
- central scheduling
 - summary of options 309
- Certificate Authorities
 - root certificates 45
 - Certificate Authorities 45
- changingretries option 340
- class option 341
- classic (standard) restore 198

- client
 - automatic update 3
 - client TCP/IP address other than the one for first server
 - contact 557
 - client TCP/IP port number other than the one for first server
 - contact 558
 - registering with server 91
 - setting password 91
 - size for the TCP/IP sliding window for client node 561
- client acceptor daemon
 - manage scheduler, web client, or both 457
- client acceptor service
 - configuring to manage scheduler 37
- client command options
 - overview 630
- client components
 - Windows client 6
- client management service 26
- client node proxy archive
 - overview 242
- client node proxy backup
 - scheduling 147
- client options
 - display current settings 715
 - exclude
 - exclude.archive 93
 - exclude.backup 93
 - exclude.compression 93
 - exclude.dir 93
 - exclude.file 93
 - exclude.file.backup 93
 - exclude.image 93
 - exclude.subfile 93
 - exclude.systemobject 93
 - guitreeviewafterbackup 422
 - order of processing (precedence) 314
 - overriding using command line 314
 - overview 630
 - using with commands 314
- client options file
 - creating and modifying 29
 - generating in shared directory 31
 - overview 27
 - required options for 29
 - specifying include-exclude options 92
- client options reference 321
- client scheduler
 - displaying scheduled work 254, 256
 - options for 260
 - run at startup 252
 - starting 39, 770
 - starting automatically 116
- Client Service Configuration Utility
 - commands for installing client services 283
 - configure CAD to manage existing scheduler service 282
 - create new scheduler and associate CAD to manage scheduler 282
 - options for installing client services 290
 - using to automate backups 280
- client services
 - considerations 105
- client user-options
 - creating multiple files 31
- client user-options file
 - overriding using commands 314
- client-node proxy support 145, 242
- client-server communication
 - client TCP/IP address other than the one for first server
 - contact 557
 - client TCP/IP port number other than the one for first server
 - contact 558
 - establishing 29
 - identify your workstation to the server 472
 - maximum disk I/O buffer size client uses when reading or writing files 369
 - method 347
 - name of a named pipe 471
 - query Active Directory for communication method and server with which to connect 568
 - reconnection attempts after failure 348
 - reconnection interval after failure 349
 - size for the TCP/IP sliding window for client node 561
 - size of internal TCP/IP communication buffer 555
 - specifying number of kilobytes client buffers before sending transaction to server 566
 - TCP/IP address for dsmcad 556
 - TCP/IP address of Tivoli Storage Manager server 560
 - TCP/IP port address of Tivoli Storage Manager server 559
 - TCP/IP port address on which to establish shared memory connection 525
 - whether to send small transactions to server without buffering them first 558
- clientview option 342
- closed registration
 - permissions 91
 - using 91
- Cluster configuration wizard 71
- cluster drives
 - enabling management of 345
- Cluster groups
 - backing up 71
- cluster resources
 - permissions 108
- clusterdisksonly option 342
- clustered data ONTAP 85
- clusternode option 345
- collecting diagnostic information 26
- collocatebyfilespec option 346
- command line
 - assigning description to archive 242
 - display current settings for client options 715
 - displaying
 - Euro characters in prompt 113
 - processing status 182
 - enabling 8.3 short names 386
 - ending a session 119
 - entering commands 630
 - general rules when entering options with commands 314
 - NAS file systems backup 161
 - overriding management class during archive 274
 - overview of parameters 630
 - performing image backup 158
 - performing point-in-time restore 233
 - restoring files and directories 195
 - restrictions for NAS file systems 159
 - return codes for operations 265
 - specifying file specification 631
 - specifying options file during session 478
 - starting a session 111
 - using wildcard characters 632
- command parameters
 - overview 630

- command processing, summary of options 311
- command session
 - ending 628
 - starting 628
- command-line prompt
 - displaying
 - Euro characters 113
- commands
 - archive 633
 - archive fastback 636
 - backup fastback 639
 - backup group 642
 - backup image 644
 - backup nas 648
 - backup systemstate 650
 - backup vapp 652
 - backup vm 658
 - batch mode 628
 - cancel process 670
 - cancel restore 671
 - delete access 671
 - delete archive 672
 - delete backup 674
 - delete filespace 677
 - delete group 678
 - entering 630
 - entering on command line 630
 - expire 680
 - general rules when entering options with 314
 - help 681
 - incremental 683
 - interactive (loop) mode 629
 - loop 690
 - macro 692
 - maximum file specifications permitted 631
 - monitor process 692
 - overview of parameters 630
 - preview archive 693
 - preview backup 694
 - query access 695
 - query adobjects 696
 - query archive 697
 - query backup 700
 - query backupset 703, 705
 - query filespace 707
 - query group 709
 - query image 710
 - query indexcl 712
 - query mgmtclass 713
 - query node 714
 - query options 715
 - query restore 716
 - query schedule 717
 - query session 717
 - query systeminfo 718
 - query systemstate 720
 - query vapp 721
 - query VM 725
 - restart restore 729
 - restore 730
 - restore adobjects 738
 - restore backupset 739, 744
 - restore backupset considerations 208, 742
 - restore group 746
 - restore image 748
 - restore NAS 751
 - restore systemstate 753

- commands (*continued*)
 - restore vapp 755
 - restore vm 758
 - retrieve 766
 - schedule 770
 - scheduled, enabling or disabling 260
 - selective backup 771
 - set access 775
 - set event 778
 - set netappsvm 780
 - set password 781
 - set vmtags 787
 - specifying file specification 631
 - using 625
 - using in executables 265
 - using options with 314
 - using wildcard characters 632
- commmethod option 347
- commrestartduration option 348
- commrestartinterval option 349
- communication methods
 - installable software 6
 - Shared Memory
 - Windows client 6
 - summary 296
 - TCP/IP
 - Windowss client 6
- communications
 - establishing through firewall 39
 - establishing with Secure Sockets Layer (SSL) 42
- compressalways option 350
- compression
 - disabling processing 435
 - enabling processing 435
 - include-exclude statements 435
- compression option 351
- compression processing
 - exclude from backup 435
 - exclude options 435
 - include files for 431
- compression, encryption, and subfile backup processing
 - back up 435
 - exclude from backup 435
 - exclude options 435
- concurrent backups 174
- configure
 - language for Java GUI 33
- configure the client for data deduplication 57
- configuring
 - client acceptor-managed scheduler 37
 - journal engine service 45
 - open file support 83
 - optional tasks 27
 - required tasks 27
 - the client scheduler 36
- Configuring cluster protection 71
- configuring the web client 34
- console option 352
- console window
 - displaying
 - Euro characters 113
- containing quotation marks 114
- control files 578
- copy destination attribute 272
- copy frequency attribute 270
- copy group name attribute 269
- copy groups 268

copy groups *(continued)*

archive 268

backup 268

copy mode parameter

absolute 272

modified 272

copy serialization attribute 271

copy type attribute 270

createnewbase 354

createnewbase option 354

D

data deduplication 53

data deduplication client configuration 57

data deduplication files

exclude 59

datacenter option 356

datastore option 356

date format

specifying 356

dateformat option 356

dedupcachepath option 359

dedupcachesize option 359

deduplication option 360

default domain

excluding domains from backup 127, 371

default management class 267

default policy domain 267

delete

file space 232

NAS or client objects 341

delete access command 671

delete archive command 672

delete backup command 674

delete group command 678

delete individual backups from server file space 129

deleted file systems 187

deletefiles option 361

deleting

authorizations 229

individual archives from server file space 244, 672

individual backups from server file space 674

deleting file spaces 232

description option 362

detail option 363

diagnostics

options 313

diffsnapshot option 364

diffsnapshotname option 366

directories

assigning management class for 367

excluding 93

excluding from backup processing 400

incremental backup processing overview 136

processing during incremental-by-date 140

restoring from command line 195

restoring from GUI 195

specifying on command line 631

directory

archiving 240

dirmc option 367

dirsonly option 368

disability 791

disablenqr option 368

discretionary access control list (permissions)

back up 188

disk space requirements

client 5

Windows client 6

diskbuffsize option 369

diskcachelocation option 370

displaying

archive information 697

online help 119

restartable restore sessions 716

scheduled events 717

session information 717

domain

back up using the GUI 126

include for full vm backups 381

include for image backup 374, 378

include for incremental backup 371

include for NAS image backup 375

specifying drives in the default 126

domain list

using universal naming convention to specify 188

domain option 371

domain.image option 374, 378

domain.nas option 375

domain.vcd option 376

domain.vmfull option 381

downloading maintenance updates 26

drive label

using to backup removable media 187

DSM_CONFIG environment variable 32

DSM_DIR environment variable 32

DSM_LOG environment variable 32

dsm.opt file

creating and modifying 29

required options for 29

specifying a drive specification using wildcards 98

dsm.smp file

copying to dsm.opt 29

location 29

dsmc command

using options 113

dsmcutil utility

commands for installing client services 283

options for installing client services 290

overview 283

dsmerlog.pru file 396

dsmerror.log file 396

dsmsched.log 515, 518

dsmwebcl.log 515, 518

duplicate file names

avoiding 193

dynamic and shared serialization 271

E

enable8dot3namesupport option 386

enablearchiveretentionprotection option 388

enablededupcache option 389

enableinstrumentation option 390

enablelanfree option 392

encrypting data during archive 134

encrypting data during backup 134

encryption

multiple clients under a single node name 324

of file data 134

saving encryption key password 394

encryption processing

determine encryption cipher used in current session 134

- encryption processing (*continued*)
 - excluding files from 400
 - include files for 431
 - query systeminfo command 718
- encryptiontype option 393
- encryptkey option
 - encryptkey=generate 394
 - encryptkey=prompt
 - encryptkey=save 394
- enhanced query schedule 254
- enhanced query schedule command 717
- environment variables 32
- error log
 - pruning 398
 - specifying path and file name 397
- error processing, summary of options 312
- error recovery
 - VMware virtual machines 219
- errorlogmax option 396
- errorlogname option 397
- errorlogretention option 396, 398
- event logging
 - scheduler 256
- event-based policy retention protection
 - archive 277
 - backup 277
- exclude
 - EXCLUDE.VMDISK 405
- exclude data deduplication files 59
- exclude options 400
 - exclude.archive 93
 - exclude.backup 93
 - exclude.compression 93
 - exclude.dir 93
 - exclude.file 93
 - exclude.file.backup 93
 - exclude.image 93
 - exclude.subfile 93
 - exclude.systemobject 93
 - preview 101
 - processing 102
 - wildcard characters 98, 100
- exclude.image option 93
- EXCLUDE.VMDISK 405
- excluding files
 - remotely accessed 97
 - system files 96
 - using wildcard characters 100
 - wildcard characters 98
- excluding system objects 149
- executable file
 - return codes from 265
- expire command 680

F

- failover
 - client 60
 - configuration and use 60
 - configuring the client 63
 - determining the replication status 65
 - disabling 66
 - other components 63
 - requirements 61
 - restore 228
 - restrictions 62
 - retrieve 228

- fbbranch option 406
- fbclient option 407
- fbpolicyname option 409
- fbreposlocation option 410
- fbserver option 411
- fbvolumename option 413
- features
 - Windows client 6
- file
 - restoring active and inactive versions 736
- file names
 - avoiding duplicate 193
- file space
 - delete 232, 677
 - determining fsID 363
 - excluding 93
 - NAS or client objects 341
 - performing an image backup 644
- file specification
 - maximum allowed on commands 631
- file systems
 - deleted 187
 - excluding from backup processing 400
 - GPFS, multinode cluster environment 371
 - image backup of 153
- file-level VM backup
 - restore 219
- filelist option 414
- filename option 417
- files
 - archive
 - directory tree 240
 - archive a list of 242, 414
 - archived, overriding management class 274
 - archiving 633
 - archiving more than one file specification 242
 - assigning management classes 186
 - backing up Microsoft Dfs 190
 - backing up open 185
 - binding management classes to 275
 - compressing during archive or backup 351
 - definition of changed 136
 - delete after archive 361
 - delete individual archives from server file space 244, 672
 - delete individual backups from server file space 674
 - encryption 134
 - excluding groups 98, 100
 - include-exclude
 - creating in Unicode format 430
 - including groups 98, 100
 - managing growth during compression 350
 - maximum file size for operations 135
 - processing include-exclude 102
 - query archive information 697
 - query backup information 700
 - query user access 695
 - renaming file spaces that are not Unicode to
 - Unicode-enabled 333, 683, 771
 - reset Windows archive attribute after backup 507
 - restoring files belonging to another node 231
 - restoring from command line 195
 - restoring from GUI 195
 - restoring to another workstation 232
 - retrieving
 - archives using command line 245
 - files belonging to another node 231
 - to another workstation 232

- files (*continued*)
 - sorting list of 118
- filesonly option 418
- firewall
 - establishing communications through 39, 423, 559
 - specifying TCP/IP ports for the web client 622
 - using web client through 622
 - whether server or client initiates sessions through 524
- fixed drives
 - backing up 188
- force incremental backup 322
- forcefailover option 419
- format and language
 - summary of options 310
- fromdate option 420
- fromnode option 420
- fromtime option 421
- full backups, creating 170
- full incremental
 - comparing with journal-based, incremental-by-date 141
 - description 136
 - when to use 141
- full VM backup
 - restore 213
 - full VM backup 213
- fuzzy backup 271

G

- getting started
 - changing your password 105
 - client scheduler 105
 - command-line session 105
 - displaying online help 105
 - ending a session 105
 - GUI session 105
 - sorting file lists 105
 - web client session 105
- GPFS file system
 - multinode cluster environment 371
- graphical user interface
 - changing password 116
 - delete individual files or images from server file space 674
 - displaying active and inactive backup versions 118, 194
 - displaying online help 119
 - displaying processing status 182
 - enabling for local backupset restore 457
 - enabling local backup set 206
 - ending a session 119
 - performing image backup 157
 - restore files and directories 195
 - starting a session 110
 - to back up objects 125
- group backup
 - display active and inactive objects 428
 - display all members of 526
 - overview 144
 - specify name of group 422
 - specify virtual file space name for 576
 - specifying full or differential 463
- groupname option 422
- GUI
 - ending a session 119
 - overriding management class during archive 274
 - performing point-in-time restore 233
 - starting a session 110

- guitreeviewafterbackup option 422

H

- hardware requirements
 - Windows client 6
- help
 - displaying online 119
 - Internet resources 119
 - online forum 119
 - service and technical support 119
- help command 681
- host option 423
- httpport option 423
- Hyper-V
 - backing up virtual machines 178
 - backup limitations 180

I

- IBM Knowledge Center xiii
- ieobjtype option 425
- ifnewer option 426
- image
 - restoring 203
 - using chkdsk to repair 203
 - using chkdsk tool to repair 748
 - using fsck to repair 203, 748
- image backup
 - configuring online image backup 82
 - considerations 154
 - deleting 674
 - excluding files from 400
 - file systems or logical volumes 644
 - include files for; assign management class to 431
 - include.dedup 431
 - incremental-by-date image backup 157
 - offline and online 153
 - perform 153
 - point-in-time restore 647
 - revoke access 671
 - specifying selective or incremental 463
 - using command line 158
 - using the GUI 157
 - using with file system incremental 157
 - using with incremental-by-date 156
 - with incremental backup 155, 647
- image backup, considerations 154
- imagegapsize option 427
- imagetofile option 428
- inactive backup versions
 - displaying 118, 194, 700
 - restoring 194
- inactive option 428
- inclexcl option 429
- include option
 - management class 273
 - processing 102
 - wildcard characters 98, 100
- include VM templates in back ups 592
- include-exclude list
 - creating 92
 - preview 101
 - query order of processing 712
 - size restriction 102
 - to control processing 134

- include-exclude options file
 - specifying path and file name of 429
 - Unicode-enabled file spaces 429
- include-exclude processing
 - options for 93
 - overview 93
- include.vm option 438
- include.vmsnapshotattempts option 442
- include.vmtsmvss option 444
- incrbydate option 445
- incremental and selective commands with snapshotroot
 - option 149
- incremental backup
 - associating local snapshot with server file space 537
 - back up new and changed files with modification date later than last backup 445
 - by date 127
 - client command line 127
 - client domain 371
 - client Java GUI 125
 - command line 127
 - description 136
 - directories, processing overview 136
 - GPFS, multinode cluster environment 371
 - memory-conserving algorithm 462
 - new and changed files 136
 - new and changed files with modification date later than last backup 445
 - of directories
 - processing overview 136
 - process a list of files 414
 - with image backup 155, 647
- incremental command 683
 - journal-based backup 687
- incremental option 446
- incremental-by-date
 - command line 127
 - comparing with incremental, journal-based 141
 - description 140
 - of directories
 - processing overview 140
 - when to use 141
- incremental-by-date backup 140
 - client command line 127
 - client Java GUI 125
 - using with image backup 156
- incrthreshold option 447
- input strings
 - containing blanks 114
- install
 - backup-archive scheduler service 279
- installation requirements
 - client 5
- installation types for the Windows client 10
- installing
 - overview 1
- installing from DVD 9
- installing the client management service 26
- installing the Windows client 10
- installing Tivoli Storage Manager
 - from DVD 9
 - silent installation 19
- instant access scenario 214
- instant restore scenario 214
- instrlogmax option 448
- instrlogname option 449

- instrumentation log
 - collecting performance information 390
 - controlling the size 448
 - specifying path and file name to store performance information 449
- interactive mode 629
- interactive session
 - ending 690
 - starting 112, 690
 - using 690

J

- Java GUI
 - configuration restrictions 111
- journal based backup
 - restoring 140
- journal based backups
 - restoring 140
- journal database files
 - errorlog 47
 - journaldir 47
 - NlsRepos 47
- journal engine service
 - configuring 45
- journal-based backup 138, 687
 - comparing with incremental, incremental-by-date 141
 - excluding directories 95
 - excluding files 95
 - include-exclude options
 - journal-based backup 95
 - performing traditional full incremental, instead of 474, 687
 - specifying how to respond to unsuccessful expire of object 447
 - when to use 141
- journaled file space
 - specifying directories with active files to expire 447
- journalpipe 47
- journalpipe option 450
- JournalSettings stanza 47

K

- keyboard 791
- Knowledge Center xiii

L

- LAN-based image backup
 - online and offline image backup 644
- LAN-free data movement 392
 - enabling communications for 130, 451, 453
 - options 131
 - prerequisites 130
 - shared memory port for 452
- lanfreecommmethod option 451
- lanfreeshmport option 452
- lanfreessl option 453
- lanfreetcppport option 453
- lanfreetcpserversaddress option 454
- language for Java GUI
 - configure 33
- language option 455

- last access date
 - specifying whether to update during backup or archive 136, 489
- latest option 456
- limitations
 - Hyper-V support 180
- limited bandwidth
 - back up 133
- local backup set
 - enabling GUI for local restore 206
- local snapshot
 - associating local snapshot with server file space 149
- localbackupset option 457
- log
 - See also* schedule log
 - controlling the size 448
 - DSM_LOG environment variable 397, 449, 517
 - error log, pruning 396
 - errorlogname option 397
 - errorlogretention option 397
 - instrlogmax option 448
 - intrlogname option 449
 - schedlogname option 517, 770
 - schedlogretention option 517, 770
 - specifying path and file name 397, 449, 517, 770
 - web client 515
- logical volume
 - image backup of 153
 - restoring 203
- logs
 - dsmsched.log 518
 - dsmsched.pru 518
 - dsmwebcl.log 518
 - dsmwebcl.pru 518
 - truncating application logs 444
- loop command 690

M

- macro command 692
- maintenance 332
 - auto-update 3
- manageservices option 457
- management class
 - assigning 186
- management classes 134
 - assigning to directories 274, 367
 - assigning to files 273
 - binding archive files to 242
 - binding to files 275
 - default 268
 - displaying 269
 - displaying information about 713
 - overriding during archive processing 274
 - overriding the default 273
 - processing 273
 - questions to consider 273
 - selecting for files 273
 - specifying with include option 273
 - using management class, example 273
- maxcmdretries option 459
- mbobjrefreshthresh 460
- mbpctrefreshthresh 461
- memory requirements
 - Windows client 6
- memoryefficientbackup option 462

- messages
 - displaying on screen 574
 - specifying language type 455
 - stop displaying 496
- Microsoft Cluster Server cluster nodes
 - FAQs 79
 - installing scheduler service 70
 - installing Tivoli Storage Manager 70, 79
- Microsoft Dfs junction
 - restore 736
- Microsoft Dfs trees and files
 - back up 190
- migrating backup-archive clients 1
- migration
 - web client 2
 - web client language files 2
- mobile dial-up support 116
- mode option 463
- modes
 - batch 628
 - interactive (loop) 629
- monitor option 468
- monitor process command 692
- myprimaryserver option 468
- myreplicationserver option 469

N

- Named Pipe communication method
 - options 297
- namedpipename option 471
- NAS
 - assigning management class to file systems 431
 - backing up file systems 159
 - deleting file spaces 232, 677
 - query node command 714
 - restore file systems 235, 751
 - restore NAS command 751
 - specifying full or differential backup 463
- NAS file servers using CIFS
 - backup and restore 163
- NAS file systems backup
 - command line 161
 - web client
 - GUI 160
- nasnodename option 471
- Net Appliance
 - backing up CIFS share definitions 182
- netapp file server 85
- Network Attached Storage (NAS)
 - backup file systems 159
- Network Data Management Protocol (NDMP) 7
- Network File System (NFS)
 - backup file systems 182
- network shares
 - backing up 109
 - making shares visible to the client 109
 - shares
 - making shares visible to client 109
- network-attached storage (NAS)
 - display nodes for which admin ID has authority 714
- network-attached storage (NAS)
 - backup file systems 648
 - cancel backup and restore processes 670, 692
 - deleting file spaces 232, 677
 - display file spaces on server 707
 - excluding files from backup 400

- network-attached storage (NAS) *(continued)*
 - monitoring backup or restore operations 468
 - querying file system images belonging to 700
 - restore file systems 235, 751
 - specifying for query 567
 - specifying node name for operations 471
 - specifying whether to save table of contents for each file system backup 563
- networked file systems
 - include-exclude statements 95
 - networked file systems 95
- new for backup-archive client V7.1.6 xvii
- no query restore 198
- node
 - specifying type to query 567
- node name 29
- node name, setting 29
- nodename option 232, 472
- nojournal option 474
- non-standard error conditions 219
- noprompt option 474
- nrtablepath option 475
- NTFS
 - Restoring volume mount points 735
- NTFS file spaces
 - back up 188
- NTFS/ReFS
 - backing up volume mount points 688
- NTFS/ReFS mounted volumes
 - Backing up data on 689
- numberformat
 - specifying 476
- numberformat option 476

O

- offline image backup 153
- online help
 - displaying 119
 - online forum 119
 - service and technical support 119
- online image backup 153
 - specifying gap size of striped volumes 427
- open file support 774
 - for backup operations 123
 - include files for 431
 - installing and configuring 83
 - overview 123
 - snapshot 240
- open registration
 - permissions 91
 - using 91
- operating system re-installation
 - Windows 202
- operating system requirements
 - clients 5
- optfile option 478
- options 546
 - absolute 322
 - adlocation 323
 - archive, summary 297
 - archmc 323
 - asnodename 324
 - asrmode 327, 620
 - auditlogging 328
 - auditlogname 330
 - authorization options 311

- options *(continued)*
 - autodeploy 332
 - autofsrename 333
 - backmc 335
 - backup
 - excluding system state 400
 - backup, summary 297
 - backupsetname 336
 - basesnapshotname 337
 - cadlistenonport 338
 - casesensitiveaware 339
 - central scheduling, summary 309
 - changingretries 340
 - class 341
 - clientview 342
 - clusterdiskonly 342
 - clusternode 345
 - collocatebyfilespec 346
 - command processing, summary 311
 - commmethod 347
 - commrestartduration 348
 - commrestartinterval 349
 - communication, summary 296
 - compressalways 350
 - compression 351
 - console 352
 - createnewbase 354
 - datacenter 356
 - datastore 356
 - dateformat 356
 - dedupcachepath 359
 - dedupcachesize 359
 - deduplication 360
 - deletefiles 361
 - description 362
 - detail 363
 - diagnostics 313
 - diffsnapshot 364
 - diffsnapshotname 366
 - dirmc 367
 - dironly 368
 - disablenqr 368
 - diskbuffsize 369
 - diskcachelocation 370
 - domain 371
 - domain.image 374, 378
 - domain.nas 375
 - domain.vcd 376
 - domain.vmfull 381
 - enable8dot3namesupport 386
 - enablearchiveretentionprotection 388
 - enablededupcache 389
 - enableinstrumentation 390
 - enablelanfree 392
 - encryptiontype 393
 - encryptkey
 - encryptkey=generate 394
 - encryptkey=prompt 394
 - encryptkey=save 394
 - errorlogmax 396
 - errorlogname 397
 - errorlogretention 398
 - exclude
 - exclude.archive 93, 400
 - exclude.backup 93, 400
 - exclude.compression 93, 400
 - exclude.dir 93, 400

options (*continued*)

exclude (*continued*)

- exclude.encrypt 400
- exclude.file 93, 400
- exclude.file.backup 93, 400
- exclude.fs.nas 400
- exclude.image 93, 400
- exclude.subfile 400
- exclude.systemobject 93
- wildcard characters 98, 100

exclude.dedup 400

EXCLUDE.VMDISK 405

fbbranch 406

fbclient 407

fbpolicyname 409

fbreposlocation 410

fbserver 411

fbvolumename 413

filelist 414

filename 417

filesonly 418

forcefailover 419

format and language, summary 310

fromdate 420

fromnode 420

fromtime 421

general rules when entering with commands 314

groupname 422

guitreeviewafterbackup 422

host 423

httpport 423

ieobjtype 425

ifnewer 426

imagegapsize 427

imagetofile 428

inactive 428

inclexcl 429

include

- wildcard characters 98, 100

include.archive 431

include.backup 431

include.compression 431

include.encrypt 431

include.file 431

include.fs 431

include.fs.nas 431

include.image 431

include.subfile 431

include.systemstate 431

include.vm 438

include.vmsnapshotattempts 442

include.vmtsmvss 444

incrbydate 445

incremental 446

incrthreshold 447

instrlogmax 448

instrlogname 449

journalpipe 450

lanfreecommmethod 451

lanfreeshmport 297, 452

lanfreessl 453

lanfreetcpport 453

lanfreetcpserveraddress 454

language 455

latest 456

localbackupset 457

managedservices 457

options (*continued*)

- maxcmdretries 459
- mbobjrefreshthresh 460
- mbpctrefreshthresh 461
- memoryefficientbackup 462
- mode 463
- monitor 468
- myprimaryserver 468
- myreplicationserver 469
- namedpipename 471
- nasnodename 471
- nodename 472
- nojournal 474
- noprompt 474
- nrtablepath 475
- numberformat 476
- optfile 478
- order of processing (precedence) 314
- password 478
- passwordaccess 480
- pick 481
- pitdate 482
- pittime 483
- postnschedulecmd 484
- postschedulecmd 484
- postsnapshotcmd 486
- prenschedulecmd 487
- preschedulecmd 487
- preservelastaccessdate 489
- preservepath 490
- presnapshotcmd 492
- querschedperiod 494
- querysummary 495
- quiet 496
- quotesareliteral 497
- replace 499
- replserverguid 500
- replservername 501
- replsslport 503
- repltcpport 504
- repltcpserveraddress 505
- resetarchiveattribute 507
- resourceutilization 508
- restore and retrieve, summary 306
- retryperiod 511
- revokeremoteaccess 512
- runasservice 512
- schedcmddisabled 513, 514
- schedlogmax 515
- schedlogname 517
- schedlogretention 518
- schedmode 519
- schedrestretrdisabed 521
- scrolllines 521
- scrollprompt 522
- sessioninitiation 524
- shmport 525
- showmembers 526
- skipmissingsyswfiles 527
- skipntpermissions 528
- skipntsecuritycrc 528
- snapdiff 83, 529
- snapdiffhttps 534
- snapshotproviderfs 535
- snapshotproviderimage 536
- snapshotroot 537
- specifying in commands 314

options (continued)

- srvoptsetencryptiondisabled 540
- srvprepostscheddisabled 540
- srvprepostsnapdisabled 541
- ssl 542
- sslrequired 544
- stagingdirectory 547
- subdir 548
- subfilebackup 549
- subfilecachepath 550
- subfilecachesize 551
- system state
 - exclude from backup processing 400
- systemstatebackupmethod 552
- tapeprompt 553
- tcpadminport 554
- tcpbuffsize 555
- tcpcadaddress 556
- tcpclientaddress 557
- tcpclientport 558
- tcpnodelay 558
- tcpport 559
- tcpserveraddress 560
- tcpwindowsize 561
- timeformat 562
- toc 563
- todate 565
- totime 565
- transaction processing, summary 312
- txnbytelimit 566
- type 567
- usedirectory 568
- useexistingbase 569
- usereplicationfailover 570
- v2archive 570
- vappmc 571
- vcdhost 572
- vcdpw 573
- vcduser 574
- verbose 574
- verifyimage 575
- virtual machine exclude options 404
- virtual machine include options 437
- virtualfsname 576
- virtualnodename 576
- vmautostartvm 577
- vmbackdir 578
- vmbacknodelete 579
- vmbackupmailboxhistory 580
- vmbackuptype 581
- vmbackupupdateguid 666
- vmbackvcbtransport 582
- vmchost 583
- vmcpw 584
- vmcuser 586
- vmdatastorethreshold 587
- vmdefaultvpportgroup 588
- vmdefaultdvswitch 589
- vmdefaultnetwork 590
- vmdiskprovision 591
- vmenabletemplatebackups 592
- vmexpireprotect 593
- vmiscsiadapter 594
- vmiscsiserveraddress 595
- vmlimitperdatastore 596
- vmlimitperhost 597
- vmlist 598

options (continued)

- vmmaxparallel 598
- vmmaxvirtualdisks 600
- vmmc 601
- vmmountage 601
- vmnoprdmdisks 602
- vmnovrdmdisks 603
- vmpreferdagpassive 604
- vmprocessvmwithprdm 606
- vmprocesswithindependent 605
- vmrestoretype 607
- vmskipctlcompression 609
- vmskipmaxvirtualdisks 610
- vmskipmaxvmdks 611
- vmstoragetype 611
- vmtagdatamover 612
- vmtempdatastore 615
- vmtimeout 621
- vmverifyifaction 615
- vmverifyiflatest 617
- vmvstortransport 618
- vssaltstagingdir 620
- web client, summary 313
- webports 622
- wildcardsareliteral 623

options file

- ASR recovery
- Windows 151

owner security information (SID)

- back up 188

P

- parallel backups 174, 596, 597, 598

parameters

- yes and no, alternatives 321

partial incremental

- incremental by date, running 127

password

- changing 116, 781
- number of characters 116
- setting 478
- setting for client 91
- specifying whether to generate automatically or set as user prompt 480
- using 111
- valid characters 116

password option 478

passwordaccess option 480

performance

- improving speed of backups, restores, archives, retrieves 297
- restore operations 197
- transaction options 312
- transaction processing 566

performing traditional full incremental backup 687

Persistent Storage Manager 164

- back up 164

pick option 481

pitdate 482

pittime option 483

point-in-time restore

- image backup 647

policies, storage management 267

policy domains

- default policy domain 267
- standard policy domain 267

- policy sets
 - active policy set 267
- portable media
 - restoring backup sets 206
- postschedulecmd option 484
- postsnapshotcmd option 486
- preferences editor
 - excluding domains from backup 126
- preschedulecmd option 487
- preschedulecmd option 487
- preservelastaccessdate option 489
- preservepath option 490
- Presnapshotcmd option 492
- preview
 - include-exclude list 101
- preview archive command 693
- preview backup command 694
- primary group SID
 - back up 188
- processing options
 - authorization 311
 - backup and archive 297
 - central scheduling 309
 - communication 296
 - diagnostics 313
 - error processing 312
 - format and language 310
 - overview 295
 - restore and retrieve 306
 - specifying in commands 314
 - transaction processing 312
 - using 295
 - web client 313
- processing time 125
- Protecting cluster disks 71
- proxied session restrictions 145, 146, 242
- publications xiii

Q

- query
 - amount of information that displays on screen 521
 - backups, establish point-in-time 482, 483
 - based on date and time of backup, archive 420, 421
 - description for 362
 - display active and inactive objects 428
 - files for another node 420
 - group
 - command 709
 - display members of 526
 - include-exclude list 712
 - NAS or client objects 341
 - nodes to which client has proxy authority 242
 - nodes to which client has proxy node authority 145
 - process directories only (not files) 368
 - scrolling preferences after displaying information on screen 522
 - system information 718
 - system state 720
- query access command 695
- query adobjects command 696
- query archive command 697
- query backup command 700
- query backupset command 703, 705
- query filespace command 707
- query group command 709
- query image command 710

- query inclexcl command 712
- query mgmtclass command 713
- query node command 714
- query options command 715
- query restore command 716
- query schedule
 - enhanced 254
- query schedule command 717
- query schedule command, enhanced 717
- query session command 717
- query systeminfo command 718
 - encryption processing 718
- query systemstate command 720
- query vapp command 721
- query VM command 725
- queryschedperiod option 494
- querysummary option 495
- quiesce applications 444
- quiet option 496
- quotesareliteral option 497

R

- raw logical volume
 - image backup of 153
 - restoring 203
- reanimate
 - tombstone objects 224
- rebinding files to a different management class 276
- ReFS
 - backing up volume mount points 688
 - Restoring volume mount points 735
- ReFS file spaces
 - back up 188
- registering
 - client with server 91
 - using closed registration 91
 - using open registration 91
- remote network connection
 - establishing 116
- remotely accessed files
 - excluding 97
 - UNC names 97
- removable media
 - back up 187
- replace option 499
- replserverguid option 500
- replservername option 501
- replsslport option 503
- repltcpport option 504
- repltcpserveraddress option 505
- resetarchiveattribute option 507
- resourceutilization option 508
- restart restore command 729
- restartable restore 198
- restartable restore sessions, display 716
- restore 200, 748
 - active and inactive file versions 736
 - active directory objects 224, 225, 226, 228
 - active version 194
 - ASR (Automated System Recovery) files 201
 - ASR recovery mode 327, 620
 - authorizing another user 229
 - backup set
 - supported tape devices 739, 744
 - backup sets
 - overview 206

- restore (*continued*)
 - backups, establish point-in-time 482, 483
 - based on date and time of backup 420, 421
 - classic (also known as standard) 198
 - create list of backup versions to 481
 - directories 195
 - display active and inactive objects 428
 - during failover 228
 - estimating processing time 125
 - files 195
 - files and directories 195
 - files belonging to another node 231
 - files for another node 420
 - from file spaces that are not Unicode-enabled 737
 - from portable media
 - overview 206
 - from system state backup 224
 - group
 - command 746
 - GUI, displaying active and inactive versions 118
 - image 203
 - considerations 748
 - enable detection of bad sectors on target volume 575
 - to a file 428
 - using chkdsk tool to repair 203
 - using fsck tool to repair 203
 - image, suppress confirmation prompt 474
 - improving speed using share memory 297
 - inactive version 194
 - large number of files 197
 - list of files 414
 - local backup set using the GUI 206
 - logical volume 203
 - Microsoft Dfs junction 736
 - Microsoft Dfs trees and files 203
 - modify CAD and agent services 228
 - most recent backup version 456
 - NAS file systems 235
 - command line 237
 - web client 235
 - NAS files and directories using web client 236
 - Net Appliance CIFS shares 211
 - no query 198
 - overview 193
 - primary tasks 193
 - process directories only (not files) 368
 - raw logical volume 203
 - replace existing file with latest backup 426
 - restartable 198
 - restrictions and limitations 226
 - sorting file list 118
 - sparse files 737
 - sparse files to a non-NTFS or non-ReFS file system 737
 - sparse files, size restriction for 737
 - standard (also known as classic) 198
 - starting a web client session 115
 - summary of options 306
 - system state 753
 - to different workstation 576
 - using commands 195
 - using fsck tool to repair 748
 - using GUI and command line 225
 - using the GUI 195
 - using universal naming convention names 194
 - VMware Consolidated Backup 212
 - whether to prompt before overwriting existing files 499
 - workstation, to another 232
 - restore adobjects command 738
 - restore backupset command 739, 744
 - restore backupset command considerations 208, 742
 - restore command 730
 - using multiple 197
 - restore full VM backup
 - VCB backups 222
 - restore group command 746
 - restore image command 748
 - restore maximum file size 135
 - restore NAS command 751
 - restore procedure
 - ASR 202
 - Windows 202
 - restore systemstate command 753
 - restore vapp command 755
 - restore vm command 758
 - Restoring data on mounted volumes
 - NTFS 735
 - Restoring data on NTFS mounted volumes 735
 - Restoring NTFS or ReFS volume mount points 735
 - restoring point-in-time 233
 - using command line 233
 - using GUI 233
 - restoring your system
 - ASR recovery mode
 - Windows 202
 - restrictions
 - asnodename option 324
 - asnodename session settings 326
 - runasservice and encryptkey 512
 - runasservice and passwordaccess 512
 - runasservice and replace 512
 - within a proxied session 145, 146, 242
 - retain extra versions attribute 270
 - retain only versions attribute 271
 - retain versions attribute 272
 - retention grace period
 - archive 268, 276
 - backup 268, 276
 - retrieve
 - archive copies 245
 - archive files by name 245
 - authorizing another user 229
 - based on date and time of archive 420, 421
 - description for 362
 - during failover 228
 - files belonging to another node 231
 - files for another node 420
 - improving speed using share memory 297
 - list of files 414
 - primary tasks 239
 - process directories only (not files) 368
 - replace existing file with latest archive if existing file is newer 426
 - sorting file list 118
 - starting a web client session 115
 - summary of options 306
 - to different workstation 576
 - whether to prompt before overwriting existing files 499
 - workstation, to another 232
 - retrieve command 766
 - retrieve maximum file size 135
 - retrieving
 - archives using command line 245
 - retryperiod option 511
 - return codes for operations 265

- revokeremoteaccess option 512
- runasservice option 512
- running a snapshot difference backup
 - with HTTPS 143
- running a snapshot differential backup
 - with HTTPS 143

S

SAN

- restoring backup sets using 743

scenarios

- instant access, from the command line 214
- instant restore, from command line 214

- schedcmddisabled option 513, 514

- schedlogmax option 515

- schedlogname option 517

- schedlogretention option 518

- schedmode option 519

- schedrestretrdisabled option 521

- schedule command 770

schedule log

- controlling the size 515
- specifying number of days to keep entries and whether to
 - save pruned entries 518
- specifying path and file name to store schedule log
 - information 517

scheduled (automated) backups

- closing files before backup 185
- displaying scheduled work 254, 256
- options for 260
- process commands after backup 484
- process commands before backup 487
- restart applications after backup 185
- starting 39

scheduled commands

- enabling-disabling 260

- scheduled events, displaying 717

scheduled services

- disabling scheduled commands 513, 514
- restrictions for NAS file systems 159

scheduler

- configuring 36
- displaying scheduled work 254, 256
- event logging 256
- number of hours between contacts to server for scheduled
 - work 494
- number of minutes between attempts to process scheduled
 - commands 511
- options for 260
- polling mode or prompted mode 519
- starting 39
- whether server or client initiates sessions through
 - firewall 524
- whether to disable execution of restore or retrieve
 - operations 521

scheduler comparison

- CAD versus traditional scheduler 36

scheduler service

- installing on Microsoft Cluster Server cluster nodes 70
- installing on Veritas Cluster Server cluster nodes 70

- Scheduler Service Configuration Utility 37

- scheduler wizard 37

scheduling

- client node proxy backup 145, 147

- scrolllines option 521

- scrollprompt option 522

Secure Sockets Layer (SSL)

- establishing communications with 42

security information

- bypass processing of 528

- selective backup 537, 771

- client command line 127

- client Java GUI 125

- command line 127

- overview 127, 144

- selective command 771

- self-contained application protection 444

serialization

- copy serialization

- dynamic 271

- shared static 271

- static 271

server

- communicating with 29

- establishing communications through firewall 39

- establishing communications with 29

- establishing communications with Secure Sockets Layer
 - (SSL) 42

- query Active Directory for communication method and

- server with which to connect 568

- TCP/IP address of Tivoli Storage Manager server 560

- TCP/IP port address for 559

server options

- Sslfipsmode 543

- service and technical support 119

- service recovery settings 252

- session information, displaying 717

- sessioninitiation option 524

- set access command 775

- restore-retrieve authorization 229

- set event command 778

- set netappsvm 85

- set password command 781

- set vmtags command 787

setting

- environment variables

- DSM_CONFIG 32

- DSM_DIR 32

- DSM_LOG 32

- user privileges 115

- shared dynamic serialization 271, 340

- shared memory communication method

- options 297

- shared static serialization 271, 340

- shmport option 525

- showmembers option 526

- silent installation 19

- skipmissingsyswfiles option 527

- skipntpermissions option 528

- skipntsecuritycrc option 528

- snapdiff option 83, 529

- snapdiffhttps option 534

snapshot

- open file support 240

- snapshot difference 83, 529

- with HTTPS 142

- snapshot differential backup

- with HTTPS 142

- snapshot differential backup with HTTPS connection 534

- snapshot-differential-incremental backup 529

- snapshotproviderfs option 535

- snapshotproviderimage option 536

- snapshotroot option 537

- snapshotroot option with incremental and selective commands 149
- Software updates 26
- sparse files
 - restore size restriction 737
 - restoring 737
 - restoring to a non-NTFS or non-ReFS file system 737
- specifying whether to update last access date 489
- srvoptsetencryptiondisabled option 540
- srvprepostscheddisabled option 540
- srvprepostsnapdisabled option 541
- SSL (Secure Socket Layer)
 - establishing communications with 42
- ssl option 542
- Sslfipsmode option 543
- sslrequired option 544
- stagingdirectory option 547
- standard (classic) restore 198
- standard management class
 - copy destination 272
 - copy frequency 270
 - copy group name 269
 - copy mode
 - absolute 272
 - modified 272
 - copy serialization 271
 - copy type 270
 - deduplicate data attribute 273
 - default values 269
 - retain extra versions 270
 - retain only version 271
 - retain versions 272
 - versions data deleted
 - active versions 270
 - inactive versions 270
 - versions data exists 270
- standard policy domain 267
- start the client scheduler at startup 252
- starting
 - automatically 116
 - overview 1
- starting a session
 - batch mode 112
 - interactive mode 112
- static serialization 271
- storage
 - displaying restartable restore sessions 716
- Storage Agent
 - for LAN-free data movement 130
 - using for LAN-free data movement 392
- storage area network
 - for LAN-free data movement 130
 - restoring backup sets using 392, 743
 - using for LAN-free data movement 392
- storage management policies 267
 - assigning management classes to files 186
 - copy groups 268
 - default management class 267
 - display on backup-archive client or web client GUI 186
 - include-exclude list 268
 - management classes 268
 - policy domains
 - default 267
 - standard 267
 - policy sets
 - active policy set 267
- subdir option 548

- subdirectories
 - archive 242
 - include in backup 127
- subfilebackup option 549
- subfilecachepath option 550
- subfilecachesize option 551
- support
 - gathering system information for 352, 417, 718
- swing-enabled browser
 - necessary to run web client 115
- syntax diagram
 - reading xiv
 - repeating values xiv
 - required choices xiv
- system access control list (auditing information)
 - back up 188
- system files
 - excluding 96
- system information
 - gathering 352, 417
- system recovery
 - Windows 202
- system restore 201
- system state
 - assigning management class 95, 149, 431
 - back up 149, 650
 - display active and inactive objects 428
 - exclude from backup processing 95, 400
 - query 720
 - restore 753
 - restore from backup set 739, 744
 - restoring 200
- systemstatebackupmethod option 552

T

- tagging support in VMware 612
- tapeprompt option 553
- tasks
 - closed registration 91
 - open registration 91
- TCP/IP communication method
 - options 296
- tcpadminport option 554
- tcpbuffsize option 555
- tcpcadaddress option 556
- tcpclientaddress option 557
- tcpclientport option 558
- tcpnodelay option 558
- tcpserveraddress option 560
- tcpwindowssize option 561
- time format
 - specifying 562
- timeformat option 562
- Tivoli Storage Manager
 - client components
 - Windows client 6
 - communication methods
 - Windows client 6
 - environment prerequisites 6
 - FAQs 79
 - hardware, disk space, memory requirements
 - Windows client 6
 - installation requirements 6
 - installing on Microsoft Cluster Server cluster nodes 70, 79
 - installing on Veritas Cluster Server cluster nodes 70
 - online forum 119

- Tivoli Storage Manager *(continued)*
 - password 111
 - upgrading from earlier versions of the product 1
- Tivoli Storage Manager client
 - authentication 108
- Tivoli Storage Manager FastBack configuration 68
- Tivoli Storage Manager FastBack configuration wizard 9, 69
- Tivoli Storage Manager FastBack data backup 181
- Tivoli Storage Manager FastBack data restore 181
- Tivoli Storage Manager FastBack installation requirements 8
- Tivoli Storage Manager password
 - using 111
- toc option 563
- todate option 565
- tombstone objects
 - preserving attributes 227
 - reanimate 224
- totime option 565
- traditional full incremental backup 138
- transaction processing
 - summary of options 312
 - txnbytelimit option 566
- troubleshooting
 - troubleshooting Windows client installations 25
 - Windows client installations 25
- txnbytelimit option 566
- type option 567

U

- UAC 109
- UNC
 - back up shared files and directories using 189
 - set domain list using 188
- UNC names
 - excluding files 97
 - remotely accessed files 97
 - restore files 194
- Unicode
 - pre-backup considerations 130, 131
 - renaming file spaces that are not Unicode to
 - Unicode-enabled 333, 683, 771
 - restore from file spaces that are not Unicode-enabled 737
- universal naming convention
 - restore 194
 - using to specify domain list 188
- updates_622_client 650, 753
- updating the client automatically 3
- upgrading backup-archive clients 1
- upgrading the backup-archive client from earlier versions of the product 1
- usedirectory option 568
- useexistingbase option 569
- user account control 109
 - effects on network shares 109
- user privileges
 - setting 115
- usereplicationfailover option 570
- using multiple sessions 182

V

- v2archive option 570
- vappmc option 571
- vcdhost option 572
- vcdpw option 573

- vcduser option 574
- verbose option 574
- verifyimage option 575
- Veritas Cluster Server cluster nodes
 - FAQs 79
 - installing scheduler service 70
 - installing Tivoli Storage Manager 70, 79
- versions data
 - deleted attribute 270
 - exists attribute 270
- virtual machine
 - exclude options 404
 - include options 437
- virtualfsname option 576
- virtualnodename option 576
- VM 164
- vmautostartvm 577
- vmbackdir option 578
- vmbacknodelete option 579
- vmbackupmailboxhistory 580
- vmbackuptype option 581, 601
- vmbackupupdateguid 666
- vmbackvcbtransport 582
- vmchost option 583
- vmcpw option 584
- vmctlmc option
 - options
 - vmctlmc 585
- vmcuser option 586
- vmdatastorethreshold
 - option 587
- vmdefaultdvportgroup option 588
- vmdefaultdvswitch option 589
- vmdefaultnetwork option 590
- vmdiskprovision 591
- vmenabletemplatebackups option 592
- vmexpireprotect option 593
- vmiscsiadapter 594
- vmiscsiserveraddress option 595
- vmlimitperdatastore option 596
- vmlimitperhost option 597
- vmlioption 598
- vmmaxparallel option 598
- vmmaxvirtualdisks 600
- vmmountage option 601
- vmnoprmdisks 602
- vmnovrmdisks 603
- vmpreferdagpassive option 604
- vmprocessvmwithprdm 606
- vmprocesswithindependent 605
- vmrestoretype option 607
- mskipctlcompression option 609
- mskipmaxvirtualdisks 610
- mskipmaxvmdks 611
- vmstoragetype option 611
- vmtagdatamover
 - option 612
- vmtempdatastore option 615
- vmtimeout option 621
- vmverifyifaction 615
- vmverifyiflatest 617
- vmvstortransport option 618
- VMware backup 173
- VMware Consolidated Backup
 - restoring data 212
- VMware virtual machine backups 170
- types 164

- VMware virtual machines backups
 - file-level backups, preparing for 171
- Volume Shadowcopy Service (VSS)
 - configuring for online image backup 83
 - configuring for open file support 83
- VSS (see Volume Shadowcopy Service) 83
- vssaltstagingdir option 620
- vStorage backup server
 - off-host backup 167

- WinPE CD
- Windows 202

W

- web client
 - configuration overview 33
 - configuring 34
 - configuring in cluster environment 70
 - enable to run in a swing-enabled browser 115
 - establishing communications through firewall 39, 423
 - GUI 160
 - NAS
 - file systems backup 160
 - restrict administrator from accessing client running web client 512
 - restrictions for NAS file systems 159
 - specifying TCP/IP port address for 423
 - starting 115
 - summary of options 313
 - supported browsers 115
 - unsupported functions 121
 - using through a firewall 622
- web client configuration overview 33
- webports option 622
- whether to compute CRC for comparison of Windows security information 528
- wildcard characters
 - guidelines 632
 - include or exclude files 97
 - include or exclude groups of files 98
 - specifying a drive specification in dsm.opt 98
 - to include or exclude groups of files 100
- wildcardsareliteral option 623
- Windows archive attribute
 - reset after backup 507
- Windows client
 - client components 6
 - communication methods 6
 - disk space requirements 6
 - forced reboot 10
 - hardware requirements 6
 - initial installation 11
 - installation prerequisites 10
 - installation types 11, 15, 19
 - initial installation 10
 - modify an installed client 10, 23
 - reinstallation 10, 23
 - silent installation 10
 - uninstalling 10, 23
 - upgrade installation 10
 - installing 10
 - memory requirements 6
 - reinstallation 19
 - upgrade installation 15
- Windows components
 - installable 6
- Windows security information
 - whether to compute CRC for comparison of 528
- Windows supported file systems 7



Product Number: 5608-E01
5608-E02
5608-E03

Printed in USA